GoLA EXPOSITORY: THE PFAFFIAN AND ITS COMBINATORICS

ETHAN ZHOU

1. The Pfaffian

Definition 1.1. A matrix $A = (a_{ij})$ is *skew-symmetric* iff $A = -A^T$, i.e. all $a_{ij} = -a_{ji}$.

One popular application of a skew-symmetric matrix is that the cross product $\mathbf{a} \times \mathbf{b}$ of two vectors $\mathbf{a}$ and $\mathbf{b}$ can be represented as a matrix multiplication $\mathbf{a}_\times \mathbf{b}$ where

$$\mathbf{a}_\times = \begin{pmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{pmatrix}.$$

We are more interested in the determinants and combinatorics. Notice that an $n \times n$ skew-symmetric matrix A satisfies $\det(A) = \det(A^T) = \det(-A) = (-1)^n \det(A)$. Therefore, if n is odd, $\det(A) = -\det(A) = 0$. We narrow in on matrices with even order $2n$. For skew-symmetric matrices of even order is defined a special polynomial that is similar to the determinant.

Definition 1.2. Let $A = (a_{ij})$ be a $2n \times 2n$ skew-symmetric matrix. Its *Pfaffian* is

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

For comparison,

$$\det(A) = \sum_{\sigma \in S_{2n}} \mathrm{sgn}(\sigma) \prod_{k=1}^{2n} a_{k,\sigma(k)}.$$

The Pfaffian's definition as it relates to permutations is about pairing elements together, like $\{(\sigma(1), \sigma(2)), \ldots, (\sigma(2n-1), \sigma(2n))\}$; multiplying their corresponding matrix entries, the $a_{\sigma(2k-1),\sigma(2k)}$ to a product for that permutation; then accumulating with sign these products of the permutations. We divide by 2^n, because in swapping the order of the elements $\sigma(2k-1)$ and $\sigma(2k)$, i.e. those in any one of the n pairs, results in a duplicate term: Calling that permutation with one swap σ', we have $\mathrm{sgn}(\sigma') = -\mathrm{sgn}(\sigma)$ and $a_{\sigma(2k),\sigma(2k-1)} = -a_{\sigma(2k-1),\sigma(2k)}$, so the resulting term from σ' in the sum is $(-1)^2 = 1$ times the resulting term from σ, meaning it is the same. Then, we also divide in the Pfaffian by $n!$ to eliminate duplicates obtained by permuting the same set of n pairings. Notice that swapping two pairings does not change the sign of the corresponding permutation, as it involves two swappings of two elements. If we wish to avoid the duplicates from the get-go, we may prefer limiting the permutations. Let P_{2n} be the set of permutations

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

where $\{\{i_1, j_1\}, \{i_2, j_2\}, \ldots, \{i_n, j_n\}\}$ is a partition of $\{1, 2, \ldots, 2n\}$ into pairs (i_α, j_α) so that i is an increasing sequence and $i_\alpha < j_\alpha$ for $1 \leq \alpha \leq n$. Essentially, this is restricting to one possible order the otherwise ($2^n n!$) permutations that are equivalent in the calculation Pfaffian. Now, $i_\alpha = \pi(2\alpha - 1)$ and $j_\alpha = \pi(2\alpha)$. Then,

$$\mathrm{pf}(A) = \sum_{\pi \in P_{2n}} \mathrm{sgn}(\pi) \prod_{\alpha=1}^{n} a_{i_\alpha, j_\alpha}.$$

Theorem 1.3. *Some properties of the Pfaffian of a $2n \times 2n$ skew-symmetric matrix A are the following:*

(1) $\mathrm{pf}(A)^2 = \det(A)$. *This is Cayley's Theorem on Pfaffians.*
(2) Adding a multiple of a row k and column k to row l column l respectively does not change the Pfaffian.
(3) Swapping rows k and l and columns k and l of A changes the sign of the Pfaffian.
(4) Multiplying row k and column k by factor scales the Pfaffian by the same amount.
(5) For a matrix B of equal order, $\mathrm{pf}(B^T AB) = \det(B)\,\mathrm{pf}(A)$.

These allow quicker calculations of the Pfaffian. The most important is the first.

Proof of Cayley's Theorem. Any permutation may be decomposed into disjoint cycles. Consider a term σ whose said decomposition involves a cycle of odd length. Then, consider permutation σ_f that is σ but with that cycle of odd length in reverse. If there is more than one cycle of odd length, arbitrarily pick one. For example, if $\sigma = (138)(25467)$, then we might say $\sigma_f = (138)(76452)$ (and there are two more that are $(831)(25467)$ and $(831)(76452)$ that pair with each other). So we write

$$\begin{cases} \sigma = \tau\sigma' \\ \sigma_f = \tau^{-1}\sigma', \end{cases}$$

where τ is a cycle of odd length and σ' is the rest of σ. Now, consider the σ and σ_f terms in the determinant polynomial. Since both cycles are of the same length, $\mathrm{sgn}(\sigma) = \mathrm{sgn}(\sigma_r)$. For all elements $k \notin \tau$, we have $a_{k,\sigma(k)} = a_{k,\sigma_f(k)}$. Now, all we have to compare are the factors for elements in τ, the $a_{\tau(k),\tau(k+1)}$. Let L be τ's odd length, and let $\tau(i) = \tau(i \bmod L)$ for simplicity. Remember that $a_{ij} = -a_{ji}$, so

$$\begin{aligned} \prod_{k=1}^{L} a_{\tau(k),\tau(k+1))} &= (-1)^L \prod_{k=1}^{L} a_{\tau(k+1),\tau(k)} \\ &= -\prod_{k=1}^{L} a_{\tau^{-1}(k),\tau^{-1}(k+1)}. \end{aligned}$$

In our example,

$$a_{25}a_{54}a_{46}a_{67}a_{72} = (-1)^5 a_{52}a_{45}a_{64}a_{76}a_{27} = -a_{76}a_{64}a_{45}a_{52}a_{27}.$$

Thus, we see that

$$\mathrm{sgn}(\sigma_f) \prod_{k=1}^{2n} a_{k,\sigma_f(k)} = -\,\mathrm{sgn}(\sigma) \prod_{k=1}^{2n} a_{k,\sigma(k)},$$

i.e. the σ and σ_f terms are opposite. This means all terms of permutations that induce a cycle with odd length cancel in the determinant of a skew-symmetric matrix.

Then, we must consider the rest of the terms. These correspond to permutations composed only of cycles with even length. Now recall the permutations π that correspond to the

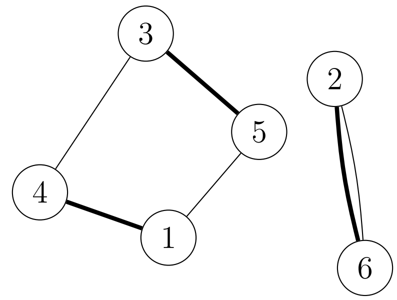

Figure 1. Two partitions into pairs, namely $\{\{1,4\},\{3,5\},\{2,6\}\}$ (thick) and $\{\{1,5\},\{3,4\},\{2,6\}\}$ (thin), together forming cycles (1435) and (26)

partitions of $\{1,\ldots,2n\}$ into pairs $\{i_\alpha, j_\alpha\}$. We take any two such permutations; call them π_P and π_Q for partitions P and Q. For any pair in one partition, say P, there is unique pair in the other partition Q that includes the first element and another pair in Q that includes the second; each element in the pairs of Q also appears in a unique pair in P. Going back and forth between P and Q, we string together a chain of elements that together make a cycle. One example is illustrated in Figure 1. All the cycles together make a permutation, denoted σ_{PQ}. Essentially, each of the two partitions into pairs P and Q take every other factor of each cycle in σ_{PQ}. For instance, in the figure, the corresponding term of the determinant is

$$\begin{aligned} a_{14}a_{26}a_{35}a_{43}a_{51}a_{62} &= (a_{14}a_{43}a_{35}a_{51})(a_{26}a_{62}) \\ &= (a_{14}a_{35}a_{26})(a_{43}a_{51}a_{62}) \\ &= -(a_{14}a_{26}a_{35})(a_{15}a_{26}a_{34}) \end{aligned}$$

where the second equality groups the factors into cycles, where the third equality groups the factors by partition, and where the fourth equality straightens the partitions so the indices are in order according the permutations associated with the partitions, the minus sign coming from multiple applications of $a_{ij} = -a_{ji}$.

Since these cycles must alternate between partitions, their lengths are all even. We therefore see a one-to-one correspondence between all such ordered pairs (P, Q) and permutations composed of even cycles. Thus, we also see a one-to-one correspondence between all ordered pairs of nonzero terms in the Pfaffian A and all nonzero terms in the determinant of A. In one term of $\det(A)$, each index appears exactly once as the row index and exactly once as a column index. In $\mathrm{pf}(A)^2$, between two partitions, each index appears exactly twice and not in the same pair; we reverse the order of certain pairs so each index appears exactly once as a row index and column index. The corresponding terms are then also equal up to sign.

$$\left|\prod_{i=1}^{2n} a_{i,\sigma_{PQ}(i)}\right| = \left|\prod_{(i,j)\in P\cup Q} a_{ij}\right|$$

Now, we will show the signs of the terms are equal. Consider a partition P' that is P with exactly two elements swapped between two pairs, so for two pairs $(\{i_\alpha, j_\alpha\}, \{i_\beta, j_\beta\})$ in P, we instead have in P' the pairs $(\{i_\alpha, j_\beta\}, \{i_\beta, j_\alpha\})$ or the pairs $(\{i_\alpha, i_\beta\}, \{j_\alpha, j_\beta\})$. Every permutation is composition of transpositions on the identity permutation, so if we can show that a transposition does not change the sign between the permutation's corresponding terms

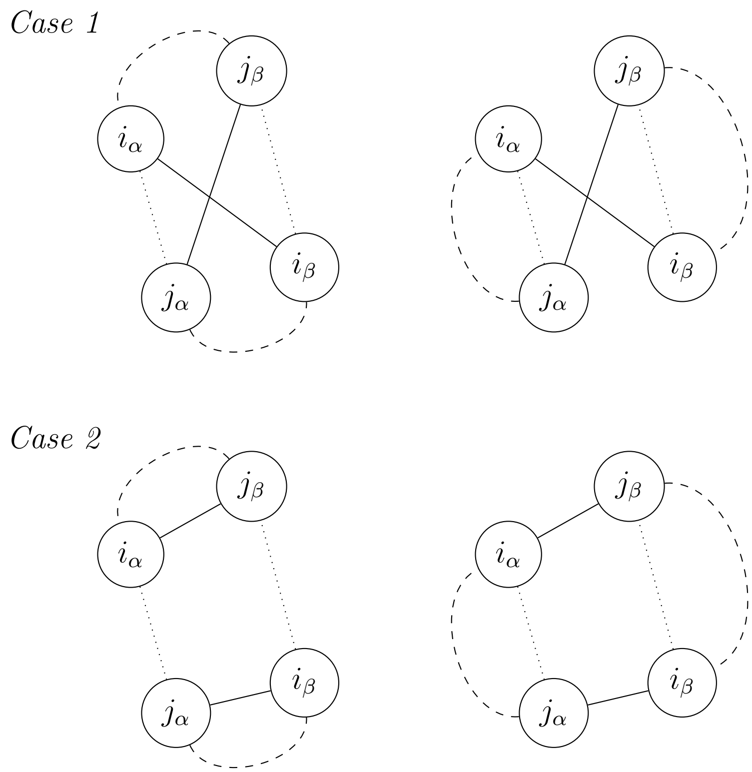

Figure 2. Cases of transpositions, where dotted edges are the removed connections, solid edges are the new connections, and dashed edges are the remaining edges and elements in the affected cycles.

in $\det(A)$ and $\mathrm{pf}(A)^2$, then transitively, we show that all corresponding terms have the same sign.

Case 1: $(\{i_\alpha, i_\beta\}, \{j_\alpha, j_\beta\})$. If the original pairs $(\{i_\alpha, j_\alpha\}, \{i_\beta, j_\beta\})$ are of the same cycle, then $(\{i_\alpha, j_\beta\}, \{i_\beta, j_\alpha\})$ are also both of a cycle of the same length. In the first subcase, without loss of generality[1], if the cycle of some length L in the original permutation π_P was

$$(i_1 j_1 i_2 j_2 \dots i_{\alpha-1} j_{\alpha-1} i_\alpha j_\alpha i_{\alpha+1} j_{\alpha+1} \dots i_{\beta-1} j_{\beta-1} i_\beta j_\beta i_{\beta+1} j_{\beta+1} \dots i_L j_L),$$

then the new cycle in $\pi_{P'}$ is something like

$$(i_1 j_1 i_2 j_2 \dots i_{\alpha-1} j_{\alpha-1} i_\alpha i_\beta j_{\beta-1} i_{\beta-1} \dots j_{\alpha+1} i_{\alpha+1} j_\alpha j_\beta i_{\beta+1} j_{\beta+1} \dots i_L j_L),$$

which is the first cycle but with elements between j_α and i_β inclusively in reverse order. We transposed two elements of π_P, so $\mathrm{sgn}(\pi_{P'}) = -\mathrm{sgn}(\pi_P)$. In the decomposition of σ_{PQ}, all cycles are still of the same length, so $\mathrm{sgn}(\sigma_{P'Q}) = \mathrm{sgn}(\sigma_{PQ})$, but an odd number of "edges" in the cycle reversed direction, which must be rectified on the $\mathrm{pf}(A)^2$ side of the equation (by making $a_{ij} = -a_{ji}$ substitutions) to keep the first indices less than second indices; thus, in total we multiply by -1 once for the transposition and then another odd number of times, resulting in multiplication by 1, so the signs of the corresponding terms in $\det(A)$ and $\mathrm{pf}(A)^2$ are both unchanged.

[1]The indices of the following i and j are not necessarily $1, 2, 3, \dots, L$ in order like this, but some ordered subset of the elements 1 through n. Generally, the indices are $\omega(1), \omega(2), \omega(3), \dots, \omega(L)$ for some permutation ω on the elements included in the cycle.

Otherwise, if the original pairs are of different cycles, then the new pairs join those cycles, so we lose two cycles and gain one cycle with all these cycles being of even length, so $\operatorname{sgn}(\sigma_{P'Q}) = (-1)^{-2+1}\operatorname{sgn}(\sigma_{PQ}) = -\operatorname{sgn}(\sigma_{PQ})$. However, this time, the number of "edges" that reverse direction is even, so we multiply the $\operatorname{pf}(A)^2$ by -1 once and then an even number of times. So, both corresponding terms' signs are flipped. In the first case, the relationship between the terms' signs is preserved between terms.

Case 2: $(\{i_\alpha, j_\beta\}, \{i_\alpha, j_\beta\})$. If the two original pairs $(\{i_\alpha, j_\alpha\}, \{i_\beta, j_\beta\})$ are of the same cycle, then $(\{i_\alpha, j_\beta\}, \{i_\beta, j_\alpha\})$ are of two different, disjoint cycles. So, $\operatorname{sgn}(\sigma_{P'Q}) = -\operatorname{sgn}(\sigma_{PQ})$. Both new cycles are even in length, so their directions of traversal do not matter and we do not have to worry about other factors of -1, though we still have $\operatorname{sgn}(\pi_{P'}) = -\operatorname{sgn}(\pi_P)$ from the transposition. The signs of the corresponding terms in $\det(A)$ and $\operatorname{pf}(A)^2$ are both flipped.

Otherwise, if the original pairs are of different cycles, then the new pairs join them into one cycle. Similarly, $\operatorname{sgn}(\sigma_{P'Q}) = (-1)^{-2+1}\operatorname{sgn}(\sigma_{PQ}) = -\operatorname{sgn}(\sigma_{PQ})$, and $\operatorname{sgn}(\pi_{P'}) = -\operatorname{sgn}(\pi_P)$. There are no reversed edges. So, both corresponding terms' signs are flipped.

We see no matter which transposition we make, the resulting corresponding terms of the have the same sign. Since corresponding terms of $\det(A)$ and $\operatorname{pf}(A)^2$ completely identical, we have $\operatorname{pf}(A)^2 = \det(A)$.

■

2. Perfect Matchings

Definition 2.1. A *matching* of a graph is a subset of the edges such that each vertex only appears in one edge of the subset. A *perfect matching* is a matching such that each vertex appears exactly once.

We find Pfaffians beneficial in counting the number of perfect matchings of a graph. For Pfaffian graphs, the calculation becomes polynomial time.

Definition 2.2. An *adjacency matrix* $A = (a_{ij})_{\substack{1\le i\le n\\ 1\le j\le n}}$ of a graph $G = (V, E)$ is a matrix representation keep track of which vertices are adjacent, defined, given some ordering of the graph's n vertices, with

$$a_{ij} = \begin{cases} 1 & v_i v_j \in E, \\ 0 & \text{otherwise.} \end{cases}$$

This matrix is symmetric, but we find a skew-symmetric matrix from instead defining a *directed adjacency matrix* $D = (d_{ij})_{\substack{1\le i\le n\\ 1\le j\le n}}$ so

$$d_{ij} = \begin{cases} 1 & v_i v_j \in E, \\ -1 & v_i v_j \in E, \\ 0 & \text{otherwise.} \end{cases}$$

Not every graph is directed, but an *orientation* assigns of direction to each of its edges.

Consider the Pfaffian of the directed adjacency matrix D of graph G. Immediately, we see the Pfaffian of an odd-order skew-symmetric matrix is 0, and the number of perfect matchings of an odd-order graph is also 0. In even-order cases, a perfect matching is a pairing of the vertices in G as are paired the entries of D in $\mathrm{pf}(D)$. Iff the permutation involves an edge $\{\pi(i), \pi(j)\} \notin E(G)$, then $a_{\pi(i),\pi(j)} = 0$ makes the whole term

$$\mathrm{sgn}(\pi) \prod_{k=1}^{n} a_{i_k, j_k} = 0.$$

Thus, there is exactly one nonzero ± 1 term for each perfect matching of G, since it sums over all distinct pairings. If we pick a good orientation, then all terms will have the same sign, and the Pfaffian will have a magnitude counting the number of perfect matchings of G.

Definition 2.3. A *central cycle* C is a cycle such that $G \setminus C$ has a perfect matching, where $G \setminus C$ denotes the graph G with the vertices of C removed. An *oddly oriented cycle* is a cycle such that there is an odd count of edges oriented in either direction. A *Pfaffian orientation* is an orientation such that all central cycles of even length are oddly oriented.

Proposition 2.4. *The Pfaffian of the directed adjacency matrix D of a graph G with Pfaffian orientation counts the number of perfect matchings.*

Proof. We already know that $\mathrm{pf}(D)$ has the correct number of terms, but we want to show that all terms have the same sign.

Take a central cycle C of even length e. By the definition of central, $G \setminus C$ has a perfect matching; therefore, a perfect matching of G may be the union of a perfect matching of C (remember C is even in length) and a perfect matching of $G \setminus C$. There are two disjoint perfect matchings M_1 and M_2 of C. Notice that each matching takes every other edge of C, so $M_1 \cap M_2 = \emptyset$ and $M_1 \cup M_2 = E(C)$. Remember the set of permutations P defined earlier. The two matchings correspond to the permutations

$$\mu_1 = \begin{pmatrix} 1 & 2 & 3 & 4 & \ldots & e-1 & e \end{pmatrix}$$
$$\mu_2 = \begin{pmatrix} 1 & e & 2 & 3 & \ldots & e-2 & e-1 \end{pmatrix}.$$

We want to compare their terms in the Pfaffian.

Notice μ_2 is an cycle permutation on the last $e-1$ elements, so

$$\mathrm{sgn}(\mu_2) = (-1)^{(e-1)-1} = (-1)^e = 1.$$

Therefore, $\mathrm{sgn}(\mu_1) = \mathrm{sgn}(\mu_2) = 1$. What remains to be compared is the products of matrix entries of D accessed by the (i_α, j_α) obtained from permutations, explicitly

$$y_1 = \prod_{\alpha=1}^{\frac{1}{2}e} a_{\mu_1(2\alpha-1), \mu_1(2\alpha)} = d_{1,2} d_{3,4} \cdots d_{e-1,e}$$
$$y_2 = \prod_{\alpha=1}^{\frac{1}{2}e} a_{\mu_2(2\alpha-1), \mu_2(2\alpha)} = d_{1,e} d_{2,3} \cdots d_{e-2,e-1}$$

for the respective μ. All referenced entries but $d_{1,e}$ lie on the superdiagonal, and each product takes every other entry. From definition, $d_{i_\alpha, j_\alpha} = -1$ iff there is an edge oriented from i_α to j_α, and $d_{i_\alpha, j_\alpha} = 1$ iff there is an edge oriented from j_α to i_α. Let k be the number of edges oriented "clockwise" from v to $v+1$ modulo e. The number of entries equal to -1 that are above the superdiagonal would be k, except the case $v = e$ ruins this, because $(e, 1)$ is the only edge whose clockwise orientation is from the higher-indexed vertex to the lower one. To fix this, we may consider $d_{n,1}$ instead of $d_{1,n}$ then substitute $d_{1,n} = -d_{n,1}$ in the second product so that $y_2 = -d_{e,1}d_{2,3} \cdots d_{e-2,e-1}$.[2] Of the edges of the first perfect matching, exactly some k_1 are oriented clockwise, so $y_1 = (-1)^{k_1}$. Then, exactly $k_2 = k - k_1$ edges of the second perfect matching are oriented clockwise, so $y_2 = -(-1)^{k_2}$.

Suppose C is not oddly oriented. Then $k = k_1 + k_2$ is even, so both k_1 and k_2 have the same parity. Then $y_1 = -y_2$, and not all terms in the Pfaffian have the same sign. ■

Not every graph has a Pfaffian orientation; graphs that do have a Pfaffian orientation are called *Pfaffian graphs.*

Proposition 2.5. *A graph is Pfaffian if it is free of $K_{3,3}$ as a minor.*

The question of which graphs are Pfaffian is generally tricky. Aside from $K_{3,3}$, there are infinitely many such basic non-Pfaffian graphs.

Proposition 2.6. *Every planar graph is Pfaffian.*

The Pfaffian orientation of a planar graph may be constructed by first orienting any spanning tree and then orienting the remaining edges accordingly so each face has an odd number of edges in a direction. Systematically, we take the dual graph; remove from it the edges whose vertices' corresponding faces in the original graph share an edge that is in the spanning tree of the original graph, i.e. those "crossing" an edge of the spanning tree of the original graph; and then traverse edges of the resulting spanning tree of the dual graph from leaves to root, orienting the corresponding edges in the original graph (which are not in the spanning tree and therefore not already oriented). This the process used in the FKT Algorithm. Since calculating the determinant is a trivial task for computers, we then find $\sqrt{\det(D)}$, which is the number of perfect matchings!

[2]Or, perhaps, consider the "entry" $d_{0,1} = d_{n,1}$, and everything lines up along the superdiagonal. Then, simply multiply the entries "on" the superdiagonal, so $y_2 = d_{0,1}d_{2,3} \cdots d_{e-2,e-1}$ In either case, the number of -1s in the considered entries is then k.