Holographic Algorithms

Markus Bläser
Universität des Saarlandes

Draft—October 6, 2016 and forever

1 The FKT algorithm

The Fisher–Kasteleyn–Tempereley algorithm is a fascinating algorithm, which you rarely find in algorithm books. It was invented independently by Temperley and Fisher [TF61] as well as Kasteleyn [Kas61] in the context of statistical physics.

1.1 Matchings

A graph G is a tuple (V, E) of sets. The node set V will always be finite in this article. We call G undirected, if the edge set $E \subseteq \binom{V}{2}$ consists of unordered pairs of nodes. G is called directed, if $E \subseteq V \times V$ consists of ordered pairs of nodes. In an undirected graph, the degree $\deg(v)$ of a node is the number of edges that are incident with v, that is $\#\{e \in E \mid v \cap e \neq \emptyset\}$. In directeded graphs, the indegree of a node v is the number of edges pointing to v, that is, $\#\{u \mid (u, v) \in E\}$. In the same way, the outdegree of a node is the number of edges leaving v. In multigraphs, the edges are counted with multiplicities. For further informations, the reader is referred to any of the many books on graph theory, e.g. [Die05].

Definition 1.1 *Let $G = (V, E)$ be a graph, $M \subseteq E$.*

1. *M is a* matching *of G if the degree of every node in (V, M) is at most one.*
2. *M is* maximal *if for every $e \in E \setminus M$, $M \cup \{e\}$ is not a matching in G.*
3. *M is* maximum *of for every matching M', $\#M' \leq \#M$.*
4. *M is perfect, if every node in (V, M) has exactly degree 1.*
5. *$\mathcal{M}(G)$ ($\mathcal{PM}(G)$) denotes the set of all (perfect) matchings of G.*

A matching is maximal, if it cannot be extended by adding a new edge. A matching is maximum, if it has maximum cardinality among all matchings. Figure 1 shows a matching that is maximal but not maximum. Every perfect matching is a maximum matching. If a graph $G = (V, E)$ has a perfect matching, then obviously $\#V$ is even.

Given a graph G, we would like to answer the following questions:

1. Has G a perfect matching?

Fig. 1. The matching (thick edges) on the left-hand side is maximal but not maximum. The right-hand side shows a maximum (and perfect) matching in the same graph.

2. Compute a maximum matching!

3. Count the number of perfect matchings!

The first two tasks are easy, there are a polynomial time algorithms for it, for instance Edmond's blossom algorithm [Edm65]. The third task turns out to be hard, it is hard for the class #P, as shown by Valiant [Val79].

1.2 A brief introduction to counting complexity

Instead of deciding languages $L \subseteq \{0,1\}^*$, counting complexity investigates the complexity of computing functions $\{0,1\}^* \to \mathbb{N}$.

Definition 1.2 *#P is the class of all functions $f : \{0,1\}^* \to \mathbb{N}$ such that there is a polynomial time bounded nondeterministic Turing machine M, such that $f(x)$ is the number of accepting paths of M.*

Natural functions in #P are the functions #SAT or #3-SAT that map a given a formula in CNF or 3-CNF, respectively, to its number of satisfying assignments. A corresponding Turing machine just nondeterministically guesses an assignment, evaluates the formula at the assignment and accepts if this assignment is satisfying. Every satisfying assignment corresponds to exactly one accepting path.

If we want to compare the computational complexity of functions $\{0,1\}^* \to \mathbb{N}$, reductions are helpful.

Definition 1.3 *Let $f, g : \{0,1\}^* \to \mathbb{N}$ be two functions. Let $s : \{0,1\}^* \to \{0,1\}^*$ be polynomial time computable and $t : \mathbb{N} \to \mathbb{N}$ be polynomial time computable (when the numbers are encoded in binary).*

1. *s is called a* parsimonious reduction *from f to g if $f(x) = g(s(x))$ for all $x \in \{0,1\}^*$.*

2. *f is called parsimoniously reducible to g if there exists a parsimonious reduction from f to g. In this case, we write $f \leq_{par} g$.*

3. *(s,t) is called a* counting reduction *from f to g, if $f(x) = t(g(s(x)))$ for all $x \in \{0,1\}^*$.*

4. *f is called counting reducible to g if there exists a counting reduction from f to g. In this case, we write $f \leq_c g$.*

Parsimonious and counting reduction fulfill the desired properties of reductions: First it is easy to check that these notions of reducibility are transitive. Second, if f is reducible to g and g is easy, that is, $g \in \mathsf{FP}$, where FP denotes the class of all functions $\{0,1\}^* \to \mathbb{N}$ that are deterministically polynomial time computable, then $f \in \mathsf{FP}$, too. A function f is called $\#\mathsf{P}$-hard if every function in $\#\mathsf{P}$ is reducible to f.

It is quite easy to show that #SAT and #3-SAT are $\#\mathsf{P}$ hard under parsimonious reductions. Just a little care is needed to ensure that every accepting path corresponds to exactly one satisfying assignment.

For a graph G, let

$$\mathrm{PerfMatch}(G) = \#\mathcal{PM}(G)$$

denote the number of perfect matchings of G. We can interpret PerfMatch as a function $\{0,1\}^* \to \mathbb{N}$ by encoding graphs appropriately (and return the value 0 whenever the input is not a valid encoding of a graph).

Theorem 1.4 (Valiant [Val79]) PerfMatch *is $\#\mathsf{P}$-hard under counting reductions.*

Valiant actually stated his theorem only for another notion of reductions, so-called Turing reductions. Counting reductions were introduced by Zankó [Zan91].

Note that since we can decide in polynomial time whether a graph has a perfect matching, PerfMatch cannot be $\#\mathsf{P}$-hard under parsimonious reductions, unless $\mathsf{P} = \mathsf{NP}$.

The problem of counting perfect matchings even stays hard if the graph is bipartite (also shown by Valiant). Recall that a graph is bipartite, if its node set can be partioning into two disjoint sets U and V such that every edge has one node in U and the other in V. If a bipartite graph has a perfect matching, then necessarily $\#U = \#V$. Moreover, every matching M defines a bijection $U \to V$ in the natural way. If $\{u,v\}$ is an edge in M with $u \in U$ and $v \in V$, then this bijection maps u to v. If we further identify U and V with $\{1,\ldots,n\}$, then every matching M defines a permutation in S_n.

For an $n \times n$-matrix $A = (a_{i,j})$ let the *permanent* of A be defined by

$$\operatorname{per} A = \sum_{\pi \in S_n} a_{1,\pi(1)} \cdots a_{n,\pi(n)}$$

With a bipartite graph G, we can associate a (bipartite) adjacency matrix A, the rows are identified with nodes from U and the columns with nodes from V and $a_{u,v} = 1$ if $\{u, v\}$ is an edge of G. Otherwise the entry $a_{u,v} = 0$. If we set $\operatorname{per} G := \operatorname{per} A$, then by the correspondance between matchings and permutations,

$$\operatorname{per} G = \operatorname{PerfMatch} G.$$

In this way, Valiants result can also be reformulated as the permanent being #P-hard.

1.3 Determinants and Pfaffians

While the permanent is hard, there is a similar polynomial that is easy to compute, the determinant

$$\det A = \sum_{\pi \in S_n} \operatorname{sgn}(\pi) a_{1,\pi(1)} \cdots a_{n,\pi(n)}.$$

There are efficient algorithms for computing the determinant, for instance, Gaussian elimination.

When one sees the similarity between permanent and determinant, the first idea is to somehow reduce the computation of the permanent to that of the determinant. Of course, this cannot be that easy. There can be cancellations between the different terms in the determinant. For instance, consider the complete bipartite graph $K_{n,n}$ with n nodes on each side. Its adjacency matrix has a 1 in each position, therefore its determinant is 0 (if $n > 1$), although $K_{n,n}$ has many matchings, namely $n!$. The next idea is to replace some of the 1s in the adjacency matrix by a -1 and try to cancel the signs. It turns out that this astonishingly works in the case of planar graphs!

Let us briefly recall some facts about permutations. A *transposition* is a permutation that leaves all but two elements $i \neq j$ fixed and exchanges i with j. The set of all transpositions on n elements generates S_n. With each permutation π, we can associate a directed graph. The nodes are the numbers $\{1, \ldots, n\}$ and the edges are $\{(i, \pi(i)) \mid 1 \leq i \leq n\}$. This graph consists of node-disjoint directed cycles counting self-loops as cycles. Such a graph is also called a *cycle cover*.

Fact 1.5 *The sign* $\operatorname{sgn}(\pi)$ *of a permutation* $\pi \in S_n$ *can be defined as:*

- $(-1)^k$ *where* $\pi = \tau_1 \circ \cdots \circ \tau_k$ *for transpositions* $\tau_1, \ldots, \tau_k$*, (This is well defined, that is, independent of the concrete decomposition of* π *into transpositions.)*

- $(-1)^{n-c}$ *where* c *is the number of cycles in* π*,*

- $(-1)^i$, *where* i *is the number of inversions of* π, *that is* $i = \#\{(i,j) \mid i < j \wedge \pi(i) > \pi(j)\}$.

A matrix A is *symmetric*, if $A = A^T$. It is *skew-symmetric*, if $A = -A^T$.

Definition 1.6 *Let* A *be a skew-symmetric* $2n \times 2n$*-matrix. The* Pfaffian *of* A *is*

$$\mathrm{pf}(A) = \frac{1}{2^n n!} \sum_{\sigma \in S_{2n}} \mathrm{sgn}(\sigma) \prod_{i=1}^{n} a_{\sigma(2i-1),\sigma(2i)}.$$

If A *has odd size, then* $\mathrm{pf}(A) = 0$.

Note that many permutations σ produce the same term $\mathrm{sgn}(\sigma) \prod_{i=1}^{n} a_{\sigma(2i-1),\sigma(2i)}$: With every such term, we can associate a partition of $\{1,\ldots,2n\}$ into unordered pairs, namely, the pairs $\{\sigma(2i-1),\sigma(2i)\}$, $1 \le i \le n$. Two permutations that have the same partition will produce the same term. If we swap two pairs, this will not change the term, since swapping a pair means composing σ with two transpositions, so the sign of the permutation will not change. If we swap the order of two elements within a pair, then the sign will changes. However $a_{\sigma(2i-1),\sigma(2i)}$ is replaced by $a_{\sigma(2i),\sigma(2i-1)} = -a_{\sigma(2i-1),\sigma(2i)}$, so this effect is cancelled. There are $n!$ ways to order the pairs and 2^n ways to order the elements within the pairs, so $2^n n!$ permutations generate the same term.

Let Π be the set of all partitions of $\{1,\ldots,2n\}$ into unordered pairs. We write the pairs of a partition $\alpha \in \Pi$ in a canonical way, namely as $(i_1,j_1),\ldots,(i_n,j_n)$ with $i_k < j_k$ for all k and $i_1 < i_2 < \cdots < i_n$. With α, we associate the permutation

$$\pi_\alpha = \begin{pmatrix} 1 & 2 & 3 & 4 & \ldots & 2n-1 & 2n \\ i_1 & j_1 & i_2 & j_2 & \ldots & i_n & j_n \end{pmatrix}.$$

We have

$$\mathrm{pf}(A) = \sum_{\alpha \in \Pi} \mathrm{sgn}(\pi_\alpha) \prod_{i=1}^{n} a_{\pi_\alpha(2i-1),\pi_\alpha(2i)}.$$

Note that we can we can think of α being a perfect matching in a graph with nodes $\{1,\ldots,2n\}$. Namely, the pairs of α are the edges of the matching. If M is the corresponding perfect matching, we will write π_M instead of π_α. If you take A to be the adjacency matrix of a graph G (made skew symmetric by replacing each 1 below the main diagonal by -1), then each perfect matching of G contributes a 1 or -1 in the Pfaffian of A.

It turns out that for the Pfaffian, there is an efficient algorithm by reduction to the determinant. We start with two helpful lemmas. Let K_n denote the complete graph with n nodes, i.e., $K_n = (\{1,\ldots,n\}, \binom{\{1,\ldots,n\}}{2})$.

Observation 1.7 *The union of two matchings is a(n undirected) cycle cover with only even cycles. Here we consider the union as a multigraph, that is, edges that are contained in both matchings form a cycle of length two in the union.*

Lemma 1.8 *There is a bijection between the permutations in S_{2n} with only even cycles and $\mathcal{PM}(K_{2n}) \times \mathcal{PM}(K_{2n})$.*

Proof. Let σ be the a permutation with only even cycles. Let C be a (directed) cycle in σ. We distribute the edges of C among the two matchings M and N: Let v be the smallest node in C. Let $e = (v, v')$ be the edge leaving v. We put e into M, the next edge on the cycle into N, the edge after that into M again and so on. Since every cycle has even length, this results in two matchings.

Let σ and σ' be the two different permutations with corresponding cycle covers C and C' and let (M, N) and (M', N') be their images under the mapping defined above. If an edge in C connects two nodes that are not connected in C' and vice versa, then their will be at least one edge that is contained in one of M or N but not in M' or N' and vice versa. Therefore $(M, N) \neq (M', N')$. Otherwise, the cycles in C and C' are the same, only their direction may differ. Since $\sigma \neq \sigma'$, the direction of at least one cycle differs. But the edges of this cycle are put into M are put into N' and those that are put into N are put into M'. Therefore $(M, N) \neq (M', N')$ and the mapping is injective.

Since the union of two matchings is a collection of even cycles, the mapping is surely surjective. ∎

An *orientation* of a graph G is a directed graph $\vec{G}$ that is obtained from G by giving each edge in G a direction. That is, for each edge $\{i, j\}$ in G we either add (i, j) or (j, i) to $\vec{G}$. Let M be a matching in G. We now extend the definition of π_M: Let $(i_1, j_1), \ldots, (i_n, j_n)$ be the edges in $\vec{G}$ of M with $i_1 < \cdots < i_n$. Then π_M is the permutation

$$\begin{pmatrix} 1 & 2 & 3 & 4 & \ldots & 2n-1 & 2n \\ i_1 & j_1 & i_2 & j_2 & \ldots & i_n & j_n \end{pmatrix}.$$

Note that before, the "direction" of the edges were given by the fact that we assumed that $i_k < j_k$. Now the direction if given by the orientation $\vec{G}$.

Lemma 1.9 *Let M and N be two perfect matchings of a graph G. Then σ be the unique permutation with only even cycles that is mapped to (M, N) under the mapping from Lemma 1.8. Then*

$$\operatorname{sgn}(\pi_M) \operatorname{sgn}(\pi_N) = (-1)^k$$

where k is the number of cycles in σ with an odd number of edges pointing in the same direction. (Note that since all cycles in σ are even, this is well defined.)

Proof. First of all we prove that if the lemma is true for some orientation, then it is true for any orientation. Consider some edge e.

- If $e \notin M \cup N$, then reversing e has neither an effect on σ nor M nor N.
- If $e \in M \cap N$, then e is contained in a cycle of length two of σ. The parity of π_M and π_N changes, but the number of cycles with an odd number of edges pointing in the same direction stays the same.
- Finally, we consider the case $e \in M \setminus N \cup N \setminus M$. W.l.o.g. we can assume that $e \in M \setminus N$. Flipping e changes the sign of π_M but leaves π_N unaffected. However, the parity of the cycles with an odd number of edges in σ changes, too.

By successively flipping an edge, we can convert one orientation into any other.

Therefore, we choose an orientation in which the edges in every cycle point in the same direction. In particular, every cycle has an even number of edges in the same direction. Let C be such a cycle. Let $i_1, \ldots, i_\ell$ be the nodes of C (in the order given by the orientation of C). One permutation, say π_M, restricted to C looks like

$$\begin{pmatrix} i_1 & i_2 & \ldots & i_{k-1} & i_k \\ i_1 & i_2 & \ldots & i_{k-1} & i_k \end{pmatrix}$$

and the other looks like

$$\begin{pmatrix} i_1 & i_2 & \ldots & i_{k-1} & i_k \\ i_2 & i_3 & \ldots & i_k & i_1 \end{pmatrix}.$$

The first one is the identity and consists of an even number of cycles of length 1, the other one of one cycle of length k. We can do the same for every other cycle of $M \cup N$. Therefore, $\operatorname{sgn}(\pi_M) = 1$ and $\operatorname{sgn}(\pi_N) = (-1)^k$. ∎

Theorem 1.10 *Let A be a skew-symmetric matrix. Then* $\operatorname{pf}(A)^2 = \det(A)$.

Proof. If A has odd size, then $\operatorname{pf}(A) = 0$ by definition. On the other hand, $\det(A) = \det(-A^T) = -\det(A^T) = \det(A)$. Therefore, $\det A = 0$.

So let A be of even size $2n \times 2n$. Consider a term of $\operatorname{pf}(A)^2$. It is of the form

$$\operatorname{sgn}(\pi_M) a_{i_1,j_1} \cdots a_{i_n,j_n} \cdot \operatorname{sgn}(\pi_{M'}) a_{i'_1,j'_1} \cdots a_{i'_n,j'_n},$$

where $\{i_1, j_1\}, \ldots, \{i_n, j_n\}$ are the edges of the perfect matching/partition M and $\{i'_1, j'_1\}, \ldots, \{i'_n, j'_n\}$ are the edges of the perfect matching/partition M'. $M \cup M'$ corresponds to a permutation σ with only even cycles. By Lemma 1.9, $\operatorname{sgn}(\pi_M) \cdot \operatorname{sgn}(\pi_{M'}) = \operatorname{sgn}(\sigma)$. Since there is a bijection between

pairs of matchings and permutations with only even cycles by Lemma 1.8, every permutation with only even cycles corresponds to exactly one term of $\mathrm{pf}(A)^2$.

Therefore, it remains to show that the overall contribution of the permutations with at least one odd cycle is zero. To this aim, we define an involution I (i.e., a self-inverse bijection) on the set with of all permutations with at least one odd cycle. We order the cycles of a permutation σ by the smallest node contained in the cycle. If the first odd cycle is a cycle of length one, then $I(\sigma) = \sigma$. Otherwise, $I(\sigma) = \sigma'$, where σ' is obtained from σ by reversing the edges in the first odd cycle. I is bijective and self-inverse by construction.

Let σ be a permutation with an odd cycle. If $I(\sigma) = \sigma$, then

$$\mathrm{sgn}(\sigma) a_{1,\sigma(1)} \cdots a_{n,\sigma(n)} = 0,$$

because $a_{i,i} = 0$ for all i. Otherwise,

$$\mathrm{sgn}(\sigma) a_{1,\sigma(1)} \cdots a_{n,\sigma(n)} + \mathrm{sgn}(I(\sigma)) a_{1,I(\sigma)(1)} \cdots a_{n,I(\sigma)(n)} = 0,$$

since the sign of σ and $I(\sigma)$ are the same, but for an odd number of factors in the product, we replaced $a_{i,j}$ by $a_{j,i} = -a_{i,j}$. ∎

1.4 Planar graphs

A *plane graph* (more accurately, a planar embedding of a graph) is a drawing of a graph $G = (V, E)$ in the Euclidean plane such that no two edges cross. While this intuition is usually enough for this article, it is not a formal definition. Formally, we have injective mapping $i : V \to \mathbb{R}^2$. A drawing of an edge $e = \{u, v\}$ is a line segment $\ell_e : [0, 1] \to \mathbb{R}$ given by $\ell_e(t) = t \cdot i(u) + (1 - t)i(v)$, that is, it connectes the points $i(u)$ and $i(v)$. (By *Fáry's theorem*, it is enough to consider line segments as drawing of edges. More complex curves, as long as they are homeomorphic to $[0, 1]$ do not help.) A graph G is called *planar*, if there is a planar embedding of the graph.

Let $H \subset \mathbb{R}^2$ be the union of the images of ℓ_e for all $e \in E$. $\mathbb{R}^2 \setminus H$ is a collection of connected regions, the so-called *faces* of the plane graph. Since H is bounded, there is exactly one unbounded region in $\mathbb{R}^2 \setminus H$, the *outer face*. Note that the faces depend on the concrete embedding chosen.[1]

Theorem 1.11 (Euler's formula) *Let G be a connected plane graph, let v be its number of nodes, e be its number of edges, and f be its number of faces. Then*

$$v - e + f = 2.$$

[1] This is a very brief introduction of planar graphs. Planar embeddings are topological objects, we replaced a lot of proofs by intuition. See [Die05] for more details.

Fig. 2. (a) A non-planar embedding of a graph. (b) A planar embedding of the same graph. (c) The faces of this embedding.

Proof. The proof is by induction on the number of edges. For an edge e, $G \setminus e$ denotes the graph that is obtained by removing e from G. Note that if G is a planar graph, then $G \setminus e$ is planar, too (by the same embedding).

Induction base: If G has no edges, then it consists of a single node, since it is connected. There is exactly one face, the outer face. We have $1-0+1 = 2$.

Induction step: If G is acyclic, then it is a tree. A tree with v nodes has $v-1$ edges. Every planar embedding of a tree has one face, the outer face. We have $v-(v-1)+1 = 2$.

If G has a cycle, then choose an edge e on a cycle. $G \setminus e$ is still connected. (The planar embedding of) $G \setminus e$ has exactly one face less than (the corresponding embedding of) G. By the induction hypothesis applied to $G \setminus e$, $v-(e-1)+(f-1) = 2$ which yields $v-e+f = 2$. ∎

1.5 Paffian orientations

Observation 1.12 *If G has a node v of degree one, then the unique egde $e = \{u, v\}$ incident with v—a so-called "dangling edge"—has to be part of every perfect matching of G. Any edge incident with u cannot be part of any perfect matching of G. So we can remove all these edges. In this way, we can successively remove any dangling edge until we get a graph in which every node has degree at least two. We will assume this in the following.*

Let G be a graph and $\vec{G}$ be an orientation of G. We call an even cycle C of G *oddly oriented (evenly oriented)* in $\vec{G}$ if the number of edges pointing in one direction is odd (even). Note that since C is even, this is well-defined: If an odd (even) number of edges is pointing in one direction, then an odd (even) number points in the other direction.

Definition 1.13 *A orientation $\vec{G}$ of a graph G is called a Pfaffian orientation, if for every even cycle C such that $G \setminus V(C)$ has a perfect matching, C is oddly oriented.*

Here, $G \setminus V(C)$ denotes the graph obtained by removing all vertices of C from G and all incident edges. The following theorem points out the importance of Pfaffian orientations. If $\vec{A}$ is the adjacency matrix of a Pfaffian orientation $\vec{G}$, then all permutations π_M will have the sign, in particular, the Pfaffian of $\vec{A}$ will compute the number of perfect matchings.

Theorem 1.14 *Let $\vec{G}$ be an orientation of a graph G. The following two statements are equivalent:*

1. *Every perfect matching M has the same sign relative to $\vec{G}$.*
2. *$\vec{G}$ is a Pfaffian orientation.*

Proof. (1) $\Rightarrow$ (2): Let C be a cycle such that $G \setminus V(C)$ has a perfect matching. Let F be a perfect matching of $G \setminus V(C)$ and decompose C into two disjoint matchings S and T. Let $M = F \cup S$ and $N = F \cup T$. By Lemma 1.9,

$$\operatorname{sgn}(\pi_M)\operatorname{sgn}(\pi_N) = (-1)^k$$

where k is the number of evenly oriented cycles in $M \cup N$. $M \cup N$ consists of the cycle C. All other cycles are 2-cycles. They are oddly oriented. Since by (1), $\operatorname{sgn}(\pi_M) = \operatorname{sgn}(\pi_N)$, we get that k is even. Therefore $k = 0$ and $\vec{G}$ is a Pfaffian orientation.

(2) $\Rightarrow$ (1): Conversely, consider two perfect matchings M and N of G. Consider a cycle C of $M \cup N$. $G \setminus V(C)$ contains a perfect matching, namely, the remaining edges of M or N. Since $\vec{G}$ is a Pfaffian orientation, the cycle C is oddly oriented. Since this is true for every cycle of $M \cup N$, the number of evenly oriented cycles in $M \cup N$ is 0. By Lemma 1.9,

$$\operatorname{sgn}(\pi_M)\operatorname{sgn}(\pi_N) = 1.$$

Therefore, π_M and π_N have the same sign. ∎

In a plane graph, the embedding defines a direction. We can talk about clockwise direction and anti-clockwise direction.

Theorem 1.15 *Let $\vec{G}$ be an orientation of a plane graph such that every face except maybe the outer face has an odd number of clockwise edges. Then $\vec{G}$ is a Pfaffian orientation.*

For the proof of the theorem, we need the following lemma.

Lemma 1.16 *Let $\vec{G}$ be as above. Let C be a cycle in $\vec{G}$, c be the number of edges in clockwise direction of C, and v_{in} be the number of vertices strictly contained in C. Then*

$$c \equiv v_{in} - 1 \mod 2.$$

Proof. Let D be the plane graph induced by the vertices of C and all interior vertices. Let e_{in} be the number of edges contained in C. Let f_{in} be the number of faces contained in C. Number the faces inside C arbitrarily. Let c_i be the number of clockwise edges of the ith face. Let k be the length of C. D has $v_{in} + k$ vertices, $e_{in} + k$ edges, and $f_{in} + 1$ faces. By Euler's formula,

$$v_{in} + k - (e_{in} + k) + (f_{in} + 1) = 2,$$

which implies

$$e_{in} = v_{in} + f_{in} - 1.$$

Every interior edge appears in exactly two faces, in one it appear in clockwise direction and in the other one in anti-clockwise direction. Therefore,

$$\sum_{i=1}^{f_{in}} c_i = c + e_{in}.$$

Since $c_i \equiv 1 \mod 2$, we have

$$\begin{aligned} f_{in} &\equiv \sum_{i=1}^{f_{in}} c_i \\ &\equiv c + e_{in} \\ &\equiv c + v_{in} + f_{in} - 1 \mod 2. \end{aligned}$$

Therefore, $c \equiv v_{in} - 1 \mod 2$. ∎

Proof of Theorem 1.15. Let C be an arbitrary even cycle. $G \setminus V(C)$ has a perfect matching. Therefore, v_{in} is even. Thus, C is oddly oriented by Lemma 1.16. ∎

To every plane graph G, we can define its *dual graph*. The nodes of the dual graph are the faces of G. Two faces are connected by an edges if they share an edge. The dual of a plane graph is planar. We get a plane embedding by placing a node of the dual graph inside the face of G and connect them by arcs. By Fáry's theorem, there is an embedding that only uses line segments as curves.

Theorem 1.17 *Every planar graph G has a Pfaffian orientation. It can be computed in polynomial time.*

Proof. Let D be the dual graph of G. Let T be a spanning tree of G. Let S be the subgraph of D that consists of all edges that do not cross an edge of T.

We claim that S is a spanning tree of D: If S contained a cycle C, then this cycle would separate the graph G into two parts, the graph induced by the nodes inside the cycle and by the nodes outside of the cycles. But

since T is a spanning tree, one of the edges of T has to cross an edge of C. Therefore, S cannot have any cycles.

Furthermore, S is connected. If S were not connected, consider one of its connected components K. Every edge of D leaving K intersects one of the edges of T. But this means that T contains a cycle, a contradiction.

We give the edges of T an arbitrary orientation. After this, every face has at least one edge without an orientation, since S is a spanning tree of the dual graph. We now start with the faces that are the leaves of S. We pick an arbitrary one, but not the outer one. Since an unrooted tree has at least two leaves, we can always achieve this. We orient the edge crossed by the edge of S in such a way, that an odd number of edges point in clockwise direction. We remove the edge of S and proceed inductively. By construction, it is clear that we end up with an orientation such that every face has an odd number of edges pointing in clockwise direction. By Theorem 1.15, this is a Pfaffian orientation. ■

1.6 The algorithm

Here is the final algorithm. The first algorithm by Temperley–Fisher and Kasteleyn only computed Pfaffian orientations on grid graphs. It was later generalized by Kasteleyn to planar graphs [Kas67].

Algorithm 1 The FKT algorithm

Input: A planar graph G
Output: The number of perfect matchings in G
1: Successively emove all dangling edges
2: Compute a Pfaffian orientation using Kasteleyn's algorithm.
3: Let $\vec{A}$ be the corresponding adjacency matrix.
4: Return $\sqrt{\det \vec{A}}$.

Algorithm 2 Kasteleyn's algorithm

Input: A planar graph G
Output: A Pfaffian orientation of G

1: Compute a spanning tree T of G.
2: Compute the dual graph D of G with respect to some planar embedding.
3: Let S be the set of all edges of D that do not cross an edge of T. (S is a spanning tree.)
4: Orient the edges of T arbitrarily.
5: **while** S contains an edge **do**
6: Let F a a leaf of S (but not the outer face).
7: Orient the one edge of F crossed by the edge of S in such a way that an odd number of edges point in clockwise direction.
8: Remove F from S.

2 The Holant framework

We introduce the *Holant framework*, see [Val08, CL07a, CLX08]. We follow the exposition of Curticapean [Cur15].

2.1 Basic definitions

For a graph G, $V(G)$ denotes its set of nodes and $E(G)$ its set of edges. For a vertex $v \in V(G)$, $I(v)$ denotes the set of indident edges.

Definition 2.1 *A* signature graph Ω *is an edge-weighted multigraph with a* vertex function $f_v : \{0,1\}^{I(v)} \to \mathbb{Q}$ *associated with each vertex in* $v \in V(\Omega)$.

The *Holant* is a particular sum over edge assignments $x \in \{0,1\}^{E(\Omega)}$. An edge $e \in E(\Omega)$ is *active* if $x(e) = 1$. Otherwise, we call it *inactive*. One can think of x as the characteristic function of the set of active edges. Therefore, we will often write x for the set of active edges $\{e \mid x(e) = 1\}$, too. For a subset $F \subseteq E(\Omega)$, $x|_F$ denotes the restriction of x to F.

Definition 2.2 *Let* Ω *be a signature graph with edge weights* $w : E(\Omega) \to \mathbb{Q}$ *and vertex functions* $f_v : \{0,1\}^{I(v)} \to \mathbb{Q}$ *for each* $v \in V(\Omega)$.

1. *Let* $x \in \{0,1\}^{E(\Omega)}$ *be an edge assignment. Then*

$$\mathrm{val}_\Omega(x) := \prod_{v \in V(\Omega)} f_v(x|_{I(v)}), \tag{2.1}$$

$$w_\Omega(x) := \prod_{e \in x} w(e) \tag{2.2}$$

   are the value *and* weight *of* x*, respectively.*

2. x satisfies Ω *if* $\mathrm{val}_\Omega(x) \neq 0$.

3. *The* Holant *of* Ω *is*

$$\mathrm{Holant}(\Omega) = \sum_{x \in \{0,1\}^{E(\Omega)}} w_\Omega(x) \cdot \mathrm{val}_\Omega(x). \tag{2.3}$$

The name Holant is a blending of the words "holographic" and "permanent/determinant" (or even immanant). There are two important special cases of Holant problems. First of all, there are unweighted signature graphs. Here the weight of every edge is 1. An second important special case are

Boolean vertex functions. Here the range of the vertex functions is simply $\{0,1\}$. In this case, we can rewrite the Holant as

$$\text{Holant}(\Omega) = \sum_{x \text{ satisfies } \Omega} w_\omega(x).$$

Now two natural questions arise:

- What interesting problems can be written as a Holant?
- When are there efficient algorithms for computing the Holant?

Here is an example that gives a first answer to both questions: Given an unweighted graph G, we define vertex functions $f_v : \{0,1\}^{I(v)} \to \{0,1\}$ by

$$f_v(x) = \begin{cases} 1 & \text{if exactly one entry of } x \text{ is } 1 \\ 0 & \text{otherwise} \end{cases}$$

Let Ω be the resulting signature graph. An assignment x satisfies Ω, if every node is incident with exactly one active edge, which means that e is a perfect matching. Since the signature graph is unweighted, all assignments have weight 1. Therefore,

$$\text{Holant}(\Omega) = \sum_{x \text{ satisfies } \Omega} 1 = \text{PerfMatch}(G).$$

We can extend PerfMatch to weighted graphs by defining the weight of a perfect matching to be the product of the weights of the edges in it. The above equation is also true for weighted graphs. Essentially PerfMatch on planar graphs is the only problem of this kind we know an efficient algorithm for and we will try to reduce other problems to it.

2.2 Signatures

In principle, every node of a signature graph can have its own vertex function. In practice, however, we will only use very few different types of vertex functions. We will make use of *signatures* to simplify this, which are functions over $\{0,1\}^d$.

Let Ω be a signature graph and $v \in V(\Omega)$ a node of degree d. We order the edges of $I(v)$ by a bijection $\sigma_v : \{1,\dots,d\} \to I(v)$. We can now view a vertex function $f_v : \{0,1\}^{I(v)} \to \mathbb{Q}$ as a function $\{0,1\}^d \to \mathbb{Q}$.

Definition 2.3 *1. A signature of arity d is a function $s : \{0,1\}^d \to I(v)$.*

2. A signature family is a sequence of signatures $(s_d)_{d\in\mathbb{N}}$, one for each arity.

Definition 2.4 *Let Ω be a signature graph, $v \in V(\Omega)$ a node of degree d and $\sigma_v : \{1, \ldots, d\} \to \mathbb{Q}$ be a bijection.*

1. *Let $x : E(\Omega) \to \{0,1\}$ be an edge assignment. $\sigma_v x \in \{0,1\}^d$ is the assignment that maps each $i \in \{1, \ldots, d\}$ to $x(\sigma_v(i))$.*
2. *$\sigma_v f_v : \{0,1\}^d \to \mathbb{Q}$ is the function defined by $\sigma_v f_v(\sigma_v x) := f_v(x)$ for all assignments $x \in \{0,1\}^{E(\Omega)}$.*

The function $\sigma_v f_v$ describes the behaviour of the f_v when just looking at the edges incident to v. Above, we have seen an example of a signature family, the family $\mathrm{HW}_{=1}$. On input $(x_1, \ldots, x_d)$,

$$\mathrm{HW}_{=1}(x_1, \ldots, x_d) = \begin{cases} 1 & \text{if exactly one entry of } x \text{ is } 1 \\ 0 & \text{otherwise} \end{cases}$$

Note that we skipped the subscript indicating the arity. We will do this when there is not any danger of confusion. "HW" stands for *Hamming weight*. Further examples of signature families that we will encounter frequently are

$$\mathrm{EQ}(x_1, \ldots, x_d) := \begin{cases} 1 & \text{if } x_1 = x_2 = \cdots = x_d \\ 0 & \text{otherwise} \end{cases}$$

$$\mathrm{ODD}(x_1, \ldots, x_d) := \begin{cases} 1 & \text{if an odd number of } x_i \text{ is } 1 \\ 0 & \text{otherwise} \end{cases}$$

$$\mathrm{EVEN}(x_1, \ldots, x_d) := 1 - \mathrm{ODD}(x_1, \ldots, x_d)$$

Definition 2.5 *1. A signature of arity d is symmetric its value only depends on the number of 1s in the given input $x \in \{0,1\}^d$.*

2. A signature family is symmetric if all its members are symmetric

All families above are symmetric.

2.3 Examples

Here are more examples of problems that we can express as Holants. The first one is the matching and the matchsum polynomial. So far, we have seen the perfect matching polynomial, which was defined as

$$\mathrm{PerfMatch}(G) = \sum_{M \in \mathcal{PM}(G)} \prod_{e \in M} w(e).$$

for a graph G with edge weights $w(e)$ for all $e \in E(G)$. This is a generalization of the function PerfMatch in the previous section, where each edge had weight 1. Note however, that also FKT algorithm can handle edge weights.

The (multivariate) *matching polynomial* is defined in the same way by summing over all matchings:

$$\mathrm{Match}(G) = \sum_{M \in \mathcal{M}(G)} \prod_{e \in M} w(e).$$

It is easy to see that we can write Match as a Holant in the same manner as PerfMatch. We just have to replace the signature family $\mathrm{HW}_{=1}$ by $\mathrm{HW}_{\leq 1}$, which is 1 when the Hamming weight of the input assignment is ≤ 1.

The (multivariate) *matching defect polynomial* or *match-sum* polynomial is defined as

$$\mathrm{MatchSum}(G) = \sum_{M \in \mathcal{M}(G)} \prod_{v \in \mathrm{usat}(M)} w(v)$$

Here $\mathrm{usat}(M)$ denotes all nodes of G that are not incident with an edge of M and every node $v \in V(G)$ has a weight $w(v)$. In $\mathrm{Match}(G)$, we record the product of the edge-weights of the edges in the matching, whereas in $\mathrm{MatchSum}(G)$, we record the product of node-weights of unmatched nodes.

We can write MatchSum as a holant problem. Consider the vertex function

$$f_w(x) = \begin{cases} w & \text{if } x = 0 \ldots 0, \\ 1 & \text{if } x \text{ contains exactly one 1,} \\ 0 & \text{otherwise.} \end{cases}$$

We equip every node v of a given graph G with the vertex function $f_{w(v)}$ (of arity $\deg(v)$). Let Ω be the resulting signature graph. If an edge assignment x is not a matching, then one of the vertex functions will be zero. Therefore, only matchings can contribute to $\mathrm{Holant}(\Omega)$. If a node is incident to exactly one active edge, then it contributes 1 to $\mathrm{val}_\Omega(x)$. If it is not matched, it contributes $w(v)$. Therefore, $\mathrm{val}_\Omega(x) = \prod_{v \in \mathrm{usat}(x)} w(v)$ and $\mathrm{Holant}(\Omega) = \mathrm{MatchSum}(G)$.

Finally, we look at a problem that is not a graph problem: Recall that #k-SAT is the problem of counting the number of satisfying solutions of a given formula ϕ in k-CNF. Assume that ϕ has variables $x_1, \ldots, x_n$ and clauses $c_1, \ldots, c_m$. For every variable x_i, we create a variable node v_i. For every clause c_j, we create a clause node u_j. A variable node v_i is connected to a clause node u_j iff x_i appears in c_j. Every variable node v_i get the signature $\mathrm{EQ}_{\deg(v_i)}$ and every clause node u_j gets the clause itself as a signature (or more precisely, its interpretation). We have to ensure that the order of the edges incident to u_j corresponds to the order of the variables in the clause.

The equality signatures ensure that every satisfying assignment x assigns the same value to all edges incident to one variable node. Therefore, x induces an assignment to the variables. Furthermore, x can only be satisfying if all clause signatures have value 1. But this is the case iff the induced assignment satisfies all clauses. Hence every satisfying assignment of Ω corresponds to

exactly one satisfying assignment of ϕ. Therefore, Holant$(\Omega) = \#k\text{-SAT}(\phi)$. Since $\#k$-SAT is #P-hard, this means that not every Holant can be efficiently evaluated.

2.4 Gates

Given a signature graph Ω, we can replace any connected vertex set $S \subseteq V(\Omega)$ by a single vertex and give this new vertex a signature in such a way, that Holant(Ω) is preserved. Given such a set S, there are edges between nodes in S, which we call internal edges, and edges with one endpoint in S, that is, edges crossing the cut $(S, \bar{S})$. The latter edges "describe" the interaction of the nodes in S we the remaining graph.

We also want to reverse this process. Given a node v , we want to replace it by some subgraph H such that Holant(Ω) is not changed. This gets interesting, we the inserted graph uses "simpler" vertex functions than the one of the replaced node. This subgraph will have internal edges and edges connected with the neighbours of v. If we look at H in isolation, these crossing edges will become "dangling" edges, that is, edges having only one endpoint. Once H is inserted into a signature graph Ω, the other endpoint will be a node of Ω.

Definition 2.6 *1. A* gate *is a signature graph Γ, possibly containing some set $D \subseteq E(\Omega)$ of dangling edges. All dangling edges have weight* 1*.*

2. The signature *of Γ is the function* $\mathrm{Sig}(\Gamma) : \{0,1\}^D \to \mathbb{Q}$ *defined by*

$$\mathrm{Sig}(\Gamma, x) = \sum_{y \in \{0,1\}^{E(\Omega)\setminus D}} w_\Gamma(xy)\,\mathrm{val}_\Gamma(xy)$$

Here xy denotes the combined assignment that assigns dangling edges values according to x and other edges values according to y.

Let Ω be a signature graph and let $S \subseteq V(\Omega)$. Let $E(S, \bar{S})$ denote the set of cut edges crossing the cut $(S, \bar{S})$.

1. We denote by Ω_S the *gate induced by S*. Ω_S is the induced subgraph $\Omega[S]$ together with the cut edges as dangeling edges.

2. The *contraction* $\Omega{\downarrow_S}$ is obtained by contracting the vertices of S into a single vertex $v{\downarrow_S}$. Every edge in $E(S, \bar{S})$ is now incident with $v{\downarrow_S}$. This might create parallel edges.

3. Conversely, we can also *insert* a gate Γ into a signature graph Ω by taking a vertex v whose degree equals the number of dangling edges of Γ. Each dangling edge of Γ is identified with on edge incident to v and

then v is replaced by Γ. Which dangling edge is identified with which edge of v will usually be clear from the context.

Lemma 2.7 *Let Ω be a signature graph and let $S \subseteq V(\Omega)$ such that all edges in $E(S, \bar{S})$ have weight one. Then* $\mathrm{Holant}(\Omega) = \mathrm{Holant}(Omega{\downarrow_S})$.

Proof. We divide $E(\Omega)$ into three sets, $C = E(S, \bar{S})$, $X = E(\Omega[S])$, and $Y = E(\Omega \setminus S)$. So C are the cut edges, X are the edges inside S and Y are the edges (completely) outside S. We have $V(\Omega{\downarrow_S}) = \bar{S} \cup \{v{\downarrow_S}\}$ and $E(\Omega{\downarrow_S}) = \tilde{C} \cup Y$. $\tilde{C}$ is the set of edges obtained from C by replacing in every edge the vertex from S by $v{\downarrow_S}$.

Let w be the weight function of Ω and f_v be the vertex functions. We have

$$\begin{aligned}
\mathrm{Holant}(\Omega) &= \sum_{x \in \{0,1\}^{E(\Omega)}} w(x) \left(\prod_{v \in S} f_v(x|_{C \cup X}) \right) \left(\prod_{v \in \bar{S}} f_v(x|_{C \cup Y}) \right) \\
&= \sum_{z \in \{0,1\}^C} \left(\sum_{x \in \{0,1\}^X} w(zx) \prod_{v \in S} f_v(zx) \right) \left(\sum_{y \in \{0,1\}^Y} w(zx) \prod_{v \in \bar{S}} f_v(zy) \right) \\
&= \sum_{z \in \{0,1\}^C} f_{v{\downarrow_S}}(z) \sum_{y \in \{0,1\}^Y} w(zy) \prod_{v \in \bar{S}} f_v(zy) \\
&= \sum_{xy \in \{0,1\}^{C \cup Y}} \prod_{v \in V(\Omega{\downarrow_S})} w(zy) f_v(zy) \\
&= \mathrm{Holant}(\Omega{\downarrow_S})
\end{aligned}$$

Above, we used that we can always add edges from C to the products, since they have weight zero. ∎

2.5 Matchgates

Recall that we want to reduce Holant problems to counting perfect matchings. You can think of counting perfect matchings as a Holant problem that only uses the vertex function $\mathrm{HW}_{=1}$. The idea is now to simulate these functions on one node by replacing the node with a subgraph only using the vertex function $\mathrm{HW}_{=1}$.

Definition 2.8 *A* matchgate Γ *(of arity k) is a gate whose vertices only have* $\mathrm{HW}_{=1}$ *as vertex functions.*

Recall that internal edges of a gate can have arbitrary weights, but all dangelling edges have weight one.

Theorem 2.9 *Let Ω be a signature graph. If there is a matchgate Γ_v with $\mathrm{Sig}_{\Gamma_v} = f_v$, then*

$$\mathrm{Holant}(\Omega) = \mathrm{PerfMatch}(G(\Omega))$$

where $G(\Omega)$ is the graph obtained from Ω by inserting Γ_v for each v.

Proof. By repeated application of Lemma 2.7. ■

The signature of a matchgate can be computed by looking at its perfect matchings. If a dangling edge is active, then its inner node cannot be matched by an internal edge. Otherwise, it has to be matched by an internal edge.

Fact 2.10 *Let Γ be a matchgate of arity d. Given an assignment $x \in \{0,1\}^d$ to the dangeling edges, let $S(x) \subseteq V(\Gamma)$ be the nodes of Γ incident with an active dangling edge. Then $\mathrm{Sig}(\Gamma, x) = \mathrm{PerfMatch}(\Gamma' \setminus S(x))$, where Γ' is Γ with all dangling edges removed.*

$\Gamma' \setminus S(x)$ only admits a perfect matching, if it has an even number of nodes. Thus the fact above implies the so-called *parity conditions*.

Fact 2.11 (Parity conditions) *If a signature f of arity d is realised by a matchgate, then at least one of the following is true:*

1. *For all $x \in \{0,1\}^d$ with odd Hamming weight, $f(x) = 0$. In this case, we call f* even.
2. *For all $x \in \{0,1\}^d$ with even Hamming weight, $f(x) = 0$. We call such an f* odd.

2.6 Planar matchgates

When we consider planar matchgates, the order of the dangling edges can be crucial. So from now on, we assume that the dangling edges are ordered by $1, \ldots, d$.

Definition 2.12 1. *A matchgate is* planar *if it admits a plane drawing with all dangling edges on the outer face. A travesal of the outer face in clockwise direction starting with the dangling edge 1 encounters the dangling edges in order $1, 2, \ldots, d$.*

2. *A signature is called* planar *if it is realized by a planar matchgate.*

Lemma 2.13 *Let Ω be a signature graph and let π be a drawing of Ω in the plane. For each $v \in V(\Omega)$, let σ_v be the clockwise ordering of $I(v)$ with respect to π. Assume that there is a matchgate Γ_v for each v realising the $\sigma_v f_v$, where f_v is the vertex function of v. Then $G(\Omega)$ admits a plane drawing.*

Fig. 3. Planar matchgates for arity 1

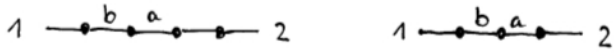

Fig. 4. Planar matchgates for arity 2

Proof. We can locally extend π to a drawing of $G(\Omega)$. Each v is replaced by Γ_v, drawn small enough. Note that the ordering of $I(v)$ and the input bits of $\sigma_v f_v$ is the same. ■

Valiant [Val08] proved that every signature of arity $d \leq 3$ that satisfies the parity condition is planar.

Lemma 2.14 *Let $d \leq 3$. If $f : \{0,1\}^d \to \mathbb{Q}$ is even or odd, then f is planar.*

Proof. If $d = 2$, then the gadget on the lefthand side in Figure 4 realizes the odd signature

$$\mathrm{Sig}(\Gamma, x) = \begin{cases} a & \text{if } x = 00 \\ 0 & \text{if } x = 01, 10 \\ b & \text{if } x = 11 \end{cases}$$

The signature on the righthand side realizes the even signature

$$\mathrm{Sig}(\Gamma, x) = \begin{cases} a & \text{if } x = 01 \\ b & \text{if } x = 10 \\ 0 & \text{if } x = 00, 11 \end{cases}$$

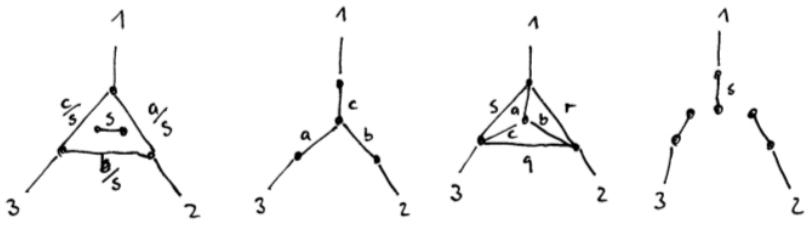

Fig. 5. Planar matchgates for arity 3

If $d = 1$, we can realize f in a similar way by a line of odd or even length, see Figure 3. If $d = 3$, the situation is more complicated. Figure 5 shows four gadgets. It can be verified that after proper substitution of the edge weights, we can realise any odd or even f by one of them. As an example, the left-most gadget has signature

$$\mathrm{Sig}(\Gamma, x) = \begin{cases} a & \text{if } x = 100 \\ b & \text{if } x = 010 \\ c & \text{if } x = 001 \\ s & \text{if } x = 111 \end{cases}$$

if $s \neq 0$. If $s = 0$, then we can use the gadget next to is can be used. The details are left to the reader. ■

3 Holographic reductions

In this chapter, we will look at planar bipartite signature graphs Ω without edge weights. That means, we can divide the nodes $V(\Omega) = V_{gen} \cup V_{rec}$ into two parts, such that all edges in $E(\Omega)$ have one end point in V_{gen} and the other in V_{rec}. We call the nodes in V_{gen} *generators* and in V_{rec} *recognisers*. We will assume that signatures of the generators come from some set G and the signatures of the recognisers from some set R. We call the set of all signatures realizable by matchgates standard signatures **STD**. From the results of the last chapter it follows that when $G, R \subseteq \mathbf{STD}$, then Holant($\Omega$) is polynomial time computable.

3.1 Reductions

A signature s is a function of arity k, $s : \{0,1\}^k \to \mathbb{Q}$ (or to any other field). Alternatively, we can view s as a vector of length 2^k. We index the entries of this vector by bitstrings of length 2^k. When s is a signature of arity 1, then this vector is $(a, b)^T$. We can write it as $ae_0 + be_1$ where e_0 and e_1 denote the standard basis.

Definition 3.1 *Let $A = (a_{i,j}) \in \mathbb{Q}^{m\times n}$ and $B = (b_{i,j}) \in \mathbb{Q}^{s\times t}$ be two matrices. The Kronecker product $A \otimes B \in \mathbb{Q}^{ms\times nt}$ is the block matrix*

$$\begin{pmatrix} a_{1,1}B & \dots & a_{1,n}B \\ \vdots & & \vdots \\ a_{m,1}B & \dots & a_{m,n}B \end{pmatrix}.$$

In general $A \otimes B$ and $B \otimes A$ are not the same matrices. However, it is easy to see that one can be transformed into the other by permutations of row and columns.

In our setting, all our vectors are of size 2^k and we index them by bitstrings. The entries of the Kronecker product of two matrices $A = (a_{x,y}) \in \mathbb{Q}^{2^k\times 2^k}$ and $B = (a_{x,y}) \in \mathbb{Q}^{2^\ell\times 2^\ell}$ can be written as follows: The entry in position (x, y) is given by

$$(a_{x',y'}b_{x'',y''})$$

where we decompose $x = x'x''$ and $y = y'y''$ into bitstrings x' and y' of length k and x'' and y'' of length ℓ, respectively. In particular, the standard basis of $\mathbb{Q}^{2^k}$ consisits of the vectors $e_x = e_{x_1} \otimes e_{x_2} \otimes \cdots \otimes e_{x_k}$.

Not every signature can be realised by matchgates. However, can one realise more signatures by a basis change? It can of course not be that easy. Our hope is to apply different basis transforms to the generators and the recognisers and hope that there effect cancels. Since signatures can have different arities and nodes share different edges, we have to treat every bit of the input of a signature in the same way by a basis change. Let b_0, b_1 be a basis of $\mathbb{Q}^2$. The change to this basis is given by the 2×2-matrix $B = (b_0\, b_1)$. Let $b_{i,j}$ be the entries of B. We will often call B a basis for short. The k-fold tensor power $B^{\otimes k}$ is defined by $B^{\otimes 0} = I$ and $B^{\otimes k} = B \otimes B^{\otimes(k-1)}$. Its entries can be written as

$$(B^{\otimes k})_{x,y} = \prod_{i=1}^{k} b_{x_i,y_i}$$

and its columns as $b_x = \bigotimes_{i=1}^{k} b_{x_i}$.

Definition 3.2 *Let B be a basis. Let G, R and G', R' be sets of signatures. We say that (G, R) admits a* holographic reduction *to (G', R') by B if*

1. *for all $g \in G$ of arity k, there is a $g' \in G'$ such that $g' = (B^{\otimes k})^{-1} g$ and*
2. *for all $r \in R$ of arity k, there is an $r' \in R'$ such that $r' = (B^{\otimes k})^T r$.*

In this case, we write $(G, R) \leq_B (G', R')$. If there is such a B, we write $(G, R) \leq_{hol} (G', R')$.

Given a bipartite signature grid Ω and a basis B, the signature grid $B\Omega$ is defined by replacing every signature $g \in G$ of arity k by $(B^{\otimes k})^{-1} g$ and $r \in R$ by $(B^{\otimes k})^T r$. Here G is the set of signatures of the generators and R is the set of signatures of the recognisers.

Theorem 3.3 (Holant theorem) $\mathrm{Holant}(\Omega) = \mathrm{Holant}(B\Omega)$.

Proof. Let $V(\Omega) = V_{gen} \cup V_{rec}$. Identify $E(\Omega)$ with $\{1, \ldots, m\}$. We order the edges in such a way that the sets $I(r)$ for all $r \in V_{rec}$ are intervals of consecutive numbers. This is possible, since Ω is bipartite.

Let $\sigma \in S_m$ be a permutation such that for all $g \in V_{gen}$, $\sigma(I(g)) = \{\sigma(e) \mid e \in I(g)\}$ are intervals of consecutive numbers. For a bitstring $x \in \{0,1\}^m$, σx is the bitstring when the positions are permutated according to σ. Let $P \in \{0,1\}^{2^m \times 2^m}$ be the corresponding permutation matrix. That is, the entry $p_{x,y}$ is 1 of $\sigma(x) = y$ and 0 otherwise.

Let

$$\rho(x) := \prod_{r \in V_{rec}} f_r(x)$$

$$\gamma(x) := \prod_{g \in V_{gen}} f_g(x)$$

and let ρ and γ denote the corresponding vectors of length 2^m. We have

$$\begin{aligned}\operatorname{Holant}(\Omega) &= \sum_{x\in\{0,1\}^m} \rho(x)\gamma(x)\\ &= \rho^T \cdot \gamma.\end{aligned}$$

Consider two recognisers r_1 and r_2 and some assignment x to the edges incident with r_1 and r_2. Since $I(r_1)$ and $I(r_2)$ are intervals of consecutive numbers, we can write $x = x_1x_2$ such that x_1 is assigned to the edges of r_1 and x_2 is assigned to the edges of r_2. We have

$$\operatorname{val}(x) = f_{r_1}(x_1) \cdot f_{r_2}(x_2) = (f_{r_1} \otimes f_{r_2})_{x_1x_2}.$$

(Recall that we identified functions with vectors.) This generalizes to an arbitrary number of recognisers, so we get

$$\rho = \bigotimes_{r\in V_{rec}} f_r.$$

We can do the same of the generators, however, here we have to reorder the edges according to σ:

$$\gamma = P \bigotimes_{g\in V_{gen}} f_g.$$

In the same way,

$$\begin{aligned}\rho' &= \bigotimes_{r\in V_{rec}} (B^{\otimes \deg(r)})^T f_r\\ \gamma' &= P \bigotimes_{g\in V_{gen}} (B^{\otimes \deg(g)})^{-1} f_g\end{aligned}$$

Since $(A \otimes B)(C \otimes B) = AC \otimes BD$ for matrices A, B, C, D, we get

$$\begin{aligned}\rho' &= (B^{\otimes m})^T \bigotimes_{r\in V_{rec}} (B^{\otimes \deg(r)} f_r = (B^{\otimes m})^T \rho\\ \gamma' &= P(B^{\otimes m})^{-1} \bigotimes_{g\in V_{gen}} f_g.\end{aligned}$$

We get

$$\begin{aligned}\operatorname{Holant}(B\Omega) &= (\rho')^T\gamma'\\ &= \rho^T(B^{\otimes m})P(B^{\otimes m})^{-1} \bigotimes_g f_g\\ &= \rho^T P(B^{\otimes m})(B^{\otimes m})^{-1} \bigotimes_g f_g\\ &= \rho^T P \bigotimes_g f_g\\ &= \rho^T\gamma\\ &= \operatorname{Holant}(\Omega).\end{aligned}$$

The third equality above follows from the lemma below. ■

Lemma 3.4 *Let B and P as above. Then*

$$(B^{\otimes m})P = P(B^{\otimes m}).$$

Proof. We show that the identity holds componentwisely. Let $x, y \in \{0,1\}^{2^m}$. $(B^{\otimes m}P)_{x,y} = (PB^{\otimes m})_{x,y}$ is equivalent to

$$\sum_{z\in\{0,1\}^{2^m}} B^{\otimes m}_{x,z} P_{z,y} = \sum_{z\in\{0,1\}^{2^m}} P_{x,z} B^{\otimes m}_{z,y}.$$

Since $P_{x,z} = 1$ for exactly one value of z, namely $z = \sigma(x)$, the equation above is equivalent to

$$B^{\otimes m}_{x,\sigma^{-1}(y)} = B^{\otimes m}_{\sigma(x),y} \iff \prod_{i=1}^{m} b_{x(i),y(\sigma^{-1}(i))} = \prod_{i=1}^{m} b_{x(\sigma(i)),y(i)}.$$

The latter product is just a reordering of the former, hence they are the same. ■

3.2 Consequences

Let G, R be sets of signatures. We write Holant$[G, R]$ for the problem of computing the Holant of bipartite graphs with the signatures of the generators taken from G and the signatures of the recognizers taken from R.

If $(G, R) \leq_{hol} (G', R')$, then the Holant theorem yields

$$\text{Holant}[G, R] \leq_{par} \text{Holant}[G', R'].$$

In particular, if $(G, R) \leq (\mathbf{STD}, \mathbf{STD})$, then

$$\text{Holant}[G, R] \leq_{par} \text{Holant}[\mathbf{STD}, \mathbf{STD}] \leq_{par} \text{PerfMatch}\,.$$

Thus, Holant$[G, R]$ restricted to planar graphs is in FP by the FKT-algorithm.

On the other hand, if Holant$[G, R]$ is #P-hard and $(G, R) \leq_{hol} (G', R')$, then we also know that Holant$[G', R']$ is #P-hard.

3.3 Example

Here is a rather surprising example. Let Pl-Rtw-Mon-k-SAT be the following problem: Given a formula in k-CNF in which every variable appears exactly twice and only positive such that the underlying graph structure of the formula (as introduced in Section 2.3) is planar. We will prove in the next Section 4 that this problem is #P-hard and that counting the solutions

modulo 2 is $\oplus$P-hard. However, we will prove now by a holographic reduction that counting the solutions modulo 7 can be done in polynomial time. If you recall the construction from Section 2.3, we need to realize the signature EQ_2 for the generators and OR_3 for the recognizers.

Consider the following basis $(n,p) := (\binom{n_0}{n_1}, \binom{p_0}{p_1}) := (\binom{1}{6}, \binom{3}{5})$. This basis seems to come out of the blue, but we will see in later sections how to systematically find such bases.

Theorem 3.5 *The basis (n,p) above simultaneously realises $[1,0,1]$ as a generator and $[0,1,1,1]$ as a recognizer.*

Proof. Let

$$B = \begin{pmatrix} n_0 & p_0 \\ n_1 & p_1 \end{pmatrix} = \begin{pmatrix} 1 & 3 \\ 6 & 5 \end{pmatrix}.$$

We have

$$B^{\otimes 2} = \begin{pmatrix} 1 & 3 & 3 & 9 \\ 6 & 5 & 18 & 15 \\ 6 & 18 & 5 & 15 \\ 36 & 30 & 30 & 25 \end{pmatrix}.$$

It is easy to check that

$$B^{\otimes 2} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 3 \\ 0 \\ 0 \\ 5 \end{pmatrix} \mod 7$$

By Lemma 2.14, we can realize $(3,0,0,5) = [3,0,5]$ as a generator. By the equation above, it follows that this matchgate realises $(1,0,0,1) = [1,0,1]$ under the basis (n,p).

Furthermore,

$$B^{\otimes 3} = \begin{pmatrix} 1 & 3 & 3 & 9 & 3 & 9 & 9 & 27 \\ 6 & 5 & 18 & 15 & 18 & 15 & 54 & 45 \\ 6 & 18 & 5 & 15 & 18 & 54 & 15 & 45 \\ 36 & 30 & 30 & 25 & 108 & 90 & 90 & 75 \\ 6 & 18 & 18 & 54 & 5 & 15 & 15 & 45 \\ 36 & 30 & 108 & 90 & 30 & 25 & 90 & 75 \\ 36 & 108 & 30 & 90 & 30 & 90 & 25 & 75 \\ 216 & 180 & 180 & 150 & 180 & 150 & 150 & 125 \end{pmatrix}$$

We have

$$(B^{\otimes 3})^T \begin{pmatrix} 0 \\ 3 \\ 3 \\ 0 \\ 3 \\ 0 \\ 0 \\ 5 \end{pmatrix} = \begin{pmatrix} 1134 \\ 1023 \\ 1023 \\ 1002 \\ 1023 \\ 1002 \\ 1002 \\ 1030 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} \mod 7.$$

By Lemma 2.14, $[0, 3, 0, 5] = (0, 3, 3, 0, 3, 0, 0, 5)$ is realisable by a planar matchgate under the standard basis. ■

Corollary 3.6 *Pl-Rtw-Mon-3-SAT is in* FP*, when counting modulo* 7*.*

4 Hardness of Pl-Rtw-Mon-3-SAT

Theorem 4.1 *Pl-Rtw-Mon-3-SAT is #P-complete and $\oplus$P-complete when we count modulo 2.*

5 Matchgate Identities

When can a standard signature be realised by a planar matchgate? Such signatures have to fulfill the so-called *matchgate identities*. Matchgate identites were introduced by Valiant [Val02]. We here follow the exposition by Cai and Gorenstein [CG14].

In the following, F will be the underlying field and e_i will denote the ith unit vector (of appropriate length, which is hopefully clear from the context).

Theorem 5.1 (Matchgate Identities) *$u \in F^{2^n}$ is the standard signature of a planar matchgate Γ of arity n iff for all $\alpha, \beta \in \{0,1\}^n$*

$$\sum_{i=1}^{\ell} (-1)^i u_{\alpha \oplus e_{\gamma_i}} u_{\beta \oplus e_{\gamma_i}} = 0, \tag{5.1}$$

where γ_i is the position of the ith nonzero bit in $\alpha \oplus \beta$ and ℓ is the Hamming weight of $\alpha \oplus \beta$.

We will first prove the theorem above for the Pfaffian. Let A be skew-symmetric $n \times n$-matrix We write $\operatorname{pf}(i_1, \ldots, i_m)$ for the Pfaffian of the principal submatrix of A obtained by selecting the rows and columns $i_1, \ldots, i_m$.

Theorem 5.2 (Grassmann–Plücker identities) *Let $K, L \in \mathbb{N}$ be odd. Let $1 \le i_1 < \cdots < i_K \le n$ and $1 \le j_1 < \cdots < j_L \le n$. Then*

$$\begin{aligned} &\sum_{\ell=1}^{L} (-1)^{\ell-1} \operatorname{pf}(j_\ell, i_1, \ldots, i_K) \operatorname{pf}(j_1, \ldots, \hat{j}_\ell, \ldots, j_L) \\ &+ \sum_{k=1}^{K} (-1)^{k-1} \operatorname{pf}(i_k, j_1, \ldots, j_L) \operatorname{pf}(i_1, \ldots, \hat{i}_k, \ldots, i_K) = 0, \end{aligned}$$

where $i_1, \ldots, \hat{i}_k, \ldots, i_K$ denotes the sequence $i_1, \ldots, i_K$ with i_k omitted.

Proof. Exercise. ∎

Lemma 5.3 *Theorem 5.1 is true when u is replaced by the Pfaffian.*

Proof. Write $I := \{i_1, \ldots, i_K\}$ and $J := \{j_1, \ldots, j_L\}$ and let $\{k_1, \ldots, k_M\} =: I \Delta J$ be the symmetric difference. We can rewrite the Grassmann-Plücker identities as

$$\sum_{m=1}^{M} (-1)^m \operatorname{pf}(I \Delta \{k_m\}) \operatorname{pf}(J \Delta \{k_m\}) = 0. \tag{5.2}$$

These are exactly the matchgate identities. ∎

For $\alpha \in \{0,1\}^n$, let Z_α be the set of output nodes with $\alpha_i = 1$. Let pf^α be the Pfaffian of the matrix of $\Gamma^\alpha := \Gamma \setminus Z_\alpha$, Fix a Pfaffian orientation of Γ and let $\delta(\alpha) = \mathrm{pf}^\alpha / u_\alpha$. Recall that u_α is the (weighted) number of perfect matchings in $\Gamma \setminus Z_\alpha$. Since we have chosen a Pfaffian orientation, u_α and pf^α can only differ in sign, that is, $\delta(\alpha) \in \{1, -1\}$.

Compare equations (5.1) and (5.2): Nonzero terms stand in bijection. They can only differ by sign. Let

$$\epsilon_\ell = \text{sign difference of the } \ell\text{th summand} \in \{1, -1\}.$$

Claim: for all $i \neq j$, $\epsilon_i \cdot \epsilon_j = 1$, that is

$$\delta(\alpha \oplus e_{\gamma_i})\delta(\alpha \oplus e_{\gamma_j}) = \delta(\beta \oplus e_{\gamma_i})\delta(\beta \oplus e_{\gamma_j}).$$

Since $\alpha_{\gamma_i} = \bar{\beta}_{\gamma_i}$ for all i, the claim follows from the following lemma.

Lemma 5.4 *For all $1 \leq i < j \leq k$ and $u, \tilde{u} \in \{0,1\}^{i-1}$, $v, \tilde{v} \in \{0,1\}^{j-i-1}$, $w, \tilde{w} \in \{0,1\}^{k-j}$ and $b, c \in \{0,1\}$:*

$$\delta(ubvcw)\delta(u\bar{b}v\bar{c}w) = \delta(\tilde{u}b\tilde{v}c\tilde{w})\delta(\tilde{u}\bar{b}\tilde{v}\bar{c}\tilde{w}) \tag{5.3}$$

Proof. We here prove it only for the case $b = c = 0$. The other cases are proven in a similar fashion. Attach a path of length 2 to each output node of Γ. This leaves its signature unchanged. A Pfaffian orientation $\vec{\Gamma}$ of Γ induces a Pfaffian orientation $\vec{\Gamma}^\alpha$ of Γ^α. Let M_α be a matching of Γ^α. Let $\mathrm{pf}_{\vec{\Gamma}^\alpha}(M_\alpha)$ be the sign of the summand corresponding to M_α in the Pfaffian and let $u_{\Gamma^\alpha}(M_\alpha)$ be the sign of the summand in the PerfMatch polynomial. Fix u, u, and w and write Γ^{00} for Γ^{u0v0w}, Γ^{11} for Γ^{u1v1w}, M^{00} for M^{u0v0w}, and M^{11} for M^{u1v1w}.

If we show that the value of the lefthand side of (5.3) is independent of u, v, and w, then we will be done. Connect the input node i and j by an edge e. Let $\Gamma^* := \Gamma^{00} \cup \{e\}$. Orient e such that we get a Pfaffian orientation $\vec{\Gamma}^*$. (e creates a new face. The other face of e is the outer face.) Let $M^* = M^{11} \cup \{e\}$.

We have

$$\frac{\mathrm{pf}_{\vec{\Gamma}^{00}}(M^{00})}{u_{\Gamma^{00}}(M^{00})} = \frac{\mathrm{pf}_{\vec{\Gamma}^*}(M^{00})}{u_{\Gamma^*}(M^{00})},$$

since M^{00} does not contain e. Furthermore,

$$\frac{\mathrm{pf}_{\vec{\Gamma}^*}(M^{00})}{u_{\Gamma^*}(M^{00})} = \frac{\mathrm{pf}_{\vec{\Gamma}^*}(M^*)}{u_{\Gamma^*}(M^*)},$$

because $\vec{\Gamma}^*$ is a Pfaffian orientation, that is, all permutations in the sum have the same sign. Finally, we have $u_{\Gamma^*}(M^*) = u_{\Gamma^{11}}(M^{11})$, since e has weight 1.

It remains to prove $\mathrm{pf}_{\vec{\Gamma}^*}(M^*) = \mathrm{pf}_{\Gamma^{11}}(M^{11})$. We can assume that $i < j$ and that i and j are the largest nodes. If M^* contains the edge (i, j), then $\mathrm{sgn}(\pi_{M^*}) = \mathrm{sgn}(\pi_{M^{11}})$. If M^* contains the edge (j, i) instead, then $\mathrm{sgn}(\pi_{M^*}) = -\mathrm{sgn}(\pi_{M^{11}})$. But in this case, we get an additional -1 from the product term. ∎

From the considerations above, the neccessity of the matchgate identities follows. It remains to prove the sufficiency.

Lemma 5.5 *If a signature $u = (u_i)$ of arity n is realisable by a matchgate, then*

1. *$\alpha \cdot u$ is realisable for any $\alpha \in \mathbb{C}$*
2. *$(u_{i \oplus x})$ is realisable for any fixed $x \in \{0,1\}^n$.*

Proof. To prove the first item, we just add an isolated edge of weight α with two internal nodes to a given matchgate. This edge belongs to perfect every matching and adds a factor of α. For the second item, we add to every external node, where the corresponding bit of x is 1, an additional edge. The other node of this edge becomes the new external node. ∎

Lemma 5.6 *If $u \in \mathbb{C}^{2^n}$ satisfies the matchgate identities, then there is a planar matchgate realising u.*

Proof. Since the all-zero signature is certainly realisable, we can assume that at least one entry of u is 1. By the previous lemma, w.l.o.g. $u_{1\ldots1} = 1$.

We start with the complete graph K_n on n vertices (which is of course not planar). We draw the vertices such that they lie in a general position, that is, there are exactly $\binom{n}{4}$ intersection points.

For each assignment α of Hamming weight $n-2$, note that there is exactly one edge left in the remaining graph. We give this edge weight $u_{e_i \oplus e_j \oplus 1\ldots1}$, where i and j are the bits of α that are zero.

Now let α be an assignment of Hamming weight $< n-2$. Since u satisfies the matchgate identities, we have

$$u_\alpha \underbrace{u_{1\ldots1}}_{=1} = \sum_{i=2}^{\ell} (-1)^i u_{\alpha \oplus e_{p_1} \oplus e_{p_i}} u_{1\ldots1 \oplus e_{p_1} \oplus e_{p_i}}$$

by the matchgate identities applied to $\alpha \oplus e_{p_1}$ and $1\ldots1 \oplus e_{p_1}$. Here, $p_1 < p_2 < \cdots < p_\ell$ are the positions of the 0s in α. Note that u_α is expressed in terms of entries of u with larger Hamming weight, Therefore, the u_α of Hamming weight $n-2$ determine the other u's. However, the Pfaffian also satisfies the matchgate identities. Therefore, $\mathrm{pf}(K_n^\alpha) = u_\alpha$.

Fig. 6. The crossover gadget

However, K_n is not planar. We replace every crossing by a planar crossover gadget X as depicted in Figure 6. The signature of X is given by

$$X_{0000} = X_{0101} = X_{1010} = -X_{1111} = 1 \tag{5.4}$$

$$X_\alpha = 0 \qquad \text{for all other } \alpha \tag{5.5}$$

Note that X is not symmetric. By replacing every crossing by X, we get a new signature graph. Every edge $\{i, j\}$ in K_n is cut into $t + 1$ edges, where t is the number of edges that $\{i, j\}$ crosses. We give one of these edges the weight of $\{i, j\}$ and all other edges the weight 1. Let $\hat{K}$ be the resulting graph. Let I denote the set of all edges inside any crossover gadget and let O be the set of all other edges in $\hat{K}$.

Every matching M of K_n^α can be extended to $\hat{K}^\alpha$ in the following way: Every edge $\{i, j\} \in M$ is replaced by the $t+1$ edges in O connecting i with j. All unmatched nodes are then matched by nodes in I: If two or four nodes of a crossover gadget are matched, then there is a unique way to match the internal unmatched nodes. If four nodes are matched, then we introduce an edge of weight -1. If no nodes of the crossover gadget are matched externally, then there are three internal matchings of the gadget, two of weight 1 and one of weight -1. The overall contribution is 1. Therefore,

$$(-1)^{c(M)} w(M) = w(M_o) \sum_{\text{internal matchings } M_i} w(M_i), \tag{5.6}$$

where M_o are the edges in O of the extended matching and $c(M)$ denotes the number of crossings of M in K_n^α.

On the other hand, every perfect matching of $\hat{K}^\alpha$ is of this form: If the external nodes of a crossover gadget are matched externally in a different way than described above, then there is no way to match the nodes internally (since the signature is zero in this case).

To conclude, note that the number of crossings of M is the sign of the matching. Therefore, the righthand side of (5.6) is exactly the contribution of M to the Pfaffian of K_n^α. Therefore, $\mathrm{pf}(K_n^\alpha) = \mathrm{PerfMatch}(\hat{K}^\alpha)$. ∎

6 Realizable Signatures

In this chapter, we analyse which signatures are realizable (under any basis). These results were proven by Cai and Lu [CL07b].

Recall that we write symmetric signatures of arity k as $[\sigma_0, \ldots, \sigma_k]$.

Theorem 6.1 *Let $[\sigma_0, \ldots, \sigma_k]$ be an even symmetric matchgate signature. Then there are r_1, r_2, not both equal to zero, such that for all even $i \geq 2$,*

$$r_1 z_{i-2} = r_2 z_i$$

Proof. Let $j > i$ be even. By the matchgate identities (with $\alpha = 1^i 10^{k-i-1}$ and $\beta = 1^i 01^{j-i-1} 0^{k-j}$) we get

$$\sigma_i \sigma_j = \sum_{h=i+2}^{j-i-1} (\pm 1) \sigma_{i+2} \sigma_{j+1}$$

(Note that α has Hamming weight $i+1$ and β has Hamming weight $j-1$ and that α and β differ in the positions $i+1, \ldots, j-i-1$.) In the above equation, there is an odd number of summands on the right-hand side with alternating signs. The first one has a plus sign, therefore, the right-hand side is simply $\sigma_{i+2}\sigma_{j-2}$. Therefore, we get

$$\sigma_i \sigma_j = \sigma_{i+2} \sigma_{j-2}. \tag{6.1}$$

When $i \leq k-4$, we can choose $j = i+4$ and get

$$\sigma_i \sigma_{i+4} = \sigma_{i+2}^2. \tag{6.2}$$

Therefore, if $\sigma_{i+2} \neq 0$, then $\sigma_i \neq 0$ and $\sigma_{i+4} \neq 0$.

We distinguish two cases: First, if any σ_i with an even index except the leftmost or rightmost one is nonzero, then we get that all even indices are nonzero and by (6.2), the quotient of two adjacent even entries is constant, i.e.,

$$\frac{\sigma_{i+2}}{\sigma_i} = \frac{\sigma_{i+4}}{\sigma_{i+2}} =: \frac{r_2}{r_1}.$$

In the second case, all inner entries with an even index are zero. If $k \leq 3$, then the statement of the theorem is trivial, since there are only two entries with an even index. Therefore, assume $k \geq 4$. Let $4 \leq k^* \leq k$ be the largest even index. By (6.1),

$$z_0 z_{k^*} = \underbrace{z_2}_{=0} \underbrace{z_{k^*-2}}_{=0} = 0.$$

Thus either $z_0 = 0$ oder $z_{k^*} = 0$ and can choose $r_1 = 1$ and $r_2 = 0$ or vice versa. ■

The above theorem can be proven for odd signatures in a similar fashion This means that there are only 4 types of symmetric signatures of arity k (up to scaling):

1. $[a^k b^0, 0, a^{k-1}b^1, 0, \ldots, a^0 b^k]$
2. $[a^k b^0, 0, a^{k-1}b^1, 0, \ldots, a^0 b^k, 0]$
3. $[0, a^k b^0, 0, a^{k-1}b^1, 0, \ldots, a^0 b^k]$
4. $[0, a^k b^0, 0, a^{k-1}b^1, 0, \ldots, a^0 b^k, 0]$

for parameters $a, b \in \mathbb{Q}$. We here use the (unusual) convention that $0^0 = 1$. Note that we can generate the signatures where all inner entries are zero by setting $a = 0$ or $b = 0$.

Now let $[\sigma_0, \sigma_1, \ldots, \sigma_n]$ be a symmetric signature over some basis b_0, b_1. We only deal with recognizers here, for generators, everything can be proven in a similar fashion (see also [CL07b]). We do a basis transform with $T = (t_i^j)$. Let $[\sigma'_0, \sigma'_1, \ldots, \sigma'_n]$ be the resulting signature. It is easy to see that symmetric signatures stay symmetric under such transforms. We can write

$$\sigma'_{k'} = \sum_{k=0}^{n} a_{k'}^k \sigma_k, \qquad 0 \le k' \le k$$

for appropriate scalars $a_{k'}^k$.

We can write the signature as

$$\sum_{i_1=0}^{1} \cdots \sum_{i_n=0}^{1} \underbrace{\tau_{i_1,\ldots,i_n}}_{=\sigma_{i_1+\cdots+i_n}} b_{i_1} \otimes \cdots \otimes b_{i_n}$$

The transformation T maps

$$b_i \mapsto t_0^i b_0 + t_1^i b_1, \qquad i = 0, 1.$$

Since we deal with recognisers, we have to take the transpose of T. We claim that

$$a_{k'}^k = \sum_{s=0}^{k} \binom{k'}{s} \binom{n-k'}{k-s} (t_1^1)^s (t_1^0)^{k'-s} (t_0^1)^{k-s} (t_0^0)^{n-k-k'+s}. \tag{6.3}$$

(Recall that $\binom{i}{j} = 0$, when $j \notin \{0, \ldots, i\}$.) $a_{k'}^k$ is the "contribution" of σ_k to $\sigma'_{k'}$ after transforming with T. We know that σ_k is the coefficient of any term $b_{i_1} \otimes \cdots \otimes b_{i_n}$ with $\sum_{j=1}^n i_j = k$. In the same way, $\sigma'_{k'}$ is the coefficient of any term $b_{i_1} \otimes \cdots \otimes b_{i_n}$ with $\sum_{j=1}^n i_j = k'$

To "reach" such a particular term with Hamming weight k' from any term with Hamming weight k, there are $\binom{k'}{s}$ choices to keeps s of the k' indices that are 1 and $\binom{n-k'}{k-s}$ choices to flip $k-s$ indices from 0 to 1. Then we have to flip $k-s$ indices from 1 to 0 and the remaining $n-k-k'+s$ indices stay 0.

We can rewrite (6.3) as

$$a^k_{k'} = (t^0_1)^{k'}(t^0_0)^{n-k'} \sum_{s=0}^{k} \binom{k'}{s}\binom{n-k'}{k-s}\left(\frac{t^1_1 t^0_0}{t^0_1 t^1_0}\right)^s \left(\frac{t^1_0}{t^0_0}\right)^k. \tag{6.4}$$

Theorem 6.2 *$[\sigma_0,\ldots,\sigma_k]$ is realisable under the basis $\beta = \{n,p\} = \{\begin{pmatrix} n_0 \\ n_1 \end{pmatrix}, \begin{pmatrix} p_0 \\ p_1 \end{pmatrix}\}$ as a recogniser iff it is of one of the following forms:*

1. $\sigma_i = \lambda((sn_0+tn_1)^{n-i}(sp_0+tp_1)^i + \epsilon(sn_0-tn_1)^{n-i}(sp_0-tp_1)^i)$
2. $\sigma_i = \lambda((n-i)n_0p_1^i n_1^{n-1-i} + ip_0p_1^{i-1}n_1^{n-i})$
3. $\sigma_i = \lambda((n-i)n_1p_0^i n_0^{n-1-i} + ip_1p_0^{i-1}n_0^{n-i})$

for rational numbers λ, s, and t, and $\epsilon = \pm 1$.

Proof. Let $T = (t^j_i) = \begin{pmatrix} n_0 & p_0 \\ n_1 & p_1 \end{pmatrix}$. Note that we take the transpose, since we deal with recognisers.

We distinguish several cases. The first one is that n is even and that the signature is even, too. By Theorem 6.1, we know that $\sigma_k = \lambda b^{n-k}c^k$ for all even k. Here we set $r_1 = b^2$ and $r_2 = c^2$. If necessary, we have to extend the ground field for this. By using (6.4), we get

$$\begin{aligned}
\sigma'_{k'} &= \lambda \sum_{k=0}^{n} \sigma_k a^k_{k'} \\
&= \lambda \sum_{k \text{ even}} b^{n-k}c^k a^k_{k'} \\
&= \lambda (t^0_1)^{k'}(t^0_0)^{n-k'} \sum_{k \text{ even}} b^{n-k}c^k \sum_{s=0}^{k} \binom{k'}{s}\binom{n-k'}{k-s}\left(\frac{t^1_1 t^0_0}{t^0_1 t^1_0}\right)^s \left(\frac{t^1_0}{t^0_0}\right)^k \\
&= \lambda (t^0_1)^{k'}(t^0_0)^{n-k'} \sum_{s=0}^{n} \binom{k'}{s} b^{k'-s}\left(\frac{ct^1_1}{t^0_1}\right)^s \left[\sum_{k \text{ even},\, k \geq s} \binom{n-k'}{k-s} b^{n-k'-k+s}\left(\frac{ct^1_0}{t^0_0}\right)^{k-s}\right].
\end{aligned}$$

We can write the expression in square brackets $[\ldots]$ as

$$\frac{1}{2}\left[\left(b+\frac{ct^1_0}{t^0_0}\right)^{n-k'} \pm \left(b-\frac{ct^1_0}{t^0_0}\right)^{n-k'}\right]$$

where the sign is a "+" when s is even and a "−" when s is odd. Call these two expressions $*_{\text{even}}$ and $*_{\text{odd}}$. We get

$$\begin{aligned}
\sigma'_{k'} &= \frac{1}{2}\lambda(t_1^0)^{k'}(t_0^0)^{n-k'}[*_{\text{even}}] \sum_{s \text{ even}} \binom{k'}{s} b^{k'-s} \left(\frac{ct_1^1}{t_1^0}\right)^s \\
&\quad - \frac{1}{2}\lambda(t_1^0)^{k'}(t_0^0)^{n-k'}[*_{\text{odd}}] \sum_{s \text{ odd}} \binom{k'}{s} b^{k'-s} \left(\frac{ct_1^1}{t_1^0}\right)^s \\
&= \frac{1}{2}\lambda(t_1^0)^{k'} \left[(bt_0^0 + ct_0^1)^{n-k'} + (bt_0^0 - ct_0^1)^{n-k'}\right] \\
&\quad \times \frac{1}{2}\left[\left(b + \frac{ct_1^1}{t_1^0}\right)^{k'} + \left(b - \frac{ct_1^1}{t_1^0}\right)^{k'}\right] \\
&\quad - \lambda(t_1^0)^{k'} \left[(bt_0^0 + ct_0^1)^{n-k'} - (bt_0^0 - ct_0^1)^{n-k'}\right] \\
&\quad \times \frac{1}{2}\left[\left(b + \frac{ct_1^1}{t_1^0}\right)^{k'} - \left(b - \frac{ct_1^1}{t_1^0}\right)^{k'}\right] \\
&= \frac{1}{2}\lambda\left[(bt_0^0 + ct_0^1)^{n-k'}(bt_1^0 + ct_1^1)^{k'} + (bt_0^0 - ct_0^1)^{n-k'}(bt_1^0 - ct_1^1)^{k'}\right].
\end{aligned}$$

Thus we reached the form as in the first item of the statement of the theorem.

In the second case, n is odd and the signature is even. In this case, we can write $\sigma_k = \lambda b^{n-1-k}c^k$ for all even k, similar to the first case. If $b \neq 0$, then we set $\lambda' := \lambda/b$ and we can proceed as in the first case.

If $b = 0$, then $\sigma_{n-1} = \lambda c^{n-1}$ and all other entries are 0. Set $\lambda' := \lambda c^{n-1}$. We have

$$\begin{aligned}
\sigma'_{k'} &= \sum_{k=0}^{n} \sigma_k a_{k'}^k \\
&= \sigma_{n-1} a_{k'}^{n-1} \\
&= \lambda'(n-k')(t_1^1)^{k'}(t_0^1)^{n-1-k'} + k'(t_1^1)^{k'-1}t_1^0(t_0^1)^{n-k'}.
\end{aligned}$$

This is the form in the second item of the statement.

The other two cases for odd signatures can be treated in a similar way. The reader is referred to the original work by Cai and Lu for more details [CL07b]. ∎

In a similar fashion, we can prove the same characterisation results for generators:

Theorem 6.3 *$[\sigma_0, \ldots, \sigma_k]$ is realisable under the basis $\beta = \{n, p\} = \{\begin{pmatrix} n_0 \\ n_1 \end{pmatrix}, \begin{pmatrix} p_0 \\ p_1 \end{pmatrix}\}$ as a generator iff it is of one of the following forms:*

1. $\sigma_i = \lambda((sp_1 - tp_0)^{n-i}(-sn_1 + tn_0)^i + \epsilon(sp_1 + tp_0)^{n-i}(-sn_1 - tn_0)^i)$

2. $\sigma_i = \lambda((n-i)p_1 n_0^i(-p_0)^{n-1-i} - in_1(n_0)^{i-1}(-p_0)^{n-i})$

3. $\sigma_i = \lambda(-(n-i)p_0(-n_1)^i(p_1)^{n-1-i} + in_0(-n_1)^{i-1}(p_1)^{n-i})$

for rational numbers λ*,* s*, and* t*, and* $\epsilon = \pm 1$*.*

We also need another characterisation of realizable symmetric signatures. We only prove it for recognisers. However, exactkly the same results hold for generators by a similar proof. We call the following cases in Theorem 6.2 *degenerate*:

- In Form 1, $sn_0 + tn_1 = 0$ or $sn_0 - tn_1 = 0$
- In Form 2, $n_1 = 0$
- In Form 3, $n_0 = 0$

In Form 1, if both $sn_0 + tn_1 = sn_0 - tn_1 = 0$, then the signature is

$$[0, 0, \ldots, 0, \lambda].$$

If only one of the terms is 0, then it looks like

$$[p, pq, pq^2, \ldots, pq^{n-1}, \lambda] \tag{6.5}$$

with arbitrary number p, q, and λ. Note that the second line subsumes the first.

In Form 2, the signatures take the form

$$[0, 0, \ldots, 0, \lambda_1, \lambda_2] \tag{6.6}$$

for arbitrary numbers λ_1 and λ_2. The same is true for Form 3.

If we are not in a degenerate case, then we can rewrite the sequence as

$$\sigma_i = A\alpha^i + B\beta^i$$

(Form 1) or as

$$\sigma_i = \alpha^i(Ai + B)$$

(Form 2 and 3). It is easy to check that these are solutions to *second-order homogeneous linear recurrences*, which are of the form $\sigma_i = a\sigma_{i-1} + b\sigma_{i-2}$ (see the appendix to this chapter).

Theorem 6.4 *A symmmetric signature* $[\sigma_0, \sigma_1, \ldots, \sigma_n]$ *is realisable (as a recogniser or generator) iff there exists constants* a*,* b*, and* c *not all equal to zero, such that for* $0 \leq k \leq n-2$*,*

$$a\sigma_k + b\sigma_{k+1} + c\sigma_{k+2} = 0$$

Proof. First assume that $[\sigma_0, \ldots, \sigma_n]$ is realisable. Then it takes one of the forms in Theorem 6.2 (or 6.3). If it is degenerate as in (6.5), then we can take $a = -q$, $b = 1$, and $c = 0$. If it is degenerate as in (6.6), then we set $a = 1$ and $b = c = 0$. In the nondegenerate cases, we have already seen above that the σ_i obey a second order recurrence relation.

For the converse direction, assume that $[\sigma_0, \ldots, \sigma_n]$ fulfills the second order linear recurrence in the statement of the theorem. If $c = b = 0$, then $a \neq 0$. In this case, $\sigma_i = 0$, $0 \leq i \leq n-2$. Thus the signature is of the form (6.6) and therefore realisable.

If $c = 0$ and $b \neq 0$, then $x_{k+1} = -a/b \cdot x_k$. We set $q := -a/b$ and conclude that the signature is of the form (6.5).

Otherwise $c \neq 0$. We get that $x_{k+2} = a_0 x_{k+1} + b_0 x_k$ with $a_0 = -b/c$ and $b_0 = -a/c$. The characteristic equation of the recurrence is $X^2 - a_0 X - b_0$. Let α and β be the two roots of this equation. If $\alpha \neq \beta$, then we can find A and B such that $\sigma_i = A\alpha^i + B\beta^i$, $0 \leq i \leq n$ (see the appendix). If $A = B = 0$, then $\sigma_i = 0$ for all i; thus the signature is trivially realisable. If $A = 0$ and $B \neq 0$, then $\sigma_i = B\beta_i$. We set $\epsilon = s = 1$, $t = 0$, $\lambda = B/2$, $(n_0, n_1) = (1, 0)$ and $(p_0, p_1) = (\beta, 1)$. It is easy to check that this is Form 1 of Theorem 6.2.

If $AB \neq 0$, then we set $\lambda = \epsilon = s = t = 1$ and get the following equations:

$$\begin{aligned} n_0 + n_1 &= \sqrt[n]{A} \\ n_0 - n_1 &= \sqrt[n]{B} \\ p_0 + p_1 &= \alpha \sqrt[n]{A} \\ p_0 - p_1 &= \beta \sqrt[n]{B} \end{aligned}$$

From the equations above, we get values of n_0, n_1, p_0, p_1 and we get that $\sigma_i = A\alpha^i + B\beta^i$ is realisable. (Note that $\alpha \neq \beta$, so we get two linearly independent vectors.)

If $\alpha = \beta$, then we can find A and B such that $\sigma_i = \alpha^i(Ai + B)$. We let $\lambda = 1$, $(n_0, n_1) = (B/n, 1)$, $(p_0, p_1) = (A\alpha + B\alpha/n, \alpha)$ and see that this signature is of Form 2. (If $A = 0$ or $\alpha = 0$, then the two vectors are linearly dependent. However, in these cases, the signature is easily seen to be realisable.) ∎

Corollary 6.5 *Over the complex numbers as well as all fields of characteristic $p > 3$, every signature $[\sigma_0, \ldots, \sigma_3]$ is realisable.*

Appendix: Second order homogeneous linear recurrences

Let $x_n = ax_{n-1}bx_{n-1}$, $n \in \mathbb{N}$, a second order homogeneous linear recurrence. We call

$$X^2 = aX + b$$

its *characteristic equation*. Let r_1 and r_2 be the roots of the characteristic equation, we call them the *characteristic roots*.

Lemma 6.6 *If r is a characteristic root, then (r^i) is a solution of the recurrence.*

Proof. We have $r^{i+2} = r^i(ar + b) = ar^{i+1} + br^i$ for all i. ∎

Lemma 6.7 *If r is a double root of the characteristic polynomial, then (ir^i) is a solution of the recurrence.*

Proof. In this case, $a^2 + 4b = 0$ and $r = a/2$. We have

$$\begin{aligned} a(i+1)r^{i+1} + bir^i &= a(i+1)\frac{a^{i+1}}{2^{i+1}} - \frac{a^2}{4}i\frac{a^i}{2^i} \\ &= (i+2)\frac{a^{i+2}}{2^{i+2}} \\ &= (i+2)r^{i+2}. \end{aligned}$$ ∎

The following lemma is obvious:

Lemma 6.8 *Any linear combination of solutions to the recurrence is again a solution.*

So if there are two distinct characteristic roots r_1, r_2, then all sequences of the form $Ar_1^i + Br_2^i$ are solutions. If there is one double root, then all sequences of the form $Air^i + Br^i$ are solutions.

These are the only solutions by the following observation: We know that

$$\begin{pmatrix} x_{i+2} \\ x_{i+1} \end{pmatrix} = \begin{pmatrix} a & b \\ 1 & 0 \end{pmatrix} \begin{pmatrix} x_{i+1} \\ x_i \end{pmatrix} = \cdots = \begin{pmatrix} a & b \\ 1 & 0 \end{pmatrix}^i \begin{pmatrix} x_1 \\ x_0 \end{pmatrix}$$

Thus the solution space is at most 2-dimensional. In particular, if we are given a concrete sequence, the values x_0 and x_1 determine the coefficients A and B above.

7 Simultaneously realisable bases

We start with a lemma due to Valiant [Val08].

Lemma 7.1 *If there is a recogniser (generator) realisable over the basis* $\{\begin{pmatrix} n_0 \\ p_0 \end{pmatrix}, \begin{pmatrix} n_1 \\ p_1 \end{pmatrix}\}$*, then there is a recogniser (generator) with the same signature over* $\{\begin{pmatrix} xn_0 \\ xp_0 \end{pmatrix}, \begin{pmatrix} yn_1 \\ yp_1 \end{pmatrix}\}$ *or* $\{\begin{pmatrix} xn_1 \\ xp_1 \end{pmatrix}, \begin{pmatrix} yn_0 \\ yp_0 \end{pmatrix}\}$ *for any* x, y *with* $xy \neq 0$.

Proof. To flip the two rows, we attach to every external node an extra edge. The other nodes of these edges become the new external nodes and the old nodes become internal nodes, see Figure 7.

To scale the rows, we add a path of length two to every external node. One edge of the path has weight x and the other one has weight y, see Figure 8. If new endnode of this path becomes the external node. If this node is externally matched, then we add an edge of weight x to the matching, if it is not externally matched, then we get an additional weight y. ∎

Definition 7.2 *We call two bases equivalent if they can be transformed into each other by exchanging the two rows and/or scaling the two rows with nonzero constants. We denote this equivalence relation by* $\sim$*. We let* $\mathcal{M} = F^{2\times 2}/\sim$ *be the set of equivalence classes, where* F *is the underlying ground field.*

Definition 7.3

1. $\mathcal{B}_{rec}([x_0, \ldots, x_n])$ *denotes the set of all bases over which* $[x_0, \ldots, x_n]$ *is realised as a recogniser (modulo* $\sim$*).*
2. $\mathcal{B}_{gen}([x_0, \ldots, x_n])$ *denotes the set of all bases over which* $[x_0, \ldots, x_n]$ *is realised as a generator (modulo* $\sim$*).*

Lemma 7.4 *Over fields of characteristic* $\neq 2$*, we have*

$$\mathcal{B}_{rec}(\lambda[a^n, a^{n-1}b, \ldots, b^n]) = \{\begin{pmatrix} a \\ n_1 \end{pmatrix}, \begin{pmatrix} b \\ p_1 \end{pmatrix}\}.$$

Fig. 7. Flipping two rows.

Fig. 8. Scaling with diagonal matrices.

Proof. First let $n = 1$. The standard signatures of arity 1 are $[\lambda, 0]$ (or $[0, \lambda]$). Over an arbitrary bases $\{\begin{pmatrix} n_0 \\ n_1 \end{pmatrix}, \begin{pmatrix} p_0 \\ p_1 \end{pmatrix}\}$, the standard signatures become $(\lambda n_0, \lambda p_0)$ (or $(\lambda n_1, \lambda p_1)$). Recall that we take the transpose, since we deal with recognisers. So we choose n_0 and p_0 such that $\lambda n_0 = a$ and $\lambda p_0 = b$. The only condition that we have is $ap_1 - bn_1 \neq 0$, that is, we have a basis.

Now, we assume that $n > 1$. By Theorem 6.2, we know how realisable signatures look like. We start with Form 1, that is,

$$x_i = \lambda[(\underbrace{sn_0 + tn_1}_{=:u_0})^{n-i}(\underbrace{sp_0 + tp_1}_{=:u_1})^i + \epsilon(\underbrace{sn_0 - tn_1}_{=:v_0})^{n-i}(\underbrace{sp_0 - tp_1}_{=:v_1})^i].$$

Thus

$$u_0^{n-i}u_1^i + \epsilon v_0^{n-i}v_1^i = a^{n-i}b^i.$$

We distinguish several cases.
Case 1: $u_0v_0 \neq 0$. Then

$$\begin{aligned} u_0^n + \epsilon v_0^n &= a^n \\ u_0^{n-2}u_1^2 + \epsilon v_0^{n-2}v_1^2 &= a^{n-2}b^2 \\ u_0^{n-1}u_1 + \epsilon v_0^{n-1}v_1 &= a^{n-1}b \end{aligned}$$

Therefore,

$$(u_0^n + \epsilon v_0^n)(u_0^{n-2}u_1^2 + \epsilon v_0^{n-2}v_1^2) = (u_0^{n-1}u_1 + \epsilon v_0^{n-1}v_1)^2.$$

We get

$$u_0^n v_0^{n-2}v_1^2 + u_0^{n-2}u_1^2 v_0^n = 2u_0^{n-1}u_1 v_0^{n-1}v_1.$$

Since $u_0v_0 \neq 0$, we can divide and get

$$\begin{aligned} u_0^2v_1^2 - 2u_0v_1u_1v_0 + u_1^2v_0^2 &= 0 \\ (u_0v_1 - u_1v_0) &= 0 \\ u_1/u_0 &= v_1/v_0 \end{aligned}$$

Set $\rho := u_1/u_0 = v_1/v_0$.

We have $x_{i+1} = \rho x_i$. This implies that $a \neq 0$ and $\rho = b/a$. (If $a = 0$, then the signature would be $[0, \ldots, 0, b^n]$. But $x_n = \rho x_{n-1}$ implies then $b = 0$.) Therfore, we know that

$$bu_0 = au_1 \qquad \wedge \qquad bv_0 = av_1,$$

that is,

$$b(sn_0 + tn_1) = a(sp_0 + tp_1) \qquad \wedge \qquad b(sn_0 - tn_1) = a(sp_0 - tp_1)$$

When we add or subtract both equations, we obtain

$$bsn_0 = asp_0 \qquad \wedge \qquad bt_1n_1 = atp_1.$$

We claim that $s = 0$ or $t = 0$. If $st \neq 0$, then dividing both equations yields $n_0/n_1 = p_0/p_1$, which means that we do not have basis, a contradiction.

Assume that $t = 0$, the case $s = 0$ is treated in the same way. We have

$$\begin{aligned} x_i &= \lambda[(sn_0)^{n-i}(sp_0)^i + \underbrace{\epsilon}_{=1}(sn_0)^{n-i}(sp_0)^i] \\ &= \lambda[2(\underbrace{sn_0}_{=a})^{n-i}(\underbrace{sp_0}_{=b})^i]. \end{aligned}$$

Therefore, all bases of the form $\{\begin{pmatrix} a \\ n_1 \end{pmatrix}, \begin{pmatrix} b \\ p_1 \end{pmatrix}\}$ will do.

Case 2: $u_1v_1 \neq 0$. This case is treated in the same way as the first.

Case 3: $u_0v_0 = 0 = u_1v_1$. We can assume w.l.o.g. that $u_0 = 0$. If $u_1 = 0$, then $s = t = 0$, since $\{\begin{pmatrix} n_0 \\ n_1 \end{pmatrix}, \begin{pmatrix} p_0 \\ p_1 \end{pmatrix}\}$ are linearly independent. Thus, $v_0 = v_1 = 0$, a contradiction. Hence $u_1 \neq 0$ and $v_1 = 0$. Therefore, our signature looks like

$$\lambda[\epsilon v_0^n, 0, \ldots, 0, u_1^n].$$

Since the inner entries are zero, we get that $ab = 0$. Since $u_1 \neq 0$, this means that $v_0 = 0$ and $u_1 = b$. It is easy to see that the statement of the lemma holds in this case.

Next, we come to Form 2 of Theorem 6.2. (Form 3 is treated in a similar way.) Here we have

$$x_i = \lambda((n-i)n_0p_1^in_1^{n-1-i} + ip_0p_1^{i-1}n_1^{n-i})$$

If $n_1 = 0$, then $a = 0$, since $x_0 = a^n = n \cdot n_0n_1^{n-1} = 0$. Then, $x_{n-1} = n_0p_1^{n-1} = a^{n-1}b = 0$. Thus, $n_0 = 0$ or $p_1 = 0$ and we have a singular "basis", a contradiction. Therefore, $n_1 \neq 0$. Then $a \neq 0$, too. (If not, then $x_1 = p_0n_1^{n-1} = a^{n-1}b = 0$. This means that $p_0 = 0$ and the equation above for x_0 yields $n_0 = 0$, a contradiction.)

Thus $a \neq 0$ and $n_0 \neq 0$. We have

$$\begin{aligned} a^n &= n \cdot n_0n_1^{n-1}, \\ a^{n-1}b &= n \cdot n_0p_1n_1^{n-2} - n_0p_1n_1^{n-2} + p_0n_1^{n-1} \\ &= n \cdot n_0p_1n_1^{n-2}\Bigg(\underbrace{\frac{p_1}{n_1}}_{=:\rho} + \underbrace{\frac{p_0}{n \cdot n_1} - \frac{p_1}{n \cdot n_1}}_{=\frac{n_1p_0 - n_0p_1}{n \cdot n_0n_1} =: c}\Bigg), \\ a^{n-2}b^2 &= n \cdot n_0n_1^{n-1}(2c\rho + \rho^2). \end{aligned}$$

Since $a^n \cdot a^{n-2}b^2 = (a^{n-1}b)^2$, we get that

$$(c+\rho)^2 = 2c\rho + \rho^2,$$

which in turn means that $c = 0$ and $n_1p_0 - n_0p_1 = 0$, a contradiction. So Form 2 is not of any use in this case. ∎

Definition 7.5 *We call* $[x_0, \ldots, x_n]$ nondegenerate *if* $\mathrm{rk}\begin{pmatrix} x_0 & \ldots & x_{n-1} \\ x_1 & \ldots & x_n \end{pmatrix} = 2$. *Otherwise, it is* degenerate.

If the rank of the matrix in the definition above is 0, then $[x_0, \ldots, x_n] = [0, \ldots, 0]$. If it is 1, then $[x_0, \ldots, x_n] = \lambda[a^n, a^{n-1}b, \ldots, b^n]$. This case has been dealt with in the previous lemma. Now, we deal with the case of rank 2.

Lemma 7.6 $\mathcal{B}_{rec}([x_0, x_1, x_2])$ *is the set of all bases* $\{\begin{pmatrix} n_0 \\ n_1 \end{pmatrix}, \begin{pmatrix} p_0 \\ p_1 \end{pmatrix}\} \in \mathcal{M}$ *such that* $x_0p_1^2 - 2x_1p_1n_1 + x_2n_1^2 = x_0p_0^2 - 2x_1p_0n_0 + x_2^2n_0^2 = 0$ *or* $x_0p_0p_1 - x_1(n_0p_1 + n_1p0) + x_2n_0n_1 = 0$.

Proof. The statement directly follows from the fact that for a standard signature of arity 2, the only constraints are the parity constraints. ∎

Now recall Theorem 6.4: $[x_0, \ldots, x_n]$ is realisable iff there are a, b, c such that for all $0 \leq k \leq n-2$,

$$ax_k + bx_{k+1} + cx_{k+2} = 0.$$

Assume that $n \geq 3$ and that the signature is nondegenerate. In this case, (a, b, c) is unique up to scaling. This is because every such vector lies in the left kernel of the matrix

$$\begin{pmatrix} x_0 & x_1 & \ldots & x_{n-2} \\ x_1 & x_2 & \ldots & x_{n-1} \\ x_2 & x_3 & \ldots & x_n \end{pmatrix}$$

and we know that the signature is nondegenerate.

Lemma 7.7 *Let* $AB \neq 0$ *and* $\alpha \neq \beta$. *Then* $\mathcal{B}_{rec}([A\alpha^i + B\beta^i \mid i = 0, \ldots, n] = \{\begin{pmatrix} 1+\omega \\ 1-\omega \end{pmatrix}, \begin{pmatrix} \alpha+\beta\omega \\ \alpha-\beta\omega \end{pmatrix}\}$, *where* $\omega^n = \pm B/A$.

Proof. We have $x_0 = A + B$ and $x_1 = A\alpha + B\beta$, that is,

$$\begin{pmatrix} x_0 \\ x_1 \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ \alpha & \beta \end{pmatrix} \begin{pmatrix} A \\ B \end{pmatrix}.$$

Since $\alpha \neq \beta$, A and B are unique. Furthermore, since $\alpha \neq \beta$, we are in Form 1, that is,

$$A\alpha^i + B\beta^i = u_0^{n-i}u_1^i + \epsilon v_0^{n-i}v_1^i$$

where

$$\begin{aligned} u_0 &= sn_0 + tn_1, \\ u_1 &= sp_0 + tp_1, \\ v_0 &= sn_0 - tn_1, \\ v_1 &= sp_0 - tp_1. \end{aligned}$$

We have $u_0 \neq 0$. Otherwise

$$\underbrace{(A+B)}_{=\epsilon v_0^n v_1}\underbrace{(A\alpha^2 + B\beta^2)}_{=\epsilon v_0^{n-2}v_1^2} = \underbrace{(A\alpha + B\beta)^2}_{=(\epsilon v_0^{n-2}v_1)^2}$$

Since $AB \neq 0$, this implies $\alpha = \beta$, a contradiction. In the same way, we get that $v_0 \neq 0$. Thus we can write

$$x_i = A\alpha^i + B\beta^i = u_0^n \left(\frac{u_1}{u_0}\right)^i + \epsilon v_0^n \left(\frac{v_1}{v_0}\right)^i$$

It is easy to see that any representation of x_i is unique (see the appendix to the previous chapter). Thus,

$$\begin{aligned} u_0^n &= A, \\ \epsilon v_0^n &= B, \\ u_1/u_0 &= \alpha, \\ v_1/v_0 &= \beta. \end{aligned}$$

We have

$$\{\begin{pmatrix} 2sn_0 \\ 2tn_1 \end{pmatrix}, \begin{pmatrix} 2sp_0 \\ 2tp_1 \end{pmatrix}\} = \{\begin{pmatrix} u_0 + v_0 \\ u_0 - v_0 \end{pmatrix}, \begin{pmatrix} u_1 + v_1 \\ u_1 - v_1 \end{pmatrix}\}.$$

Since $\alpha \neq \beta$, we get $s \neq t$. (This is proven in the same way as $u_0 \neq 0$.) Let $\omega = v_0/u_0$. Then $\omega^n = \pm B/A$. We have

$$\{\begin{pmatrix} n_0 \\ n_1 \end{pmatrix}, \begin{pmatrix} p_0 \\ p_1 \end{pmatrix}\} \sim \{\begin{pmatrix} 2sn_0 \\ 2tn_1 \end{pmatrix}, \begin{pmatrix} 2sp_0 \\ 2tp_1 \end{pmatrix}\} \sim \{\begin{pmatrix} 1+\omega \\ 1-\omega \end{pmatrix}, \begin{pmatrix} \alpha+\beta\omega \\ \alpha-\beta\omega \end{pmatrix}\}.$$

■

Corollary 7.8 *For* $\alpha \neq 0$, $\mathcal{B}_{rec}([A, A\alpha, A\alpha^2, \ldots, A\alpha^n + B]) = \{\begin{pmatrix} 1/\alpha \\ 1/\alpha \end{pmatrix}, \begin{pmatrix} 1+\omega \\ 1-\omega \end{pmatrix} \mid \omega^n = \pm B/A\}$.

Proof. The reversed signature $[A\alpha^n + B, A\alpha^{n-1}, \ldots, A]$ is a special case of the lemma, we set $\beta = 1$ and replace α by $1/\alpha$. ∎

When the characteristic roots are equal, that is, $\alpha = \beta$, then we have the following lemma.

Lemma 7.9 *Let p be the characteristic of F and let $A \neq 0$.*

1. *When $p = 0$ or p does not divide n, then $\mathcal{B}_{rec}([Ai\alpha^{i-1} + B\alpha^i \mid i = 0, \ldots, n] = \{\begin{pmatrix} 1 \\ B \end{pmatrix}, \begin{pmatrix} \alpha \\ nA + B\alpha \end{pmatrix}\}$.*
2. *When $p|n$ and $x_0 = 0$, then $B = 0$, that is, $x_i = Ai\alpha^{i-1}$, and $\mathcal{B}_{rec}([Ai\alpha^{i-1} \mid i = 0, \ldots, n] = \{\begin{pmatrix} 1 \\ n_1 \end{pmatrix}, \begin{pmatrix} \alpha \\ p_1 \end{pmatrix} \in \mathcal{M} \mid n_1, p_1 \in F\}$.*
3. *When $p|n$ and $x_0 \neq 0$, then $[Ai\alpha^{i-1} + B\alpha^i]$ is not realisable.*

Proof. In the first case, from $B = x_0$ and $A + B\alpha = x_1$, we can uniquely solve for A and B, since this is an systems of two linear equations. Now consider Form 2 of Theorem 6.2. (Form 3 will give the same bases.) We claim that $n_1 \neq 0$: Otherwise $x_i = 0$, $i = 0, \ldots, n-2$. Since $n \geq 3$, this would give $A = 0$, a contradiction.

Form 2 looks like

$$x_i = i(n_1p_0 - n_0p_1)n_1^n \left(\frac{p_1}{n_1}\right)^{i-1} + nn_0n_1^{n-1}\left(\frac{p_1}{n_1}\right)^i.$$

Since the representation is unique, we get that $(n_1p_0 - n_0p_1)n_1^n = A$, $p_1/n_1 = \alpha$ and $nn_0n_1^{n-1} = B$. Since $n_1 \neq 0$, we can scale $n_1 = 1$. Then $n_0 = B/n$, $p_1 = \alpha$ and $p_0 = A + B\alpha/n$. Now the results follows by flipping the rows of the basis.

In the second case, $x_0 = B$ yields $B = 0$. Thus $x_i = Ai\alpha^{i-1}$. If we now proceed as in the first case (using the fact that $p|n$), we can choose $n_1 = 1$ and $p_1 = \alpha$. Then the expression for x_i has the form $Ai\alpha^{i-1}$. The condition on n_0 and p_0 is just $n_1p_0 - n_0p_1 = 1$, that is, we have a basis. Again the results follows by flipping rows.

In the third case, $B \neq 0$, but the expression we get for B like in case 1 is $nn_0n_1^{n-1}$, which is zero. ∎

We shall remark that all these results can be proven for generators in a similar fashion, see [CL07a].

Now let us apply these results to our example Pl-Rtw-Mon-k-SAT. By Lemma **??**, we get that

$$\mathcal{B}_{rec}([0, 1, \ldots, 1]) = \left\{\begin{pmatrix} 1+\omega \\ 1-\omega \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \end{pmatrix} \mid \omega^k = \pm 1\right\}.$$

(We set $A = 1$, $B = -1$, $\alpha = 1$, $\beta = 0$.)

By Lemma 7.6, such a basis is in $\mathcal{B}_{gen}([1,0,1])$ (proven for generators) if $(1+\omega)^2 + 1 = (1-\omega)^2 + 1 = 0$ or $(1+\omega)(1-\omega) + 1 = 0$. It is easy to see that the first possibility has no solutions. The second possibility gives

$$\omega^2 = 2.$$

Since $\omega^k = \pm 1$, we get that $\omega^{2k} = 1$, thus $2^k - 1 = 0$. Therefore, if we want to count solutions to Pl-Rtw-Mon-k-SAT modulo p, p needs to divide $2^k - 1$.

Theorem 7.10 *Pl-Rtw-Mon-k-SAT can be solved in polynomial time when counting modulo* $2^k - 1$. *Every modulus* m *for which such a holographic algorithm exists must divide* $2^k - 1$.

Proof. The necessity for m has been demonstrated above. So we have to show that we can indeed realise the two signatures modulo $2^k - 1$.

We first assume that l is even. $[1+\epsilon 2^{k/2}, 1, \ldots, 1]$ and $[1,0,1]$ are realisable over the basis

$$\mathcal{B}_{rec}([0,1,\ldots,1]) = \left\{ \begin{pmatrix} 1+\sqrt{2} \\ 1-\sqrt{2} \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \end{pmatrix} \mid \omega^k = \pm 1 \right\}.$$

by Lemma 7.7 (set $A = 1$, $B = \pm 2^{n/2}$, $\alpha = 1$, $\beta = 0$) and Lemma 7.6 When we set $\epsilon = 1$, then we realise our signatures modulo $2^{k/2} + 1$. When we set $\epsilon = -1$, then we realise them modulo $2^{k/2} - 1$. So we can count modulo $2^{k/2} + 1$ and $2^{k/2} - 1$ using the FKT algorithm over $\mathbb{Q}(\sqrt{2})$. Using the Chinese remainder theorem, we can get the number of solutions modulo $2^k - 1 = (2^{k/2} + 1)(2^{k/2} - 1)$.

If k is odd, then set $r = 2^{(k+1)/2}$. We have $r^2 = 2 \mod 2^k - 1$ and $1 - r^k = 1 - (2^k)^{(k+1)/2} = 0 \mod 2^k - 1$. Thus we can take the basis above for $\epsilon = -1$ and replacing $\sqrt{2}$ by r. ∎

Bibliography

[CG14] Jin--Yi Cai and Aaron Gorenstein. Matchgates revisited. *Theory of Computing*, 10:167–197, 2014.

[CL07a] Jin-Yi Cai and Pinyan Lu. Holographic algorithms: From an art to a science. In *Proc. 39th Ann. ACM. Symp. on Theory of Comput. (STOC)*, pages 401–410, 2007.

[CL07b] Jin-Yi Cai and Pinyan Lu. On symmetric signatures in holographic algorithms. In *24th Annual Symposium on Theoretical Aspects of Computer Science (STACS 2007)*, volume 4393 of *Lecture Notes in Computer Science*, pages 429–440, 2007.

[CLX08] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Holographic algorithms by Fibonacci gates and holographic reductions for hardness. In *Proc. 49th Ann. IEEE Symp. on Foundations of Comput. Sci. (FOCS)*, pages 644–653, 2008.

[Cur15] Radu Curticapean. *The simple, little, and slow things count: On parameterized counting complexity.* PhD thesis, Saarland University, 2015.

[Die05] Reinhard Diestel. *Graph Theory.* Springer, 2005.

[Edm65] J. Edmonds. Paths, trees, and flowers. *Canad. J. Math.*, 17:449–467, 1965.

[Kas61] P. W. Kasteleyn. The statistics of dimers on a lattice. i. the number of dimer arrangements on a quadratic lattice. *Physica*, 27(12):1209–1225, 1961.

[Kas67] P. W. Kasteleyn. Graph theory and crystal physics. In F. Harary, editor, *Graph Theory and Theoretical Physics*, pages 43–110. Academic Press, 1967.

[TF61] H. N. V. Temperley and Michael E. Fisher. Dimer problem in statistical mechanics—an exact result. *Philosophical Magazine*, 6(68):1061–1063, 1961.

[Val79] Leslie G. Valiant. The complexity of computing the permanent. *Theoret. Comput. Sci.*, 8(2):189–201, 1979.

[Val02] Leslie Valiant. Expressiveness of matchgates. *Theoret. Comput. Sci.*, 289(1):457–471, 2002.

[Val08] Leslie G. Valiant. Holographic algorithms. *SIAM J. Comput*, 37(5):1565–1594, 2008.

[Zan91] Viktória Zankó. #P-completeness via many-one-reductions. *Int. J. of Found. Comput. Sci.*, 2:77–82, 1991.