Lectures on Nonbipartite Matchings

Michel X. Goemans

Based on scribed notes by Nick Harvey, Alex Levin, Robert Kleinberg, Dan Stratila, Aleksander Madry and others from past courses.

1 Nonbipartite Matching

Our first topic of study is matchings in graphs which are not necessarily bipartite. We begin with some relevant terminology and definitions. A *matching* is a set of edges that share no endvertices. A vertex v is *covered* by a matching if v is incident with an edge in the matching. A matching that covers every vertex is known as a *perfect matching* or a 1*-factor* (i.e., a spanning regular subgraph in which every vertex has degree 1). We will let $\nu(G)$ denote the cardinality of a maximum matching in graph G. A *vertex cover* is a set C of vertices such that every edge is incident with at least one vertex in C. The minimum cardinality of a vertex cover is denoted $\tau(G)$. The following simple proposition relates matchings and vertex covers.

Proposition 1 *If M is a matching and C is a vertex cover then $|M| \leq |C|$.*

Proof: For each edge in M, at least one of the endvertices must be in C, since C covers every edge. Since the edges in M do not share any endvertices, we must have $|M| \leq |C|$. $\square$

This proposition implies that $\nu(G) = \max_M |M| \leq \min_C |C| = \tau(G)$, so $\nu(G) \leq \tau(G)$. Kőnig showed that in fact equality holds if G is a bipartite graph with no isolated vertices. Unfortunately if G is not bipartite then we may have $\nu(G) < \tau(G)$. For example, if G is the cycle on three vertices then $\nu(G) = 1$ but $\tau(G) = 2$. We will give another upper-bound for $\nu(G)$ after introducing some more definitions.

If $G = (V, E)$ is a graph and $U \subseteq V$, $G - U$ denotes the subgraph of G obtained by deleting the vertices of U and all edges incident with them. Let $o(G-U)$ denote the number of components of $G-U$ that contain an odd number of vertices. Let M be a matching in $G-U$ and consider a component of $G-U$ with an odd number of vertices. There must be at least one unmatched vertex v in this component, since any matching necessarily covers an even number of vertices. Treating M as a matching in G, it is possible that we could increase the size of M by matching v with some vertex in U. However, we can add at most $|U|$ edges to M in this manner, since the vertices in U will eventually all be matched. Thus any matching in G must have least $o(G-U) - |U|$ unmatched vertices. This argument shows that the maximum size of a matching is upper-bounded by $(|V| + |U| - o(G-U))/2$, for any subset U. The following theorem strengthens this result.

Theorem 2 (Tutte-Berge Formula) *Let $G = (V, E)$ be a graph. Then*

$$\nu(G) = \max_M |M| = \min_{U \subset V} (|V| + |U| - o(G-U))/2,$$

where the maximization is over all matchings M in G.

Proof: We will consider the case that G is connected. If G is not connected, the result follows by adding the formulas for the individual components. The proof proceeds by induction on the order of G. If G has at most one vertex then the result holds trivially. Otherwise, suppose that G has at least two vertices. We consider two cases.

Case 1: G contains a vertex v that is covered by *all* maximum matchings. The subgraph $G - v$ cannot have a matching of size $\nu(G)$, otherwise that would give a maximum matching for G that leaves v unmatched. Thus $\nu(G - v) = \nu(G) - 1$. By induction the result holds for the graph $G - v$, so there exists a set $U' \subset V - v$ that achieves equality in the Tutte-Berge Formula. Defining $U = U' \cup \{v\}$, we see that

$$\begin{aligned}\nu(G) &= \nu(G - v) + 1 \\ &= (|V - v| + |U'| - o(G - v - U'))/2 + 1 \\ &= ((|V| - 1) + (|U| - 1) - o(G - U))/2 + 1 \\ &= (|V| + |U| - o(G - U))/2\end{aligned}$$

Case 2: For every vertex $v \in G$, there is a maximum matching that does not cover v. We will prove that each maximum matching leaves exactly one vertex uncovered. Suppose to the contrary, that is, each maximum matching leaves at least two vertices uncovered. We choose a maximum matching M and its two uncovered vertices u and v such that we minimize $d(u, v)$, the distance between vertices u and v. If $d(u, v) = 1$ then the edge uv can be added to M to obtain a larger matching, which is a contradiction.

Otherwise, $d(u, v) \geq 2$ so we may fix an intermediate vertex t on some shortest u-v path. By the assumption of the present case, there is a maximum matching N that does not cover t. Furthermore, we may choose N such that its symmetric difference with M is minimal. If N does not cover u then (N, u, t) contradicts our choice of (M, u, v). Thus N covers u and, by symmetry, v as well. Since N and M both leave at least two vertices uncovered, there exists a second vertex $x \neq t$ that is covered by M but not by N. Let (x, y) be the edge in M that is incident with x. If y is also uncovered by N then $N \cup \{(x, y)\}$ is a larger matching than N, a contradiction. So let (y, z) be the edge in N that is incident with y, and note that $z \neq x$. Then $N \cup \{(x, y)\} \setminus \{(y, z)\}$ is a maximum matching that does not cover t and has smaller symmetric difference with M than N does. This contradicts our choice of N, so each maximum matching must leave exactly one vertex uncovered. Then $\nu(G) = (|V| - 1)/2$. The Tutte-Berge Formula then follows by choosing $U = \emptyset$. □

A natural question to ask next is: Given a graph G, what is a set $U \subset V(G)$ giving equality in the Tutte-Berge Formula? Such a set is provided by the **Edmonds-Gallai Decomposition** of G. This decomposition partitions $V(G)$ into three sets: $D(G)$ is the set of all vertices v such that there is some maximum matching that leaves v uncovered, $A(G)$ is the neighbor set of $D(G)$, and $C(G)$ is the set of all remaining vertices.

Theorem 3 (Edmonds-Gallai Decomposition) *Given a graph G, let*

$$\begin{aligned}D(G) &:= \{v : \text{there exists a maximum size matching missing } v\}, \\ A(G) &:= \{v : v \text{ is a neighbor of some } u \in D(G), \text{ but } v \notin D(G)\}, \\ C(G) &:= V(G) \setminus (D(G) \cup A(G)).\end{aligned}$$

Then:

(i) $U = A(G)$ achieves the minimum on the right side of the Tutte-Berge formula,

(ii) $C(G)$ is the union of the even-sized components of $G \setminus A(G)$,

(iii) $D(G)$ is the union of the odd-sized components of $G \setminus A(G)$,

(iv) Every odd-sized component of $G \setminus A(G)$ is factor-critical. *(A graph H is* factor-critical *if for every vertex v, there is a matching in H whose only unmatched vertex is v.)*

Let $G[D(G)]$ be the subgraph of G induced by $D(G)$. The last condition says that every connected component H of $G[D(G)]$ is not only of odd cardinality but we can actually choose in it any particular vertex to be left uncovered.

The Edmonds-Gallai Decomposition of a graph can be found as a byproduct of Edmonds' algorithm for finding a maximum matching, and we will thus postpone the proof of Theorem 3 until after presenting Edmonds' algorithm.

2 Edmonds' Algorithm

Before describing this algorithm, we need some basic results about optimality of matchings. Let M be a matching in a graph G. An *alternating walk* (relative to M) is a walk P in G whose edges are alternately in M and not in M. Sometimes, we call P M-alternating to emphasize the dependence on M. If the alternating walk P has only *distinct* vertices then it is called an *alternating path* (or M-alternating path). If all internal vertices of the walk are distinct and the endpoints are identical then the alternating walk is an *alternating cycle.* And an *augmenting path* for M (or M-augmenting path) is an alternating path with both endvertices uncovered by M, see Figure 1. Let M' be the matching obtained by switching M-edges and non-M-edges along path P (i.e., $M' = M \triangle E(P)$). Then $|M'| = |M| + 1$, which explains why P is called an augmenting path.

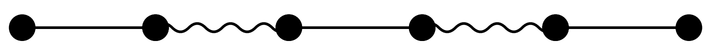

Figure 1: An M-augmenting path

Theorem 4 (Berge) *M is a maximum matching if and only if G contains no M-augmenting path.*

Proof: The "only if" direction is trivial, since any augmenting path can be used to increase the size of M. To prove the other direction, suppose that M is not maximum and let N be a maximum matching chosen with minimum symmetric difference with M. Consider the subgraph spanned by $M \cup N$. Each vertex has degree at most 2, so the subgraph is a disjoint union of paths and cycles. There are no cycles or paths with equal number of edges from N and M, since $N \triangle M$ is minimum. It follows that there is at least one component with more N-edges than M-edges. Such a component is an augmenting path for M. $\square$

Theorem 4 implies the following approach for finding a maximum matching: start with an empty matching and repeatedly find augmenting paths to increase its size. **Edmonds' Algorithm** uses this approach and gives a specific method for finding augmenting paths.

2.1 Flowers, Stems, and Blossoms

Consider a graph $G = (V, E)$ and a matching M in G. Let X be the set of uncovered vertices in G. To find an augmenting path for M, it will be helpful to define an auxiliary directed graph $\hat{G}$ with vertex set V and arc set $A = \{(u, v) \mid \exists x \in V \text{ such that } (u, x) \in E \setminus M \text{ and } (x, v) \in M\}$. Observe that a directed path in $\hat{G}$ corresponds to an (even length) alternating path in G. Furthermore, if there is an augmenting path for M then there is a directed path in $\hat{G}$ starting at a vertex in X and ending at a neighbor of X. Unfortunately, the converse does not necessarily hold: G may contain a directed path in $\hat{G}$ starting at a vertex in X and ending at a neighbor of X that does *not* correspond to an augmenting path. Such a path must necessarily have a prefix that is a *flower*, as shown in

Figure 2. The dotted arcs show a directed path in the auxiliary graph that starts at a vertex in X and ends at a neighbor of set X but does not correspond to an augmenting path. The graph contains a flower, which consists of a *stem* and a *blossom*. The stem is simply an alternating path and the blossom is an odd-length cycle.

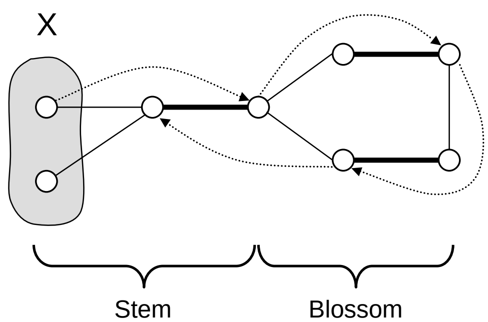

Figure 2: An M-flower. Note that the dashed edges represent edges of the auxiliary graph $\hat{G}$.

Here is a formal definition of all this botany.

Definition 1 *An M-flower is an M-alternating walk $v_0, v_1, v_2, \ldots, v_t$ (numbered so that we have $(v_{2k-1}, v_{2k}) \in M$, $(v_{2k}, v_{2k+1}) \notin M$) satisfying:*

1. $v_0 \in X$.
2. $v_0, v_1, v_2, \ldots, v_{t-1}$ *are distinct.*
3. t *is odd.*
4. $v_t = v_i$, *for an even* i.

The portion of the flower from v_0 to v_i is called the stem, *while the portion from v_i to v_t is called the* blossom.

The next lemma shows that alternating walks between exposed vertices either correspond to an augmenting path or contain a flower.

Lemma 5 *Let M be a matching in G, and let $P = (v_0, v_1, \ldots, v_t)$ be a shortest alternating walk from X to X. Then either P is an M-augmenting path, or $v_0, v_1, \ldots, v_j$ is an M-flower for some $j < t$.*

Proof: If $v_0, v_1, \ldots, v_t$ are all distinct, P is an M-augmenting path. Otherwise, assume $v_i = v_j$, $i < j$, and let j be as small as possible, so that $v_0, v_1, \ldots, v_{j-1}$ are all distinct. We shall prove that $v_0, v_1, \ldots, v_j$ is an M-flower. Properties 1 and 2 of a flower are automatic, by construction. It cannot be the case that j is even, since then $(v_{j-1}, v_j) \in M$, which gives a contradiction in both of the following cases:

- $i = 0$: $(v_{j-1}, v_j) \in M$ contradicts $v_0 \in X$.
- $0 < i < j - 1$: $(v_{j-1}, v_j) \in M$ contradicts the fact that M is a matching, since v_i is already matched to a vertex other than v_{j-1}.

This proves that j is odd. It remains to show that i is even. Assume, by contradiction, that i is odd. This means that (v_i, v_{i+1}) and (v_j, v_{j+1}) are both edges in M. Then $v_{j+1} = v_{i+1}$ (since both

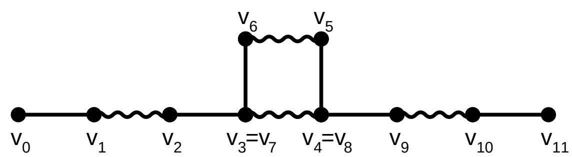

Figure 3: An alternating walk from X to X which can be shortened.

are equal to the other endpoint of the unique matching edge containing $v_j = v_i$), and we may delete the cycle from P to obtain a shorter alternating walk from X to X. (See Figure 3.) $\square$

Given a flower $F = (v_0, v_1, \ldots, v_t)$ with blossom B, observe that for any vertex $v_j \in B$ it is possible to modify M to a matching M' satisfying:

1. Every vertex of F is the endpoint of an edge of M', except v_j.
2. M' agrees with M outside of F, i.e. $M \triangle M' \subseteq F$.
3. $|M'| = |M|$.

To do so, we take M' to consist of all the edges of the stem which do not belong to M, together with a matching in the blossom which covers every vertex except v_j, as well as all the edges in M outside of F.

Whenever a graph G with matching M contains a blossom B, we may simplify the graph by *shrinking* B, a process which we now define.

Definition 2 (Shrinking a blossom) *Given a graph $G = (V, E)$ with a matching M and a blossom B, the* shrunk graph *G/B with matching M/B is defined as follows:*

- $V(G/B) = (V \setminus B) \cup \{b\}$
- $E(G/B) = E \setminus E[B]$
- $M/B = M \setminus E[B]$

where $E[B]$ denotes the set of edges within B, and b is a new vertex disjoint from V.

Observe that M/B is a matching in G, because the definition of a blossom precludes the possibility that M contains more than one edge with one but not both endpoints in B. Observe also that G/B may contain parallel edges between vertices, if G contains a vertex which is joined to B by more than one edge.

The relation between matchings in G and matchings in G/B is summarized by the following theorem.

Theorem 6 *Let M be a matching of G, and let B be an M-blossom. Then, M is a maximum-size matching if and only if M/B is a maximum-size matching in G/B.*

Proof: ($\Longrightarrow$) Suppose N is a matching in G/B larger than M/B. Pulling N back to a set of edges in G, it is incident to at most one vertex of B. Expand this to a matching N^+ in G by adjoining $\frac{1}{2}(|B|-1)$ edges within B to match every other vertex in B. Then we have $|N^+| - |N| = (|B|-1)/2$, while at the same time $|M| - |M/B| = (|B|-1)/2$ (the latter follows because B is an M-blossom, so there are $(|B|-1)/2$ edges of M in B; then M/B contains all the corresponding edges in M except those $(|B|-1)/2$). We conclude that $|N^+|$ exceeds $|M|$ by the same amount that $|N|$ exceeds $|M/B|$.

(⟸) If M is not of maximum size, then change it to another matching M', of equal cardinality, in which B is an entire flower. (If S is the stem of the flower whose blossom is B, then we may take $M' = M \triangle S$.) Note that M'/B is of the same cardinality as M/B, and b is an unmatched vertex of M'/B. Since M' is not a maximum-size matching in G, there exists an M'-augmenting path P. At least one of the endpoints of P is not in B. So number the vertices of P $u_0, u_1, \ldots, u_t$ with $u_0 \notin B$, and let u_i be the first node on P which is in B. (If there is no such node, then $u_i = u_t$.) This sub-path $u_0, u_1, \ldots, u_i$ is an (M'/B)-augmenting path in G/B. □

Note that if M is a matching in G that is not of maximum size, and B is blossom with respect to M, then M/B is not a maximum-size matching in G/B. If we find a maximum-size matching N in G/B, then the proof gives us a way to "unshrink" the blossom B in order to turn N into a matching N^+ of G of size larger than that of M. However, it is important to note that N^+ will not, in general, be a maximum-size matching of G, as the example in Figure 4 shows.

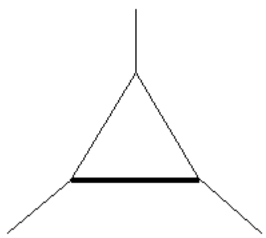

(a) The graph G with the matching M shown in bold

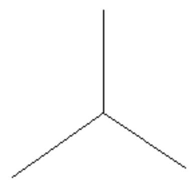

(b) The shrunk graph G/B with the matching M/B (the empty matching)

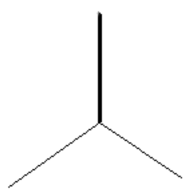

(c) The shrunk graph G/B with a maximum-size matching N

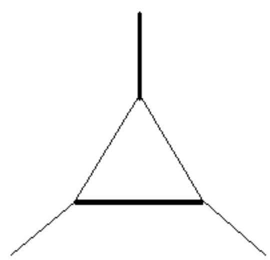

(d) The graph G with the matching N^+ obtained from N by the unshrinking process. Note that it is not a maximum-size matching

Figure 4: A maximum matching in the graph G/B does not necessarily pull back to a maximum matching in G.

2.2 Polynomial-time maximum matching algorithm

Edmonds' algorithm for computing a maximum matching is specified in Figure 5.

```
M := ∅
X := {unmatched vertices} /* Initially all of V. */
Form the directed graph Ĝ.
while Ĝ contains a directed path P̂ from X to N(X)
    Find such a path P̂ of minimum length.
    P := the alternating path in G corresponding to P̂
    if P is an M-augmenting path,
        modify M by augmenting along P.
        Update X and construct Ĝ.
    else
        P contains a blossom B.
        Recursively find a maximum-size matching M′ in G/B.
        if |M′| = |M/B|     /* M is already a max matching. */
            return M        /* Done! */
        else                /* M can be enlarged */
            Unshrink M′ as in the proof of Theorem 6,
             to obtain a matching in G of size > |M|.
             /* It is not necessarily maximal */
            Update M and X and construct the graph Ĝ.
end
```

Figure 5: Algorithm for computing a maximum matching

The correctness of the algorithm is established by Lemma 5 and Theorem 6. The running time may be analyzed as follows. We can compute X and $\hat{G}$ in linear time, and can find $\hat{P}$ in linear time (by breadth-first search). Shrinking a blossom also takes linear time. We can only perform $O(n)$ such shrinkings before terminating or increasing $|M|$. The number of times we increase $|M|$ is $O(n)$. Therefore the algorithm's running time is $O(mn^2)$. With a little more work, this can be improved to $O(n^3)$. (See Schrijver's book.) The fastest known algorithm, due to Micali and Vazirani, runs in time $O(\sqrt{n}\, m)$.

3 Edmonds-Gallai decomposition

So far we have seen the Tutte-Berge formula (Theorem 2) characterizing the size of the maximum matching, and Edmonds' algorithm for computing a maximum matching in any graph G. We willl now show that the proof of correctness of Edmonds' algorithm can be used to prove the Edmonds-Gallai decomposition Theorem 3, see Figure **??** (which, as a by-product, gives another proof of the Tutte-Berge formula).

To prove Theorem 3 for a given graph G, let us consider the maximum-size matching M that is returned by Edmonds' algorithm executed on G. Let X be the set of vertices not matched by M. Consider all the vertices which can be reached by an alternating path from $x \in X$. The first edge on such a path must lie outside of M, the second edge must lie in M, and so on, leading to a picture as in Figure 7.

Motivated by this picture, we define the following three subsets of $V(G)$:

$$\begin{aligned}
\mathsf{Even} &:= \{v : \exists \text{ an alternating path of even length from } X \text{ to } v\},\\
\mathsf{Odd} &:= \{v : \exists \text{ an alternating path from } X \text{ to } v\} \setminus \mathsf{Even},\\
\mathsf{Free} &:= \{v : \nexists \text{ an alternating path from } X \text{ to } v\}.
\end{aligned}$$

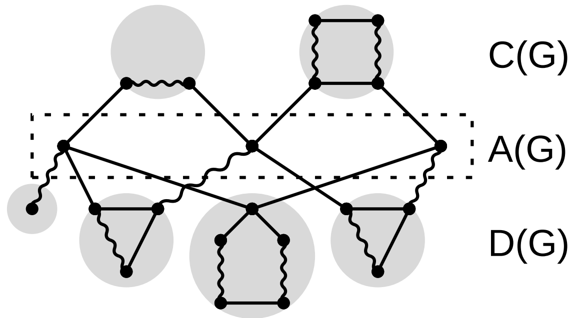

Figure 6: The Edmonds-Gallai Decomposition.

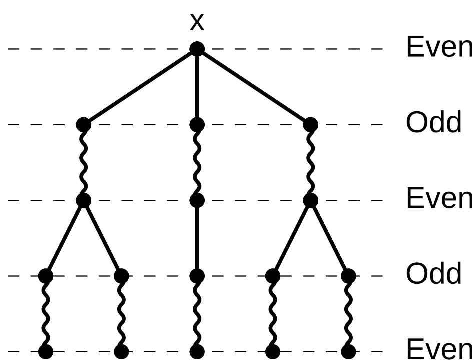

Figure 7: Vertices reachable by alternating paths from $x \in X$.

We will sometimes refer to a vertex as being even, odd, or free, according to which of these sets it belongs to. Note that in the above definitions we are interested in alternating *paths* i.e. alternating walks in which all the vertices are distinct.

We start with the following claim.

Claim 7 *If there is an edge from a vertex $u \in \mathsf{Even}$ to some v, then there is an alternating walk of odd length from X to v, and there is an alternating path from X to v.*

Proof: If $e = (u, v)$ is the edge in question, and P is an alternating path of even length from X to u, then an alternating walk of odd length from X to v is constructed as follows. If $e \in M$, then we take P and delete the final edge, which is necessarily e. If $e \notin M$, then we append e to P. If this alternating walk is not a path, it can only be because v lies on P, in which case P contains a sub-path which is an alternating path from X to v. $\square$

As a result, any vertex adjacent to a vertex from Even has to belong to either Even or Odd. This gives us the following corollary.

Corollary 8 *In G there is no edge between Even and Free.*

Let us now define the *shrunk graph* G_k to be the graph obtained in the final iteration of the execution of Edmonds' algorithm on G. More precisely, G_k is the final graph obtained from G

by repeated shrinking of blossoms performed during the course of the algorithm. Let M_k be the maximum size matching in G_k computed by the algorithm – M_k is just the matching M from which the edges of the blossoms shrunk in G_k have been removed. Note that the set of the vertices of G_k that are unmatched in M_k is still X. Notice also that all vertices of a blossom become even whenever we expand them, since the stem is an even-length alternating path from X to the base v of the blossom, and all other vertices of the blossom are reachable from v by an even-length alternating path which goes around the blossom in one of the directions (as it is odd).

Also, we claim that the vertices in $V(G_k)$ have the same classification (as even, odd, or free) no matter whether we classify them with respect to G_k and M_k, or G and M. Indeed, first consider an alternating path (of even or odd parity) from X to v in G_k. As we expand blossoms, if our alternating path went through the shrunk blossom then we can easily update the alternating path into the expanded graph *without modifying the parity of its length* as the alternating path will be entering the blossom through its base. Conversely, if we have an alternating path P in G from X to a vertex v which intersects a blossom B then consider the first time P visits a vertex of the corresponding flower. We can now replace the this prefix of P with part of the flower in such a way that we still have an alternating path and the parity of the length of the path has not changed.

By properties of the algorithm, G_k has no flowers, and M_k is a maximum matching in G_k. Therefore, G_k has no alternating walk from X to X – if such walk existed then from the previous lecture we would know that there is either an augmenting path or a flower in G_k. This fact implies the following

Claim 9 *In G_k, there is no edge between two even vertices.*

Proof: If such an edge $e = (u, v)$ existed, then by Claim 7, G_k contained an alternating walk P of odd length from X to v. But v is even, so there would also be an alternating path P' of even length from X to v. Concatenating P with the reverse of P', we would obtain an alternating walk from X to X, contradicting the definition of G_k. $\square$

It is worth noting that Claim 9 doesn't necessarily hold in G. This is because, as we already mentioned above, all the vertices of a blossom are even.

We are now ready to prove that the sets Even, Odd, and Free coincide with the sets $D(G)$, $A(G)$, and $C(G)$ from definition of Edmonds-Gallai decomposition.

Claim 10 $\mathsf{Even} = D(G) = \{v : \exists$ *a maximum-size matching missing* $v\}$.

Proof:
($\subseteq$) Certainly if v is even then there is a maximum-size matching M' missing v. Such a matching is obtained by taking an even-length alternating path P from X to v and putting $M' = M \triangle P$.

($\supseteq$) Conversely, if there exists a maximum-size matching M' missing v, then $M \triangle M'$ is a union of even-length cycles and even-length paths, and v is an endpoint of one of these paths, because it does not belong to an edge of M'. The other endpoint of this path P does not belong to an edge of M, i.e. it is an element of X. This confirms that P is an even-length alternating path from X to v. $\square$

Claim 11 $\mathsf{Odd} = A(G) = \{v : v$ *is a neighbor of some* $u \in D(G)$, *but* $v \notin D(G)\}$.

Proof:
($\subseteq$) If v is odd, then there is an alternating path of odd length from X to v. The vertex preceding v on this path must be even, thus v is a neighbor of some vertex from Even. Moreover, since it is odd then it is not in Even. But by Claim 10 $\mathsf{Even} = D(G)$, so indeed $v \in A(G)$.

($\supseteq$) The reverse inclusion follows from Claim 7, which ensures that every vertex adjacent to Even belongs to $\mathsf{Even} \cup \mathsf{Odd}$, which in conjunction with Claim 10 gives us that $v \in \mathsf{Odd}$. $\square$

Claim 12 $\mathsf{Free} = C(G) = V(G) \setminus (D(G) \cup A(G))$.

Proof: Immediate from the definition of Free, and from the preceding two claims which identify $\mathsf{Even}, \mathsf{Odd}$ with $D(G), A(G)$, respectively. $\square$

We proceed to proving the desired properties of the decomposition stated in Theorem 3. We start with property (ii) which is directly implied by the following claim asserting that not only all the vertices of $C(G)$ are matched in M, but also the edges matching them are always connecting two free vertices.

Claim 13 $|M \cap C(G)| = |C(G)|/2$.

Proof: Consider some v from $C(G)$ (which is equal to Free by Claim 12). By Corollary 8 we know that v cannot be adjacent to any even vertex, so $C(G)$ is disconnected from $D(G)$ in $G \setminus A(G)$. Moreover, v has to be matched by some edge $e = (v, u)$ in M, otherwise it would be even. However, u cannot be odd, since then we could augment the odd-length path from X to u by e which would imply that v is either odd or even. Therefore, we must have u being free as well. This implies that $M \cap C(G)$ matches all the vertices of $C(G)$ and thus has the desired size. $\square$

To establish properties (iii) and (iv) we prove the following claim.

Claim 14 *For every connected component H of $(G \setminus A(G)) \cap D(G)$:*

(a) either $|X \cap H| = 1$ and $|M \cap \delta(H)| = 0$; or $|X \cap H| = 0$ and $|M \cap \delta(H)| = 1$, where $\delta(H)$ is the set of edges with exactly one endpoint in H.

(b) H is factor-critical.

Proof: The proof is by induction on the number of blossoms which are shrunk during the execution of Edmonds' algorithm. If no blossoms are shrunk, then $G = G_k$ and the claim follows as a consequence of Corollary 8 and Claim 9 that assert that $(G \setminus A(G)) \cap D(G)$ is a union of isolated vertices (for which both $((a))$ and $((b))$ trivially hold).

Now for the induction step, suppose B is a blossom in G and that the claim holds for G/B (in which B is shrunk). In this case, B corresponds to a vertex $b \in G/B$ which has to be even, since the stem of the flower containing B corresponds to an even-length alternating path from X to b in G/B. In fact, as it was already mentioned before, in G all vertices of B are even and they have all to be in the same connected component, say H_b, of $(G \setminus A(G)) \cap D(G)$.

Clearly, since the vertices of $B \setminus \{b\}$ are all matched in M by edges inside B, neither the size of $M \cap \delta(H_b)$ nor the size of $X \cap H_b$ can increase as a result of expanding B in G/B. Thus, we see that $((a))$ holds.

Now to prove $((b))$, we note that, by inductive assumption, all connected components of $(G \setminus A(G)) \cap D(G)$ other than H_b are factor-critical. Thus, it remains to show that H_b is factor-critical as well. To this end, assume that some vertex $v \in H_b$ was removed. If $v \notin B$ then, by inductive assumption, we know that there exists a matching M' in H_b that matches all vertices of H_b except v and $B \setminus \{b\}$. But then M' can be straight-forwardly augmented inside B to match all vertices in $B \setminus \{b\}$. Similarly, if $v \in B$ then we know that there is a matching M'' that matches all vertices of H_b except B – this correspond to situation in which we remove b in G/B. But, if we remove any vertex of a blossom the rest of them can be easily matched within B, thus once again giving raise to matching that matches the whole $H_b \setminus \{v\}$. This concludes the proof.

$\square$

Having proved Claim 14, property (iii) follows since each factor-critical graph has to be odd-sized, and property (iv) is implied by Claim 13 which shows that all odd-sized connected components of $G \setminus A(G)$ are in $D(G)$.

Finally, we prove property (i).

Claim 15 $|M| = \frac{1}{2}\left[|V| + |A(G)| - o(G \setminus A(G))\right].$

Proof: We only need to show that $|M| \geq \frac{1}{2}\left[|V| + |A(G)| - o(G \setminus A(G))\right]$. Observe that

$$|M| \geq |M \cap C(G)| + |M \cap E(D(G))| + |M \cap \delta(A(G))|.$$

By Claim 13, the first term is $|C(G)|/2$. By Claim 14 part (a), the second term is $\frac{|D(G)|-o(G\setminus A(G))}{2}$ while the third term is $|A(G)|$ since every vertex of $A(G)$ is matched to a vertex of $D(G)$. Thus,

$$|M| \geq \frac{1}{2}(|C(G)| + |D(G)| + 2|A(G)| - o(G \setminus A(G))) = \frac{1}{2}\left[|V| + |A(G)| - o(G \setminus A(G))\right],$$

proving the claim. □

4 Ear-decompositions

An *ear decomposition* $G_k, G_1, \ldots, G_k = G$ of a graph G is a sequence of graphs with the first graph being simple (e.g. a vertex, edge, even cycle, or odd cycle), and each graph G_{i+1} obtained from G_i by *adding an ear*. Adding an ear is done as follows: take two vertices a and b of G_i and add a path P_i from a to b such that all vertices on the path except a and b are new vertices (present in G_{i+1} but not in G_i). An ear with $a \neq b$ is called proper (or open), and an ear with P_i having an odd (even) number of edges is called odd (even). (See Figure 4.) Several basic properties of graphs can be translated into the existence of an ear decomposition of a certain kind. Here are some examples.

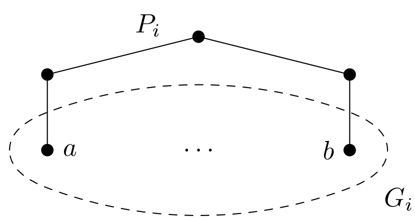

Figure 8: An even proper ear added to G_i.

Theorem 16 (Robbins, 1939 (implicit)) *G is 2-edge-connected if and only if G has an ear decomposition starting from a cycle.*

Theorem 17 (Whitney, 1932) *G is 2-connected if and only if G has a proper ear decomposition starting from a cycle.*

Proof: Obviously, any graph that has a proper ear decomposition starting from a cycle is 2-connected.

Conversely, we assume G is 2-connected, and will show by induction how to construct it starting from a cycle. First, since G is 2-connected, it contains at least one cycle, which we can take as the initial cycle.

Now, suppose we have constructed a subgraph G' of G. If $V(G') = V(G)$ and we are only missing edges, then we can add these edges as proper ears of length one. If $V(G') \subset V(G)$, then pick a vertex $v \in V(G) \setminus V(G')$. Since G is connected, there is a path P from some $a \in V(G)$ to v; since G is 2-connected, there is a path Q distinct from P from v back to some vertex $b \in V(G'), b \neq a$. Hence the paths P and Q form a proper ear from a to b containing at least one new vertex. □

Theorem 18 (Lovász, 1972) *G is factor-critical if and only if G has an odd ear decomposition starting from an odd cycle.*

Proof: If G has an odd ear decomposition, then it is factor critical, since blossoming yields a factor critical graph.

Conversely, suppose G is factor-critical. First, we establish the existence of an initial odd cycle. For any v, fix a near-perfect matching M_v that misses v. Then for an edge (u,v) the existence of M_u and M_v implies there is an alternating even path from v to u. By adding (u,v) to it we obtain an odd cycle.

Fix a vertex v. We proceed by induction; let H be the vertex set already covered by the odd ear decomposition such that no edge in M_v crosses H. Since G is connected, there is an edge $(a,b), a \in H, b \notin H, (a,b) \notin M_v$. Moreover, $M_b \triangle M_v$ contains an alternating path Q from b back to v. The first edge (w,u) to cross back into H on Q is not in M_v, by the construction of H. Therefore, we obtain an odd path from b to u, and can increase the size of H. $\square$

The two results can be combined. One can show that G is factor-critical and 2-connected if and only it has a proper ear decomposition starting from an odd cycle.

We conclude with the following theorem

Theorem 19 *Let G be a 2-connected factor-critical graph. Then the number of near-perfect matchings is at least $|E(G)|$.*

Proof: We proceed by induction on the number of odd ears. Consider a graph G', and G obtained from G' by adding an odd ear $P = (u_0, \ldots, u_k)$ of k edges. Then $|V(G)| = |V(G')| + k - 1, |E(G)| = |E(G')| + k$.

We can obtain $|E(G')|$ near-perfect matchings by taking $(u_1, u_2), \ldots, (u_{k-2}, u_{k-1})$ into the matching, and then generating $|E(G')|$ near perfect matchings in G'. Moreover, we can obtain $k-1$ by matching all vertices on P except $u_j, j = 1, \ldots, k$, and then taking a near-perfect matching on G' that misses either u_0 (if j is odd) or u_k (if j is even). The final matching is obtained by taking the matching missing u_k, but not u_0, removing the edge matching u_k in G' and adding the edge matching u_k in P. $\square$