Algebraic Approach to Matchings

Lecturer: Michel X. Goemans

These notes on an algebraic approach to solve the matching problem are based on past notes scribed by Zhenyu Liao. Our goal is to derive an algebraic test for deciding if a graph $G = (V, E)$ has a perfect matching. We may assume that the number of vertices is even since this is a necessary condition for having a perfect matching. First, we will define a few basic needed notations.

Definition 1 *A skew-symmetric matrix A is a square matrix which satisfies $A^T = -A$, i.e. if $A = (a_{ij})$ we have $a_{ij} = -a_{ji}$ for all i, j.*

For a graph $G = (V, E)$ with $|V| = n$ and $|E| = m$, we construct a $n \times n$ skew-symmetric matrix $A = (a_{ij})$ with an entry $a_{ij} = -a_{ji}$ for each edge $(i, j) \in E$ and $a_{ij} = 0$ if (i, j) is not an edge; the values a_{ij} for the edges will be specified later.

Recall:

Definition 2 *The determinant of matrix A is*

$$\det(A) = \sum_{\sigma \in S_n} \operatorname{sgn}(\sigma) \prod_{i=1}^{n} a_{i\sigma(i)}$$

where S_n is the set of all permutations of n elements and the $\operatorname{sgn}(\sigma)$ is defined to be 1 if the number of inversions in σ is even and -1 otherwise.

Note that for a skew-symmetric matrix A, $\det(A) = \det(-A^T) = (-1)^n \det(A)$. So if n is odd we have $\det(A) = 0$. Consider K_4, the complete graph on 4 vertices, and thus

$$A = \begin{pmatrix} 0 & a_{12} & a_{13} & a_{14} \\ -a_{12} & 0 & a_{23} & a_{24} \\ -a_{13} & -a_{23} & 0 & a_{34} \\ -a_{14} & -a_{24} & -a_{34} & 0 \end{pmatrix}.$$

By computing its determinant one observes that

$$\det(A) = (a_{12}a_{34} - a_{13}a_{24} + a_{14}a_{23})^2.$$

First, it is the square of a polynomial $q(a)$ in the entries of A, and moreover this polynomial has a monomial precisely for each perfect matching of K_4. This is not a coincidence and this will be proved and exploited in this lecture.

Given any permutation σ of S_n where n is even, we can associate to it a perfect matching M by letting

$$M = \{(\sigma(2i-1), \sigma(2i)) : 1 \leq i \leq n/2\}.$$

Notice that this map is surjective; several permutations will correspond to the same matching. Given a permutation σ and associated matching M and a skew-symmetric matrix A, define the weight of M as follows:

$$wt(M) = \operatorname{sgn}(\sigma) \prod_{i=1}^{n/2} a_{\sigma(2i-1),\sigma(2i)}.$$

We claim that the weight of M is well-defined, in the sense that it is independent of the permutation defining M. If two permutations σ_1, σ_2 correspond to same M, we can obtain one from the other by a series of elementary moves consisting of (i) switching $\sigma(2k-1)$, $\sigma(2k)$ for some k (changing the order of the vertices of an edge) or of (ii) switching $\sigma(2j-1)$, $\sigma(2j)$ with $\sigma(2k-1)$ and $\sigma(2k)$ (i.e., switching two edges of the matching). For the first elementary move, the sign of permutation gets changed, while, at the same time, the corresponding a_{ij} gets multiplied by -1. The second type of elementary move has no effect on the sign of the permutation. So, for both types of moves, the weight is unaffected.

Definition 3 *Let $A = (a_{ij})$ be a $n \times n$ skew-symmetric matrix with n even. The Pfaffian of A is defined by*

$$\mathrm{Pf}(A) = \sum_{M \in \mathcal{M}_n} wt(M)$$

where $\mathcal{M}_n$ is the set of all perfect matchings.

We are now able to state our main theorem by Cayley.

Theorem 1 (Cayley 1842) *If $A_{n\times n}$ is skew symmetric with n even,* $\det(A) = (\mathrm{Pf}(A))^2$.

The core of the proof of this theorem is a bijection between permutations with only even cycles and (ordered) pairs of perfect matchings. This is stated in the next lemma.

Lemma 2 *Let $\xi_n = \{\sigma \in S_n$: every cycle has even length$\}$. Then there exists a bijection T: $\xi_n \to \mathcal{M}_n \times \mathcal{M}_n$.*

This means that $|\xi_n| = |\mathcal{M}_n|^2$. To prove the lemma, we define the map $T(\sigma) = (M, M')$ as follows. For each even cycle in σ, consider the smallest index u, and assign the edge $(u, \sigma(u))$ to M. Then assign $(\sigma(u), \sigma(\sigma(u)))$ to M'. Keep repeating this along the cycle, and for each cycle of σ. This is well defined since every cycle of σ is even. Note that T is a bijection since you can construct a permutation from a pair of matchings by reversing the construction procedure of T. More precisely, given two matchings M and M' consider their union. Now orient each cycle such that the lowest indexed vertex has an arc from M directed away from it. This produces the permutation σ. Notice that T not only maps a permutation σ to a pair (M, M') of perfect matchings, but also implicitly give an orientation of each edge in M and M'.

Example 1 *Let $n = 4$,$\xi_n = \{(12)(34), (13)(24), (14)(23), (1234), (1243), (1324), (1342), (1423), (1432)\}$. $\mathcal{M}_n = \{M_1 = (1, 2), (3, 4), M_2 = (1, 3), (2, 4), M_3 = (1, 4), (2, 3)\}$. Consider $T((1423)) = (M, M')$, by definition, the smallest index is 1 and $\sigma(1) = 4$, thus assign $(1, 4)$ to M. Then assign $(4, \sigma(4)) = (4, 2)$ to M'. Continuing we get $(2, 3) \in M$ and $(3, 1) \in M'$. Finally, $M = \{(1, 4), (2, 3)\}$ (corresponding to M_3) and $M' = \{(4, 2), (3, 1)\}$ (corresponding to M_2).*

Proof: By definition $\det(A) = \sum_{\sigma \in S_n} \mathrm{sgn}(\sigma) \prod_{i=1}^{n} a_{i\sigma(i)}$. We will first divide this summation into three parts.

For those permutations such that $\sigma(i) = i$ for some i, the contribution to the sum is zero since $a_{ii} = 0$.

Consider $\sigma \in S_n \setminus \xi_n$, i.e $\sigma = (C_1)(C_2)...(C_k)$ where the C_i's are the cycles of σ. Taking the odd cycle (if any) in σ involving the element with smallest index, say C_i, and by only reversing C_i, we get a new permutation σ' with the property that $(\sigma')' = \sigma$, $\mathrm{sgn}(\sigma') = \mathrm{sgn}(\sigma)$ but $\prod_{j \in C_i} a_{j,\sigma(j)} = -\prod_{j \in C_i} a_{j,\sigma'(j)}$. So the contribution of such pairs will cancel in the determinant.

Therefore,

$$\det(A) = \sum_{\sigma \in \xi_n} \mathrm{sgn}(\sigma) \prod_{i=1}^{n} a_{i\sigma(i)}. \tag{1}$$

Now, consider the Pfaffian:

$$(Pf(A))^2 = \sum_{(M,M') \in \mathcal{M}_n \times \mathcal{M}_n} wt(M)wt(M').$$

Take now the permutation σ with $T(\sigma) = (M, M')$. Let $\sigma = (C_1)(C_2)...(C_k)$, where each cycle is even, by assumption. To get a permutation α corresponding to M, we write down the elements of each cycle in turn, starting from the lowest index (of each cycle). We see that

$$wt(M) = \operatorname{sgn}(\alpha) \prod_{i=1}^{n/2} a_{\alpha(2i-1)\alpha(2i)}.$$

For M', we use as representative permutation α' the one corresponding to composing α with each cycle in turn. Thus $\operatorname{sgn}(\alpha') = \operatorname{sgn}(\alpha)(-1)^k$. Also,

$$wt(M) = \operatorname{sgn}(\alpha') \prod_{i=1}^{n/2} a_{\alpha'(2i-1)\alpha'(2i)}.$$

Therefore,

$$wt(M)wt(M') = \operatorname{sgn}(\alpha)\operatorname{sgn}(\alpha') \prod_{i=1}^{n} a_{i\sigma(i)} = (-1)^k \prod_{i=1}^{n} a_{i\sigma(i)}.$$

Comparing this with (1) completes the proof of the theorem since $\operatorname{sgn}(\sigma) = (-1)^k$ (as it consists of k even cycles). □

Example 2 *In the example above (continued), we have* $C_1 = (1423)$*,*

$$\alpha = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 1 & 4 & 2 & 3 \end{pmatrix}$$

while

$$\alpha' = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 4 & 2 & 3 & 1 \end{pmatrix},$$

with $\operatorname{sgn}(\alpha) = 1$ *and* $\operatorname{sgn}(\alpha') = -1$*. Therefore,*$wt(M) = a_{14}a_{23}$ *and* $wt(M') = -a_{42}a_{31} = -a_{42}a_{31}$*, and indeed* $wt(M)wt(M') = -a_{14}a_{23}a_{31}a_{42}$ *since* $\operatorname{sgn}((1423)) = -1$*.*

In order to get a test for the existence of a perfect matching from the abve theorem, we consider the entries of A to be indeterminates: $a_{ij} = x_{ij}$ (say) and $a_{ji} = -x_{ij}$. The determinant and the Pfaffian of A are now polynomials in these indeterminates $\{x_e : e \in E\}$. We can thus derive the following known as Tutte's theorem.

Theorem 3 $\det(A(x_1, ..., x_m)) \equiv 0$ *if and only if* G *has no perfect matching.*

In order to check whether $\det(A(x_1, ..., x_m)) \equiv 0$ or not efficiently, we provide a randomized algorithm. The idea (due to Lovász) is to try random values (from a given field F) for the indeterminates x_{ij}, but we need to make sure that if the determinant of A is not identically zero (i.e. that the graph has a perfect matching) then the probability that the determinant cancels to 0 is small. For this, we use the Schwartz-Zippel lemma.

Lemma 4 (Schwartz-Zippel) *Let* $p(x_1, x_2, ..., x_m) \in F_q[X_1, ..., X_m]$ *be a non-zero polynomial with total degree no larger then* d *over a finite field with* q *elements. Let* $r = (r_1, ..., r_m)$ *be chosen uniformly from* F_q^m*. Then*

$$\mathbb{P}[p(r_1, r_2, ..., r_m) = 0] \leq d/q.$$

Proof: By induction on the number of variables m. For $m = 1$, this holds since a polynomial of degree d has no more then d roots. We assume that the lemma holds for polynomials with $m - 1$ variables. Consider now a polynomial with m variables, and we can write $p(x_1, ..., x_m) = \sum_{i=0}^{d} p_i(x_1, ..., x_{m-1})x_m^i$. Therefore, there exists i such that p_i with degree $\leq d - i$ is a non-zero polynomial and assume that i is the largest index satisfying this property. By induction, we know

$$\mathbb{P}[p_i(r_1, ..., r_{m-1}) = 0] \leq (d - i)/q.$$

If $p_i(r_1, ..., r_{m-1}) \neq 0$, then the degree of $p(r_1, r_2, \cdots, r_{m-1}, x_m)$ seen as a polynomial in x_m is i. So $\mathbb{P}[p(r_1, ..., r_m) = 0|p_i(r_1, ..., r_{m-1}) \neq 0] \leq (i/q)$. In all, we get

$$\mathbb{P}[p(r_1, r_2, ..., r_m) = 0] \leq \frac{i}{q}(1-\mathbb{P}[p_i(r_1, ..., r_{m-1}) = 0])+\mathbb{P}[p_i(r_1, ..., r_{m-1}) = 0] \leq \frac{i}{q}+(1-\frac{i}{q})\frac{d-i}{q} < \frac{d}{q}$$

which finishes the proof. □

In our case, the determinant is a polynomial of total degree n (where n is the number of vertices of the graph) and thus, if we choose $q = 2n$, the probability of not identifying that the graph has a perfect matching while it does is $\leq \frac{1}{2}$. To decrease this probability, we can repeat this random experiment k times (and thus with $c \log n$ trials, we can bound the probability of not detecting the existence of a perfect matching when there is one to $\leq n^{-c}$).

Until now, we obtain a randomized algebraic test Has-Perfect(G) to check whether there is a perfect matching in a graph G. This can be used to *find* a perfect matching in G as the property is self-reducible. First, check that whether there is a perfect matching in this graph. Then, pick any edge $e = (u, v)$. If there still exists a perfect matching in $G - e$, delete e from G and continue. If there is no perfect matching in $G - e$, add e to the matching, delete both vertices u and v (and all edges adjacent to it) and recurse on the smaller graph. This will terminate with a perfect matching in G.

Here are some additional results that were briefly discussed in lecture. Harvey has shown that a more clever algorithm allows to find the perfect matching in time $O(n^\omega)$ where this is the time to multiply two $n \times n$ matrices ($\omega \sim 2.3$). One can also use the algebraic approach discussed here to find the largest matching in a graph. Indeed, one can show:

Theorem 5 (Tutte)

$$\max_{\textit{matching } M} |M| = \operatorname{rank}(A(x_1, \cdots, x_m)).$$

The algorithms as described are Monte Carlo (one-sided error). One can also make them Las Vegas (i.e. these are algorithms that always give the correct answer and run in expected polynomial time) by constructing the Edmonds-Gallai decomposition. This was done by Karloff (1986). Geelen (2000) has shown further that the algebraic approach can be made purely deterministic.

Exact Matching

There is an extension of maximum or perfect matching for which the only known approach is the randomized, algebraic approach discussed in these lectures. In the Exact Matching, or Red-Blue Matching, we are given a graph $G = (V, E)$ in which the edges are colored blue or red, we are also given an integer k, and the problem is to decide whether a perfect matching exists with precisely k red edges (and find one if such a matching exists). This is one of the very few problems for which a randomized polynomial-time algorithm is known (even in RNC) while no deterministic polynomial-time algorithms are known. To use the algebraic approach, consider the skew symmetric matrix A with indeterminates where

$$a_{ij} = \begin{cases} \pm x_{ij} & (i,j) \in E \text{ blue} \\ \pm x_{ij} z & (i,j) \in E \text{ red} \\ 0 & \text{otherwise.} \end{cases}$$

Observe that $\det(A[x_1, x_2, \cdots, x_m, z])$ is now a polynomial

$$q(x_1, x_2, \cdots, x_m, z) = \sum_{i=0}^{n} q_i(x_1, \cdots, x_m) z^i.$$

Similarly,

$$\mathrm{Pf}(A) = \sum_{j=0}^{n/2} p_j(x_1, x_2, \cdots, x_m) z^j.$$

From Cayley's theorem 1, we derive that

Corollary 6 *G has a perfect matching with exactly k red edges iff $p_k(x_1, \cdots, x_m) \not\equiv 0$.*

How can we find p_k? For given values for $x_1, \cdots, x_m$, we can evaluate all the q_i's by computing $\det(A)$ in $n+1$ values for z and using extrapolation (as $\det(A)$ is a polynomial of degree $\leq n$ in z. And once we have evaluated all the p_i, we can also evaluate all the q_j's (by taking the square root of a polynomial in one varaiable). Thus, for any values $x_1, \cdots, x_n$, we can evaluate $p_k(x_1, \cdots, x_n)$. But p_k is a polynomial of total degree at most $n/2$. Thus if we choose the x_i's uniformly at random from a field with q values, the probability of finding $p_k(r_1, \cdots, r_m) = 0$ while $p_k(x_1, \cdots, x_m) \not\equiv 0$ is at most $n/2q$, and we can use the same approach as before. As mentioned earlier, it is an open question to find a deterministic algorithm for this problem.

Counting Perfect Matchings

In some cases, we can also use Cayley's theorem to count the number of perfect matchings. Remember that $\mathrm{Pf}(A)$ has a term $wt(M)$ for each perfect matching M. If all $a_{ij} = \pm 1$ for the edges $(i,j) \in E$ then $wt(M) = \pm 1$. If we could choose the sign of a_{ij} so that all perfect matchings have weight $+1$ (or all -1) then we could easily count the perfect matchings as $\sqrt{\det(A)}$. Notice by the way that $\det(A)$ changes as we flip the sign on an edge. To keep track of the signs on the edges, it is convenient to orient the edges of E into $\vec{E}$ and let

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

Definition 4 *$(V, \vec{E})$ is a Pfaffian orientation of (V, E) if for all perfect matchings M and M', we have $wt(M) = wt(M')$.*

If you consider the cycle on 4 vertices $1-2-3-4-1$ and take the (cyclic) orientation for which $1-2-3-4-1$ is a directed cycle then the two perfect matchings have opposite weights and this is not a Pfaffian orientation. However, if you switch the orientation of $(4,1)$ to $(1,4)$ then one can check this is a Pfaffian orientation. The importance of Pfaffian orientations is due to the following which follows directly from Cayley's theorem.

Theorem 7 *If $G = (V, E)$ has a Pfaffian orientation $\vec{E}$ with associated skew-symmetric matrix A then the number of perfect matchings equals $\sqrt{\det(A)}$.*

Unfortunately, not all graphs admit a Pfaffian orientation. For example, $K_{3,3}$ does not. The fact that not all graphs have a Pfaffian orientation is not surprising as the problem of counting the number of perfect matchings is $\#P-$hard in general. Testing if a bipartite graph admits a Pfaffian orientation is known to be polynomial, but the question is open for general graphs.

The main result we will be discussing is that planar graphs admit a Pfaffian orientation. This is a classical result of Kastleyn, a statistical physicist.

Theorem 8 (Kastleyn 1963) *Every planar graphs has a Pfaffian orientation.*

Before we prove this theorem, we need to better understand which orientations are Pfaffian. In the definition of Pfaffian orientations, we only need to consider two perfect matchings M and M' that differ in only one cycle $C = M\Delta M'$. Observe that such a cycle C has (i) $|C|$ even and also is such that (ii) $G[V \setminus V(C)]$ admits a perfect matching. The converse is also true: any cycle C satisfying (i) and (ii) arises from the symmetric difference of two perfect matchings. Now, how do $wt(M)$ and $wt(M')$ differ? Consider two permutations σ and σ' associated with M and M' such that σ' is obtained by composing σ with C. Thus, $\operatorname{sgn}(\sigma') = -\operatorname{sgn}(\sigma)$. Furthermore, $wt(M)wt(M') = \operatorname{sgn}(\sigma)\operatorname{sgn}(\sigma')(-1)^k$ where k is the number of backward oriented arcs along C. Thus, this number must be odd for the orientation to be Pfaffian (and thus also the number of forward oriented arcs will also be odd). Summarizing,

Lemma 9 *An orientation $\vec{E}$ of $G = (V, E)$ is Pfaffian if for every cycle C with $G \setminus V(C)$ having a perfect matching, the number of backward oriented arcs along C is odd.*

For planar graphs, we will be able to exhibit an even stronger type of orientation.

Theorem 10 *Let $G = (V, E)$ be a planar graph with a given planar embedding. Suppose G admits an orientation $\vec{E}$ such that for every internal face F, the number of forward arcs in clockwise traversal is odd. Then $\vec{E}$ is Pfaffian.*

Proof: We first claim that for any (simple) cycle (not necessarily corresponding to a face), the number f of forward clockwise arcs plus the number k of vertices inside C is odd. The proof of the claim is by induction on k. For $k = 0$, this is the hypothesis of the theorem. Now suppose C encloses more than one face. One can view C as being obtained by gluing together 2 cycles C_1 and C_2 such that the faces enclosed by C_i partition the faces enclosed by C. Let f_1, k_1 (resp. f_2, k_2) be the number of forward clockwise arcs and internal vertices for C_1 (resp. C_2). By induction, we can assume that $f_1 + k_1 = 1 \pmod 2$ and $f_2 + k_2 = 1 \pmod 2$. Consider the path P shared by C_1 and C_2 and let b its number of internal vertices. Observe that $k = k_1 + k_2 + b$ and $f = f_1 + f_2 - (b+1)$ as every arc along P is either clockwise forward in C_1 or in C_2. Thus $k + b = k_1 + k_2 + f_1 + f_2 - 1 = 1 \pmod 2$.

Now that the claim is proved, it is easy to show that the orientation is Pfaffian. Indeed, consider any even cycle C such that $G \setminus V(C)$ has a perfect matching. Using planarity, this means that the number of vertices inside C is even (as they must be perfectly matchable) and therefore the claim implies that the number of forward arcs in the orientation is odd. This means that the orientation is Pfaffian. $\square$

For planar graphs, it is easy to find an orientation that satisfies the above lemma. Take any tree, and orient its edges in any way. Now consider all the remaining edges in order in which they close faces (i.e. they are minimal in the sense that no other edge inside it has not been considered yet) and orient it appropriately for the enclosed face to have an odd number of forward clockwise arcs. This orientation is thus Pfaffian.