

Chapter 37

Matchings II

By Sarel Har-Peled, December 30, 2014^①

Version: 1.01

37.1. Maximum Size Matching in a Non-Bipartite Graph

The results from the previous lecture suggests a natural algorithm for computing a maximum size (i.e., matching with maximum number of edges in it) matching in a general (i.e., not necessarily bipartite) graph. Start from an empty matching M and repeatedly find an augmenting path from an unmatched vertex to an unmatched vertex. Here we are discussing the unweighted case.

Notations. Let T be a given tree. For two vertices $x, y \in V(T)$, let τ_{xy} denote the path in T between x and y . For two paths π and π' that share an endpoint, let $\pi \parallel \pi'$ denotes the path resulting from concatenating π to π' . For a path π , let $|\pi|$ denote the number of edges in π .

37.1.1. Finding an augmenting path

We are given a graph G and a matching M , and we would to compute a bigger matching in G . We will do it by computing an augmenting path for M .

We first observe that if G has any edge with both endpoints being free, we can just add it to the current matching. Thus, in the following, we assume that for all edges, at least one of their endpoint is covered by the current matching M . Our task is to find an augmenting path in M .

We start by collapsing the unmatched vertices to a single vertex s , and let H be the resulting graph. Next, we compute an **alternating BFS** of H starting from s . Formally, we perform a BFS on H starting from s such that for the even levels of the tree the algorithm is allowed to traverse only edges in the matching M , and in odd levels the algorithm traverses the unmatched edges. Let T denote the resulting tree.

An augmenting path in G corresponds to an odd cycle in H with passing through the vertex s .

Definition 37.1.1. An edge $uv \in E(G)$ is a **bridge** if the the following conditions are met: (i) u and v have the same depth in T , (ii) if the depth of u in T is even then uv is free (i.e., $uv \notin M$), and (iii) if the depth of u in T is odd then $uv \in M$.

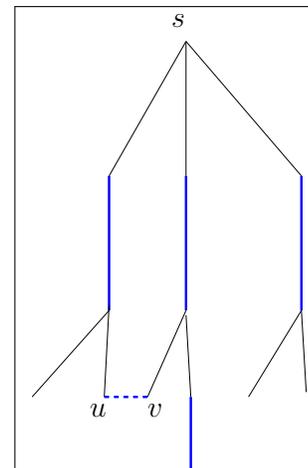


Figure 37.1: A cycle in the alternating BFS tree.

^①This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, USA.

Note, that given an edge uv we can check if it is a bridge in constant time after linear time preprocessing of T and G .

The following is an easy technical lemma.

Lemma 37.1.2. *Let v be a vertex of G , M a matching in G , and let π be the shortest alternating path between s and v in G . Furthermore, assume that for any vertex w of π the shortest alternating path between w and s is the path along π .*

Then, the depth $d_T(v)$ of v in T is $|\pi|$.

Proof: By induction on $|\pi|$. For $|\pi| = 1$ the proof trivially holds, since then v is a neighbor of s in G , and as such it is a child of s in T .

For $|\pi| = k$, consider the vertex just before v on π , and let us denote it by u . By induction, the depth of u in T is $k - 1$. Thus, when the algorithm computing the alternating BFS visited u , it tried to hang v from it in the next iteration. The only possibility for failure is if the algorithm already hanged v in earlier iteration of the algorithm. But then, there exists a shorter alternating path from s to v , which is a contradiction. ■

Lemma 37.1.3. *If there is an augmenting path in G for a matching M , then there exists an edge $uv \in E(G)$ which is a bridge in T .*

Proof: Let π be an augmenting path in G . The path π corresponds to an odd length alternating cycle in H . Let σ be the shortest odd length alternating cycle in G (note that both edges in σ that are adjacent to s are unmatched).

For a vertex x of σ , let $d(x)$ be the length of the shortest alternating path between x and s in H . Similarly, let $d'(x)$ be the length of the shortest alternating path between s and x along σ . Clearly, $d(x) \leq d'(x)$, but we claim that in fact $d(x) = d'(x)$, for all $x \in \sigma$. Indeed, assume for the sake of contradiction that $d(x) < d'(x)$, and let π_1, π_2 be the two paths from x to s formed by σ . Let η be the shortest alternating path between s and x . We know that $|\eta| < |\pi_1|$ and $|\eta| < |\pi_2|$. It is now easy to verify that either $\pi_1 \parallel \eta$ or $\pi_2 \parallel \eta$ is an alternating cycle shorter than σ involving s , which is a contradiction.

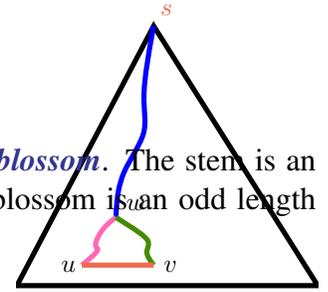
But then, take the two vertices of σ furthest away from s . Clearly, both of them have the same depth in T , since $d(u) = d'(u) = d'(v) = d(v)$. By Lemma 37.1.2, we now have that $d_T(u) = d(u) = d(v) = d_T(v)$. Establishing the first part of the claim. See Figure 37.1.

As for the second claim, observe that it easily follows as σ is created from an alternating path. ■

Thus, we can do the following: Compute the alternating BFS T for H , and find a bridge uv in it. If M is not a maximal matching, then there exists an augmenting path for G and by Lemma 37.1.3 there exists a bridge. Computing the bridge uv takes $O(m)$ time.

Extract the paths from s to u and from s to v in T , and glue them together with the edge uv to form an odd cycle μ in H ; namely, $\mu = \tau_{su} \parallel uv \parallel \tau_{vs}$. If μ corresponds to an alternating path in G then we are done, since we found an alternating path, and we can apply it and find a bigger matching.

But μ , in fact, might have common edges. In particular, let π_{su} and π_{sv} be the two paths from s to u and v , respectively. Let w be the lowest vertex in T that is common to both π_{su} and π_{sv} .



Definition 37.1.4. Given a matching M , a **flower** for M is formed by a **stem** and a **blossom**. The stem is an even length alternating path starting at a free vertex v ending at vertex w , and the blossom is an odd length (alternating) cycle based at w .

Lemma 37.1.5. Consider a bridge edge $uv \in G$, and let w be the least common ancestor (LCA) of u and v in T . Consider the path π_{sw} together with the cycle $C = \pi_{wu} \parallel uv \parallel \pi_{vw}$. Then π_{sw} and C together form a flower.

Proof: Since only the even depth nodes in T have more than one child, w must be of even depth, and as such π_{sw} is of even length. As for the second claim, observe that $\alpha = |\pi_{wu}| = |\pi_{vw}|$ since the two nodes have the same depth in T . In particular, $|C| = |\pi_{wu}| + |\pi_{vw}| + 1 = 2\alpha + 1$, which is an odd number. ■

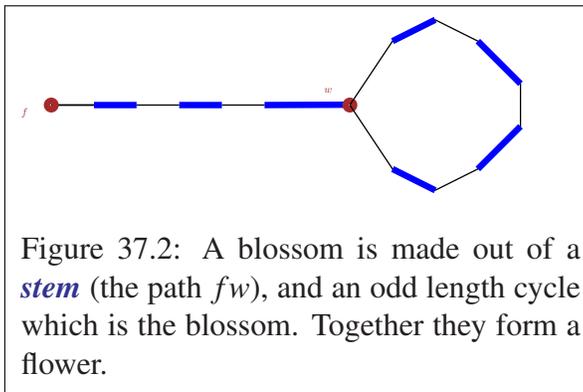


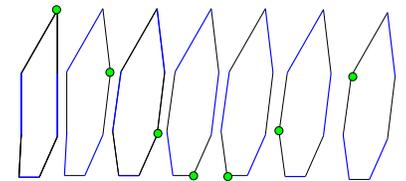
Figure 37.2: A blossom is made out of a **stem** (the path fw), and an odd length cycle which is the blossom. Together they form a flower.

Let us translate this blossom of H back to the original graph G . The path s to w corresponds to an alternating path starting at a free vertex f (of G) and ending at w , where the last edge in the stem is in the matching, the cycle $w \dots u \dots v \dots w$ is an alternating odd length cycle in G where the two edges adjacent to w are unmatched.

We can not apply a blossom to a matching in the hope of getting better matching. In fact, this is illegal and yield something which is not a matching. On the positive side, we discovered an odd alternating cycle in the graph G . Summarizing the above algorithm, we have:

Lemma 37.1.6. Given a graph G with n vertices and m edges, and a matching M , one can find in $O(n + m)$ time, either a blossom in G or an augmenting path in G .

To see what to do next, we have to realize how a matching in G interact with an odd length cycle which is computed by our algorithm (i.e., blossom). In particular, assume that the free vertex in the cycle is unmatched. To get a maximum number of edges of the matching in the cycle, we must at most $(n - 1)/2$ edges in the cycle, but then we can rotate the matching edges in the cycle, such that any vertex on the cycle can be free. See figure on the right.



Let G/C denote the graph resulting from collapsing such an odd cycle C into single vertex. The new vertex is marked by $\{C\}$.

Lemma 37.1.7. Given a graph G , a matching M , and a flower B , one can find a matching M' with the same cardinality, such that the blossom of B contains a free (i.e., unmatched) vertex in M' .

Proof: If the stem of B is empty and B is just formed by a blossom, and then we are done. Otherwise, B was as stem π which is an even length alternating path starting from from a free vertex v . Observe that the matching $M' = M \oplus \pi$ is of the same cardinality, and the cycle in B now becomes an alternating odd cycle, with a free vertex.

Intuitively, what we did is to apply the stem to the matching M . See Figure 37.3. ■

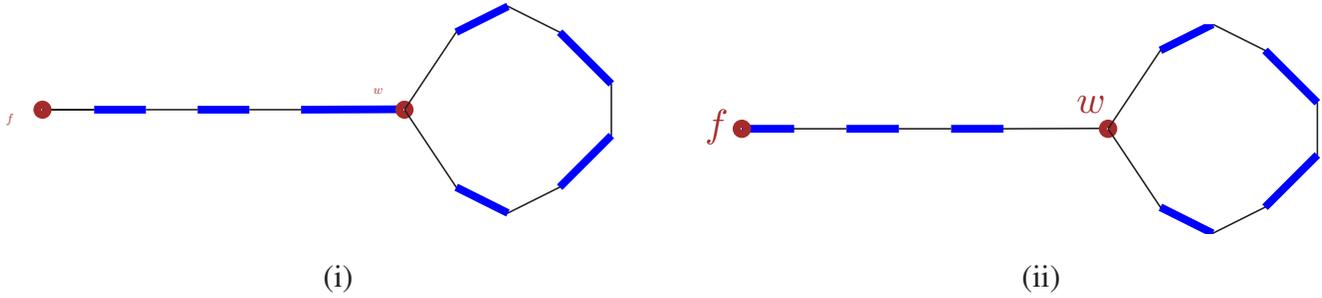


Figure 37.3: (i) the flower, and (ii) the invert stem.

Theorem 37.1.8. *Let M be a matching, and let C be a blossom for M with an unmatched vertex v . Then, M is a maximum matching in G if and only if $M/C = M \setminus C$ is a maximum matching in G/C .*

Proof: Let G/C be the collapsed graph, with $\{C\}$ denoting the vertex that correspond to the cycle C .

Note, that the collapsed vertex $\{C\}$ in G/C is free. Thus, an augmenting path π in G/C either avoids the collapsed vertex $\{C\}$ altogether, or it starts or ends there. In any case, we can rotate the matching around C such that π would be an augmenting path in G . Thus, if M/C is not a maximum matching in G/C then there exists an augmenting path in G/C , which in turn is an augmenting path in G , and as such M is not a maximum matching in G .

Similarly, if π is an augmenting path in G and it avoids C then it is also an augmenting path in G/C , and then M/C is not a maximum matching in G/C .

Otherwise, since π starts and ends in two different free vertices and C has only one free vertex, it follows that π has an endpoint outside C . Let v be this endpoint of π and let u be the first vertex of π that belongs to C . Let σ be the path $\pi[v, u]$.

Let f be the free vertex of C . Note that f is unmatched. Now, if $u = f$ we are done, since then π is an augmenting path also in G/C . Note that if u is matched in C , as such, it must be that the last edge e in π is unmatched. Thus, rotate the matching M around C such that u becomes free. Clearly, then σ is now an augmenting path in G (for the rotated matching) and also an augmenting path in G/C . ■

Corollary 37.1.9. *Let M be a matching, and let C be an alternating odd length cycle with the unmatched vertex being free. Then, there is an augmenting path in G if and only if there is an augmenting path in G/C .*

37.1.2. The algorithm

Start from the empty matching M in the graph G .

Now, repeatedly, try to enlarge the matching. First, check if you can find an edge with both endpoints being free, and if so add it to the matching. Otherwise, compute the graph H (this is the graph where all the free vertices are collapsed into a single vertex), and compute an alternating BFS tree in H . From the alternating BFS, we can extract the shortest alternating cycle based in the root (by finding the highest bridge). If this alternating cycle corresponds to an alternating path in G then we are done, as we can just apply this alternating path to the matching M getting a bigger matching.

If this is a flower, with a stem ρ and a blossom C then apply the stem to M (i.e., compute the matching $M \oplus \rho$). Now, C is an odd cycle with the free vertex being unmatched. Compute recursively an augmenting path π in G/C . By the above discussing, we can easily transform this into an augmenting path in G . Apply this augmenting path to M .

Thus, we succeeded in computing a matching with one edge more in it. Repeat till the process get stuck. Clearly, what we have is a maximum size matching.

37.1.2.1. Running time analysis

Every shrink cost us $O(m + n)$ time. We need to perform $O(n)$ recursive shrink operations till we find an augmenting path, if such a path exists. Thus, finding an augmenting path takes $O(n(m + n))$ time. Finally, we have to repeat this $O(n)$ times. Thus, overall, the running time of our algorithm is $O(n^2(m + n)) = O(n^4)$.

Theorem 37.1.10. *Given a graph G with n vertices and m edges, computing a maximum size matching in G can be done in $O(n^2m)$ time.*

37.2. Maximum Weight Matching in A Non-Bipartite Graph

This the hardest case and it is non-trivial to handle. There are known polynomial time algorithms, but I feel that they are too involved, and somewhat cryptic, and as such should not be presented in class. For the interested student, a nice description of such an algorithm is presented in

Combinatorial Optimization - Polyhedral and efficiency
by Alexander Schrijver
Vol. A, 453–459.

The description above also follows loosely the same book.