Chapter 24

Cardinality nonbipartite matching

In this chapter we consider maximum-cardinality matching, with as key results Tutte's characterization of the existence of a perfect matching (implying the Tutte-Berge formula for the maximum-size of a matching) and Edmonds' polynomial-time algorithm to find a maximum-size matching. As in Section 16.1, we call a path P an *M*-augmenting path if P has odd length and connects two vertices not covered by M, and its edges are alternatingly out of and in M. By Theorem 16.1, a matching M has maximum size if and only if there is no M-augmenting path. We say that a matching M covers a vertex v if v is incident with an edge

We say that a matching M covers a vertex v if v is incident with an edge in M. If M does not cover v, we say that M misses v.

In this chapter, graphs can be assumed to be simple.

24.1. Tutte's 1-factor theorem and the Tutte-Berge formula

A basic result of Tutte [1947b] characterizes graphs that have a perfect matching. Berge [1958a] observed that it implies a min-max formula for the maximum size of a matching in a graph, the Tutte-Berge formula.

Call a component of a graph odd if it has an odd number of vertices. For any graph G, let

(24.1) o(G) := number of odd components of G.

Let $\nu(G)$ denotes the maximum size of a matching. Then:

Theorem 24.1 (Tutte-Berge formula). For each graph G = (V, E),

(24.2)
$$\nu(G) = \min_{U \subseteq V} \frac{1}{2} (|V| + |U| - o(G - U)).$$

Proof. To see \leq , we have for each $U \subseteq V$:

(24.3)
$$\nu(G) \le |U| + \nu(G - U) \le |U| + \frac{1}{2}(|V \setminus U| - o(G - U))$$
$$= \frac{1}{2}(|V| + |U| - o(G - U)).$$

We prove the reverse inequality by induction on |V|, the case $V = \emptyset$ being trivial. We can assume that G is connected, since otherwise we can apply induction to the components of G.

First assume that there exists a vertex v covered by all maximum-size matchings. Then $\nu(G-v) = \nu(G) - 1$, and by induction there exists a subset U' of $V \setminus \{v\}$ with

(24.4)
$$\nu(G-v) = \frac{1}{2}(|V \setminus \{v\}| + |U'| - o(G-v-U')).$$

Then $U := U' \cup \{v\}$ gives equality in (24.2), since

(24.5) $\nu(G) = \nu(G - v) + 1 = \frac{1}{2}(|V \setminus \{v\}| + |U'| - o(G - v - U')) + 1$ $= \frac{1}{2}(|V| + |U| - o(G - U)).$

So we can assume that there is no such v. In particular, $\nu(G) < \frac{1}{2}|V|$. We show that there exists a matching of size $\frac{1}{2}(|V|-1)$, which implies the theorem (taking $U := \emptyset$).

Indeed, suppose to the contrary that each maximum-size matching M misses at least two distinct vertices u and v. Among all such M, u, v, choose them such that the distance dist(u, v) of u and v in G is as small as possible.

If $\operatorname{dist}(u, v) = 1$, then u and v are adjacent, and hence we can augment M by the edge uv, contradicting the maximality of |M|. So $\operatorname{dist}(u, v) \geq 2$, and hence we can choose an intermediate vertex t on a shortest u - v path. By assumption, there exists a maximum-size matching N missing t. Choose such an N with $|M \cap N|$ maximal.

By the minimality of dist(u, v), N covers both u and v. Hence, as M and N cover the same number of vertices, there exists a vertex $x \neq t$ covered by M but not by N. Let $x \in e = xy \in M$. Then y is covered by some edge $f \in N$, since otherwise $N \cup \{e\}$ would be a matching larger than N. Replacing N by $(N \setminus \{f\}) \cup \{e\}$ would increase its intersection with M, contradicting the choice of N.

(This proof is based on the proof of Lovász [1979b] of Edmonds' matching polytope theorem.)

The Tutte-Berge formula immediately implies Tutte's 1-factor theorem. A *perfect matching* (or 1-*factor*) is a matching covering all vertices of the graph.

Corollary 24.1a (Tutte's 1-factor theorem). A graph G = (V, E) has a perfect matching if and only if G - U has at most |U| odd components, for each $U \subseteq V$.

Proof. Directly from the Tutte-Berge formula (Theorem 24.1), since G has a perfect matching if and only if $\nu(G) \geq \frac{1}{2}|V|$.

24.1a. Tutte's proof of his 1-factor theorem

The original proof of Tutte [1947b] of his 1-factor theorem (Corollary 24.1a), with a simplification of Maunsell [1952], and smoothed by Halton [1966] and Lovász [1975d], is as follows.

Suppose that there exist graphs G = (V, E) satisfying the condition, but not having a perfect matching. Fixing V, take such a graph G with G simple and |E| as large as possible. Let $U := \{v \in V \mid v \text{ is adjacent to every other vertex of } G\}$. We show that each component of G - U is a complete graph.

Suppose to the contrary that there are distinct $a, b, c \notin U$ with $ab, bc \in E$ and $ac \notin E$. By the maximality of |E|, adding ac to E makes that G has a perfect matching (since the condition is maintained under adding edges). So G has a matching M missing precisely a and c. As $b \notin U$, there exists a vertex d with $bd \notin E$. Again by the maximality of |E|, G has a matching N missing precisely b and d. Now each component of $M \triangle N$ contains the same number of edges in M as in N — otherwise there would exist an M- or N-augmenting path, and hence a perfect matching in G, a contradiction. So the component P of $M \triangle N$ containing d is a path starting at d, with first edge in M and last edge in N, and hence ending at a or c; by symmetry we may assume that it ends at a. Moreover, P does not traverse b. Then extending P by the edge ab gives an N-augmenting path, and hence a perfect matching in G — a contradiction.

So each component of G-U is a complete graph. Moreover, by the condition, G-U has at most |U| odd components. This implies that G has a perfect matching, contradicting our assumption.

More proofs were given by Gallai [1950,1963b], Edmonds [1965d], Balinski [1970], Anderson [1971], Brualdi [1971d], Hetyei [1972,1999], Mader [1973], and Lovász [1975a,1979b].

24.1b. Petersen's theorem

The following theorem of Petersen [1891] is a consequence of Tutte's 1-factor theorem (a graph is *cubic* if it is 3-regular):

Corollary 24.1b (Petersen's theorem). A bridgeless cubic graph has a perfect matching.

Proof. Let G = (V, E) be a bridgeless cubic graph. By Tutte's 1-factor theorem, we should show that G - U has at most |U| odd components, for each $U \subseteq V$.

Each odd component of G-U is left by an odd number of edges (as G is cubic), and hence by at least three edges (as G is bridgeless). On the other hand, U is left by at most 3|U| edges, since G is cubic. Hence G-U has at most |U| odd components.

24.2. Cardinality matching algorithm

The idea of finding an M-augmenting path to increase a matching M is fundamental in finding a maximum-size matching. However, the simple trick

for bipartite graphs, of orienting the edges based on the colour classes of the graph, does not extend to the nonbipartite case. Yet one could try to find an M-augmenting path by finding an 'M-alternating walk', but such a walk can run into a loop that cannot simply be deleted. It was Edmonds [1965d] who found the trick to resolve this problem, namely by 'shrinking' the loop (for which he introduced the term 'blossom'). Then applying recursion to a smaller graph solves the problem¹.

Let G = (V, E) be a graph, let M be a matching in G, and let X be the set of vertices missed by M. A walk $P = (v_0, v_1, \ldots, v_t)$ is called M-alternating if for each $i = 1, \ldots, t - 1$ exactly one of the edges $v_{i-1}v_i$ and v_iv_{i+1} belongs to M. Note that one can find a shortest M-alternating X - X walk of positive length, by considering the auxiliary directed graph D = (V, A) with

(24.6)
$$A := \{ (u, v) \mid \exists x \in V : ux \in E, xv \in M \}.$$

Then each *M*-alternating X - X walk of positive length yields a directed X - N(X) path in *D*, and vice versa (where N(X) denotes the set of neighbours of *X*).

An *M*-alternating walk $P = (v_0, v_1, \ldots, v_t)$ is called an *M*-flower if t is odd, v_0, \ldots, v_{t-1} are distinct, $v_0 \in X$, and $v_t = v_i$ for some even i < t. Then the circuit $(v_i, v_{i+1}, \ldots, v_t)$ is called an *M*-blossom (associated with the *M*-flower).



The core of the algorithm is the following observation. Let G = (V, E) be a graph and let B be a subset of V. Denote by G/B the graph obtained by contracting (or shrinking) B to one new vertex, called B. That is, G/B has vertex set $(V \setminus B) \cup \{B\}$, and for each edge e of G an edge obtained from eby replacing any end vertex in B by the new vertex B. (We ignore loops that may arise.) We denote the new edge again by e. (So its ends are modified,

¹ The idea of applying shrinking recursively to matching problems was introduced by Petersen [1891], and was applied in an algorithmic way by Brahana [1917].

but not its name.) We say that the new edge is the *image* (or *projection*) of the original edge.

For any matching M, let M/B denote the set of edges in G/B that are images of edges in M not spanned by B. Obviously, if M intersects $\delta(B)$ in at most one edge, then M/B is a matching in G/B. In the following, we identify a blossom with its set of vertices.

Theorem 24.2. Let B be an M-blossom in G. Then M is a maximum-size matching in G if and only if M/B is a maximum-size matching in G/B.

Proof. Let $B = (v_i, v_{i+1}, ..., v_t)$.

First assume that M/B is not a maximum-size matching in G/B. Let P be an M/B-augmenting path in G/B. If P does not traverse vertex B of G/B, then P is also an M-augmenting path in G. If P traverses vertex B, we may assume that it enters B with some edge uB that is not in M/B. Then $uv_j \in E$ for some $j \in \{i, i+1, \ldots, t\}$.

(24.7) If j is odd, replace vertex B in P by $v_j, v_{j+1}, \ldots, v_t$. If j is even, replace vertex B in P by $v_j, v_{j-1}, \ldots, v_i$.

In both cases we obtain an M-augmenting path in G. So M is not maximumsize.

Conversely, assume that M is not maximum-size. We may assume that i = 0, that is, $v_i \in X$, since replacing M by $M \triangle EQ$, where Q is the path (v_0, v_1, \ldots, v_i) , does not modify the theorem. Let $P = (u_0, u_1, \ldots, u_s)$ be an M-augmenting path in G. If P does not intersect B, then P is also an M/B-augmenting path in G/B. If P intersects B, we may assume that $u_0 \notin B$. (Otherwise replace P by its reverse.) Let u_j be the first vertex of P in B. Then $(u_0, u_1, \ldots, u_{j-1}, B)$ is an M/B-augmenting path in G/B. So M/B is not maximum-size.

Another useful observation is:

Theorem 24.3. Let $P = (v_0, v_1, \ldots, v_t)$ be a shortest *M*-alternating X - X walk. Then either *P* is an *M*-augmenting path or (v_0, v_1, \ldots, v_j) is an *M*-flower for some $j \leq t$.

Proof. Assume that P is not a path. Choose i < j with $v_j = v_i$ and with j as small as possible. So v_0, \ldots, v_{j-1} are all distinct.

If j - i would be even, we can delete v_{i+1}, \ldots, v_j from P so as to obtain a shorter M-alternating X - X walk. So j - i is odd. If j is even and i is odd, then $v_{i+1} = v_{j-1}$ (as it is the vertex matched to $v_i = v_j$), contradicting the minimality of j.

Hence j is odd and i is even, and therefore (v_0, v_1, \ldots, v_j) is an M-flower.

We now describe an algorithm (the *matching-augmenting algorithm*) for the following problem:

(24.8) given: a matching M; find: an M-augmenting path, if any.

Denote the set of vertices missed by M by X.

(24.9) If there is no *M*-alternating X - X walk of positive length, there is no *M*-augmenting path. If there exists an *M*-alternating X - X walk of positive length, choose a shortest one, $P = (v_0, v_1, \ldots, v_t)$ say. **Case 1:** *P* is a path. Then output *P*. **Case 2:** *P* is not a path. Choose *j* such that (v_0, \ldots, v_j) is an *M*-flower, with *M*-blossom *B*. Apply the algorithm (recursively) to *G*/*B* and *M*/*B*, giving an *M*/*B*-augmenting path *P* in *G*/*B*. Expand *P* to an *M*-augmenting path in *G* (cf. (24.7)).

The correctness of this algorithm follows from Theorems 24.2 and 24.3. It gives a polynomial-time algorithm to find a maximum-size matching, which is a basic result of Edmonds [1965d].

Theorem 24.4. Given a graph, a maximum-size matching can be found in time $O(n^2m)$.

Proof. The algorithm directly follows from algorithm (24.9), since, starting with $M = \emptyset$, one can iteratively apply it to find an *M*-augmenting path *P* and replace *M* by $M \triangle EP$. It terminates if there is no *M*-augmenting path, whence *M* is a maximum-size matching.

By using (24.6), path P in (24.9) can be found in time O(m). Moreover, the graph G/B can be constructed in time O(m). Since the recursion has depth at most n, an M-augmenting path can be found in time O(nm). Since the number of augmentations is at most $\frac{1}{2}n$, the time bound follows.

This implies for perfect matchings:

Corollary 24.4a. A perfect matching in a graph (if any) can be found in time $O(n^2m)$.

Proof. Directly from Theorem 24.4, as a perfect matching is a maximum-size matching.

24.2a. An $O(n^3)$ algorithm

The matching algorithm described above consists of a series of matching augmentations. Each matching augmentation itself consists of a series of two steps performed alternatingly:

(24.10) finding an *M*-alternating walk, and shrinking an *M*-blossom,

until the *M*-alternating walk is simple, that is, is an *M*-augmenting path.

Each of these two steps can be done in time O(m). Since there are at most n shrinkings and at most n matching augmentations, we obtain the $O(n^2m)$ time bound.

If we want to save time we must consider speeding up both the walk-finding step and the shrinking step. In a sense, our description above gives a brute-force polynomial-time method. The O(m) time bound for shrinking gives us time to construct the shrunk graph completely, by copying all vertices that are not in the blossom, by introducing a new vertex for the shrunk blossom, and by introducing for each original edge its 'image' in the shrunk graph. The O(m) time bound for finding an *M*-alternating walk gives us time to find, after any shrinking, a walk starting just from scratch.

In fact, we cannot do much better if we explicitly construct the shrunk graph. But if we modify the graph only locally, by shrinking the *M*-blossom *B* and removing loops and parallel edges, this can be done in time O(|B|n). Since the sum of |B|over all *M*-blossoms *B* is O(n), this yields a time bound of $O(n^2)$ for shrinking.

To reduce the O(m) time for walk-finding, we keep data from the previous walksearch for the next walk-search, with the help of an *M*-alternating forest, defined as follows.



An *M*-alternating forest

Let G = (V, E) be a simple graph and let M be a matching in G. Define X to be the set of vertices missed by M. An *M*-alternating forest is a subset F of E satisfying:

(24.11) F is a forest with $M \subseteq F$, each component of (V, F) contains either exactly one vertex in X or consists of one edge in M, and each path in F starting in X is M-alternating

(cf. Figure 24.2). For any M-alternating forest F, define

(24.12) $even(F) := \{v \in V \mid F \text{ contains an even-length } X - v \text{ path}\},\ odd(F) := \{v \in V \mid F \text{ contains an odd-length } X - v \text{ path}\},\ free(F) := \{v \in V \mid F \text{ contains no } X - v \text{ path}\}.$

Then each $u \in \text{odd}(F)$ is incident with a unique edge in $F \setminus M$ and a unique edge in M. Moreover:

(24.13) if there is no edge connecting even(F) and $even(F) \cup free(F)$, then M is a maximum-size matching.

Indeed, if there is no such edge, even(F) is a stable set in G - odd(F). Hence, setting U := odd(F):

$$(24.14) o(G - U) \ge |even(F)| = |X| + |odd(F)| = (|V| - 2|M|) + |U|,$$

and hence M has maximum size by (24.2).

Now algorithmically, we keep, next to E and M, an M-alternating forest F. We keep the set of vertices by a doubly linked list. We keep for each vertex v, the edges in E, M, and F, incident with v as doubly linked lists. We also keep the incidence functions $\chi^{\text{even}(F)}$ and $\chi^{\text{odd}(F)}$. Moreover, we keep for each vertex v of G one edge $e_v = vu$ with $u \in \text{even}(F)$, if such an edge exists.

Initially, F := M and for each $v \in V$ we select an edge $e_v = vu$ with $u \in X$ (if any). The iteration is:

(24.15) Find a vertex $v \in even(F) \cup free(F)$ for which $e_v = vu$ exists.

Case 1: $v \in \text{free}(F)$. Add uv to F. Let vw be the edge in M incident with v. For each edge wx incident with w, set $e_x := wx$.

Case 2: $v \in even(F)$. Find the X - u and X - v paths P and Q in F.

Case 2a: P and Q are disjoint. Then P and Q form with uv an M-augmenting path.

Case 2b: *P* and *Q* are not disjoint. Then *P* and *Q* contain an *M*-blossom *B*. For each edge bx with $b \in B$ and $x \notin B$, set $e_x := Bx$. Replace *G* by *G*/*B* and remove all loops and parallel edges from *E*, *M*, and *F*.

The number of iterations is at most |V|, since, in each iteration, |V| + |free(F)| decreases by at least 2 (one of these terms decreases by at least 2 and the other does not change). We end up either with a matching augmentation or with the situation that there is no edge connecting even(F) and $\text{even}(F) \cup \text{free}(F)$, in which case M has maximum size by (24.13).

It is easy to update the data structure in Case 1 in time O(n). In Case 2, the paths P and Q can be found in time O(n), and hence in Case 2a, the M-augmenting path is found in time O(n).

Finally, the data structure in Case 2b can be updated in O(|B|n) time². Also a matching augmentation in G/B can be transformed to a matching augmentation in G in time O(|B|n). Since |B| is bounded by twice the decrease in the number of vertices of the graph, this takes time $O(n^2)$ overall.

Hence a matching augmentation can be found in time $O(n^2)$, and therefore:

Theorem 24.5. A maximum-size matching can be found in time $O(n^3)$.

Proof. From the above.

The first $O(n^3)$ -time cardinality matching algorithm was published by Balinski [1969], and consists of a depth-first strategy to find an *M*-alternating forest, replacing shrinking by a clever labeling technique.

Bottleneck in a further speedup is storing the shrinking. With the disjoint set union data structure of Tarjan [1975] one can obtain an $O(nm\alpha(m,n))$ -time algorithm (Gabow [1976a]). A special set union data structure of Gabow and Tarjan [1983,1985] gives an O(nm)-time algorithm. An $O(\sqrt{n}m)$ -time algorithm was announced (with partial proof) by Micali and Vazirani [1980]. A proof was given by Blum [1990], Vazirani [1990,1994], and Gabow and Tarjan [1991] (cf. Peterson and Loui [1988]).

24.3. Matchings covering given vertices

Brualdi [1971d] derived from Tutte's 1-factor theorem the following extension of the Tutte-Berge formula:

Theorem 24.6. Let G = (V, E) be a graph and let $T \subseteq V$. Then the maximum size of a subset S of T for which there is a matching covering S is equal to the minimum value of

 $(24.16) |T| + |U| - o_T(G - U)$

over $U \subseteq V$. Here $o_T(G-U)$ denotes the number of odd components of G-U contained in T.

Proof. For any matching M in G and any $U \subseteq V$, at most |U| odd components of G - U can be covered completely by M. So M misses at least $o_T(G-U) - |U|$ vertices in T. This shows that the minimum is not less than the maximum.

To see equality, let μ be equal to the minimum. Let C be a set disjoint from V with |C| = |V| and let $C' \subseteq C$ with $|C'| = |T| - \mu$. Make a new graph H by extending G by C, in such a way that C is a clique, each vertex in C'

² For each $Z \in \{E, M, F\}$, we scan the vertices b in B, and for $b \in B$ we scan the Zneighbours w of b. If w does not belong to B and was not met as a Z-neighbour of an
earlier scanned vertex in B, we replace bw by Bw in Z. Otherwise, we delete bw from Z.

is adjacent to each vertex in V, and each vertex in $C \setminus C'$ is adjacent to each vertex in $V \setminus T$.

If H has a perfect matching M, then M contains at most $|C'| = |T| - \mu$ edges connecting T and C (since T is not connected to $C \setminus C'$). Hence at least μ vertices in T are covered by edges in M spanned by V, as required.

So we may assume that H has no perfect matching. Then by Tutte's 1-factor theorem, there is a set W of vertices of H such that H - W has at least |W| + 2 odd components (since |V| + |C| is even).

If $C' \not\subseteq W$, then H - W has only one component (since each vertex in C' is adjacent to every other vertex), a contradiction. If $C \subseteq W$, then H - W has at most |V| components, while $|W| + 2 \ge |C| + 2 = |V| + 2$, a contradiction.

So $C' \subseteq W$ and $C \setminus C' \not\subseteq W$. Then at most one component of H - W is not contained in T (since $C \setminus C'$ is a clique and each vertex in $C \setminus C'$ is adjacent to each vertex in $V \setminus T$). Let $U := W \cap V$. Then

(24.17)
$$o_T(G-U) = o_T(H-W) \ge o(H-W) - 1 > |W| \ge |C'| + |U|$$

= $|T| - \mu + |U|$,

contradicting the definition of μ .

(This theorem was also given by Las Vergnas [1975b].)

A consequence is a result of Lovász [1970c] on sets of vertices covered by matchings:

Corollary 24.6a. Let G = (V, E) be a graph and let T be a subset of V. Then G has a matching covering T if and only if T contains at most |U| odd components of G - U, for each $U \subseteq V$.

Proof. Directly from Theorem 24.6.

(This theorem was also given by McCarthy [1975].)

24.4. Further results and notes

24.4a. Complexity survey for cardinality nonbipartite matching

$O(n^2m)$	Edmonds [1965d] (cf. Witzgall and Zahn [1965])
$O(n^3)$	Balinski [1969] (also Gabow [1973,1976a], Karzanov [1976], Lawler [1976b])
$O(nm\alpha(m,n))$	Gabow [1976a]
$O(n^{5/2})$	Even and Kariv [1975], Kariv [1976] (also Bartnik [1978])
$O(\sqrt{n} m \log n)$	Even and Kariv [1975], Kariv [1976]

 \gg

Section 24.4b. The Edmonds-Gallai	decomposition	of a graph	423
-----------------------------------	---------------	------------	-----

continued	
$O(\sqrt{n}m\log\log n)$	Kariv [1976]
$O(\sqrt{n} m + n^{1.5+\varepsilon})$	Kariv [1976] for each $\varepsilon > 0$
$O(\sqrt{n}m)$	announced by Micali and Vazirani [1980], full proof in Blum [1990], Vazirani [1990,1994], and Gabow and Tarjan [1991](cf. Gabow and Tarjan [1983,1985])
$O(\sqrt{n} m \log_n \frac{n^2}{m})$	Goldberg and Karzanov [1995]

Here * indicates an asymptotically best bound in the table. (Kameda and Munro [1974] claim to give an O(nm)-time cardinality matching algorithm, but the proof contains some errors which I could not resolve.)

Gabow and Tarjan [1988a] observed that the method of Micali and Vazirani [1980] also implies that one can find, for given k, a matching of size at least $\nu(G) - \frac{n}{k}$ in time O(km). They derived that a maximum-size matching M minimizing max_{e \in M} w(e) can be found in time $O(\sqrt{n \log n} m)$. (the 'bottleneck matching problem').

Mulmuley, Vazirani, and Vazirani [1987a,1987b] showed that 'matching is as easy as matrix inversion', which is especially of interest for the parallel complexity.

24.4b. The Edmonds-Gallai decomposition of a graph

There is a canonical set U that attains the minimum in (24.2). It has the property that the odd components of G-U cover an inclusionwise minimal set of vertices, and is given by the *Edmonds-Gallai decomposition*, independently found by Edmonds [1965d] and Gallai [1963a,1964].

Let G = (V, E) be a graph. The Edmonds-Gallai decomposition of G is the partition of V into D(G), A(G), and C(G) defined as follows (recall that $N(U) := \{v \in V \setminus U \mid \exists u \in U : uv \in E\}$):

 $\begin{array}{ll} (24.18) & D(G) := \{v \in V \mid \text{there exists a maximum-size matching missing } v\}, \\ & A(G) := N(D(G)), \\ & C(G) := V \setminus (D(G) \cup A(G)). \end{array}$

It yields a 'canonical' certificate of maximality of a matching:

Theorem 24.7. U := A(G) attains the minimum in (24.2), D(G) is the union of the odd components of G-U, and (hence) C(G) is the union of the even components of G-U.

Proof. Case 1: D(G) is a stable set. Let M be a maximum-size matching and let X be the set of vertices missed by M. Then each vertex v in A(G) is contained in an edge $uv \in M$ (as $v \notin D(G)$). We show that $u \in D(G)$. Assume that $u \notin D(G)$.

Since $v \in A(G) = N(D(G))$, there is an edge vw with $w \in D(G)$. Let N be a matching missing w. Then $M \triangle N$ contains a path component starting at a vertex in X and ending at w. Let (v_0, v_1, \ldots, v_t) be this path, with $v_0 \in X$ and $v_t = w$. Then t is even and $v_i \in D(G)$ for each even i (because $M \triangle \{v_0v_1, v_2v_3, \ldots, v_{i-1}v_i\}$ is a

maximum-size matching missing v_i). Hence, assuming $u \notin D(G)$, the edge vu is not on P. So extending P by wv and vu gives a path Q. Then $M \triangle Q$ is a maximum-size matching missing u. So $u \in D(G)$.

As this is true for any $v \in A(G)$, we see that part of M matches A(G) and $D(G) \setminus X$. Hence

(24.19)
$$o(G-U) \ge |D(G)| = |X| + |A(G)| = |V| - 2|M| + |U|.$$

So U attains the minimum in (24.2), and moreover o(G - U) = |D(G)|, that is, D(G) is the union of the odd components of G - U.

Case 2: D(G) spans some edge e = uv. Let M and N be maximum-size matchings missing u and v, respectively. Then $M \cup N$ contains a path component P starting at u. If it does not end at v, then $P \cup \{e\}$ forms an N-augmenting path, contradicting the maximality of N. So P ends at v, and hence $P \cup \{e\}$ gives an M-blossom B.

Let G' := G/B and M' := M/B and let X' be the set of vertices of G' missed by M'. By Theorem 24.2, $|M'| = \nu(G')$. Then

$$(24.20) D(G') = (D(G) \setminus B) \cup \{B\},\$$

since $B \in D(G')$ and since for each $v \in V \setminus B$:

(24.21)
$$v \in D(G') \iff G'$$
 has an even-length M' -alternating $X' - v$ path $\iff G$ has an even-length M -alternating $X - v$ path $\iff v \in D(G)$.

This proves (24.20), which implies that A(G') = A(G) and C(G') = C(G). By induction, D(G') is the union of the odd components of G' - U. Hence D(G) is the union of the odd components of G - U (since $B \subseteq D(G)$ by (24.20)). Also by induction, $|M'| = \frac{1}{2}(|V'| + |U| - o(G' - U))$. Hence $|M| = \frac{1}{2}(|V| + |U| - o(G - U))$, since |V| - 2|M| = |V'| - 2|M'|.

So U = A(G) is the unique set attaining the minimum in (24.2) for which the union of the odd components of G - U is inclusionwise minimal.

Note that:

(24.22) for any U attaining the minimum in (24.2), each maximum-size matching M has exactly $\lfloor \frac{1}{2} |K| \rfloor$ edges contained in any component K of G - U, and each edge of M intersecting U also intersects some odd component of G - U.

This implies the following. Call a graph G = (V, E) factor-critical if G - v has a perfect matching for each vertex v.

Corollary 24.7a. Let G = (V, E) be a graph. Then each component K of G[D(G)] is factor-critical.

Proof. Directly from Theorem 24.7 and (24.22): if $v \in K$, then $v \in D(G)$, and hence G - v has a maximum-size matching M missing v. By (24.22), M has $\lfloor \frac{1}{2} |K| \rfloor$ edges contained in K. So K - v has a perfect matching.

The Edmonds-Gallai decomposition can be found in polynomial time, since the set D(G) of vertices missed by at least one maximum-size matching can be determined in polynomial time (with the cardinality matching algorithm). In fact,