# Faster Scaling Algorithms for General GraphMatching Problems 

HAROLD N. GABOW<br>University of Colorado, Boulder, Colorado<br>AND<br>ROBERT E. TARJAN<br>Princeton University, Princeton, NJ, and NEC Research Institute, Princeton, NJ


#### Abstract

An algorithm for minimum-cost matching on a general graph with integral edge costs is presented. The algorithm runs in time close to the fastest known bound for maximum-cardinality matching. Specifically, let $n, m$, and $N$ denote the number of vertices, number of edges, and largest magnitude of a cost, respectively. The best known time bound for maximum-cardinality matching is $O(\sqrt{n} m)$. The new algorithm for minimum-cost matching has time bound $O(\sqrt{n \alpha(m, n) \log n} m$ $\log (n N)$ ). A slight modification of the new algorithm finds a maximum-cardinality matching in $O(\sqrt{n} m)$ time. Other applications of the new algorithm are given, including an efficient implementation of Christofides' travelling salesman approximation algorithm and efficient solutions to update problems that require the linear programming duals for matching.


Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Non-numerical Algorithms and Problems-Computations on discrete structures; G.2.2 [Discrete Mathematics]: Graph Theory - graph algorithms, network problems.
General Terms: Algorithms, Design, Theory8GAugmentıng path, blossom, matching, network optimization, scaling.

## 1. Introduction

Finding a minimum-cost matching on a general graph is a classic problem in network optimization. It has many practical applications and very efficient algorithms. We present an algorithm for this problem that is almost as fast as

Harold Gabow's research was supported in part by NSF Grant Nos. DCR-851191, CCR-8815636, and AT\&T Bell Laboratories.
Robert Tarjan's research at Princeton University was supported in part by NSF Grant No. DCR-8605962 and ONR Contract No. N00014-87-K-0467. His research was also supported in part by AT\&T Bell Laboratories.
Authors' addresses: Harold N. Gabow, Department of Computer Science, University of Colorado, Boulder, CO 80309; Robert E. Tarjan, Computer Science Department, Princeton University, Princeton, NJ 08544/NEC Research Institute, Princeton, NJ 08540.
Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission. (c) 1991 ACM 0004-5411/91/1000-0815 $\$ 01.50$
the best known algorithm for the problem without costs, that is, maximumcardinality matching.

In stating resource bounds, we use $n$ and $m$ throughout this paper to denote the number of vertices and the number of edges in the given graph, respectively. When the graph has integral-valued edge costs, $N$ denotes the largest magnitude of a cost.

The fastest known algorithm for maximum-cardinality matching is due to Micali and Vazirani [23] and runs in time $O(\sqrt{n} m)$ (see also [14]). Edmonds gave the first polynomial algorithm for weighted matching [7]. At the time of the initial writing of this paper, the best known implementation of this algorithm was that of [13]; it runs in time $O\left(n\left(m \log \log \log _{d} n+n \log n\right)\right)$. where $d=\max \{m / n, 2\}$ is the density of the graph. Recently an implementation achieving time $O(n(m \mp n \log n))$ has been given [12]. These bounds for weighted matching can be substantially improved under the assumption of integral costs that are not huge: The scaling algorithm of [10] runs in time $O\left(n^{3 / 4} m \log N\right)$. We improve this last bound to $O(\sqrt{n \alpha(m, n) \log n} m \log (n N))$. We also show that for maximum-cardinality matching our algorithm runs in the same time as the algorithm of Micali and Vazirani cited above. We present two other applications: We show how to speed up Christofides' travelling salesman approximation algorithm [4] to time $O\left(n^{25}(\log n)^{1.5}\right)$; this bound is independent of the size of the input numbers. We also show how to find the linear programming dual variables for matching, which are the basis of Edmonds' algorithm. This gives efficient solutions to various matching update problems. Some more recent applications of our algorithm are mentioned in the last section.

Our algorithm is based on the approach to scaling introduced by Goldberg and Tarjan for the minimum-cost flow problem [18, 19] and applied in [16] to the assignment problem. Historically, the first approaches to scaling computed an optimum solution at each of $\log N$ scales (e.g., [7a], [20], [9, 10]). The new method computes an approximate optimum solution at each of $\log (n N)$ scales; using $\log n$ extra scales ensures that the last approximate optimum is exact. The notion of $\epsilon$-optimality $[3,25]$ turns out to be the appropriate definition of '"approximate optimum' for this scaling technique.

Applying this scaling technique to general graphs is difficult because of "blossoms." In all of the scaling algorithms mentioned above for bipartite and directed graphs, the solution to one scale gives an obvious starting point for the solution to the next. Blossoms invalidate the obvious starting point. The techniques of [10], including the notion of "shells," are used to overcome this difficulty. Nonetheless, blossoms slow our algorithm down: The algorithm of [16] finds a minimum-cost matching on a bipartite graph in time $O(\sqrt{n} m$ $\log (n N)$ ). The extra factor of $\sqrt{\log n}$ in our bound for general matching comes from errors introduced in finding the starting point; the extra factor of $\sqrt{\alpha(m, n)}$ comes from data structures for blossom manipulation.

The paper is organized as follows. Section 1.1 reviews Edmonds' weighted matching algorithm [7]; many of the ideas and routines of this algorithm are incorporated into ours. The rest of the paper presents our algorithm in a top-down fashion. Section 2 gives the main routine, Sections 3-4 give lower level subroutines. These sections also show that the algorithm is correct and give parts of the efficiency analysis. Sections 5-6 essentially complete the
efficiency analysis. Sections 7-8 give the remaining lower-level details of the algorithm. Section 9 concludes the analysis of the algorithm. Section 10 applies the algorithm to other matching problems such as minimum perfect matching. Section 11 mentions a few further applications of the algorithm.

The present section closes with notation and definitions. We use several standard mathematical conventions to simplify the efficiency analysis. Background concerning matching can be found in greater detail in [21], [22], and [27].
If $S$ is a set and $e$ an element, $S+e$ denotes $S \cup\{e\}$ and $S-e$ denotes $S-\{e\}$. For integers $i$ and $j,[i \ldots j]=\{k \mid k$ is an integer, $i \leq k \leq j\}$. The function $\log n$ denotes logarithm to the base two.
We use a hat, for example, $\hat{f}$, to emphasize that an object is a function, when necessary. We use a dot, $\cdot$, to denote the argument of a function. For example, if $f$ is a function of two variables, $f(x, \cdot)$ denotes the function of one variable mapping $y$ to $f(x, y)$. If $f$ and $g$ are real-valued functions then $f+g$ and $f g$ denote their sum and product, respectively, that is, $(f+g)(x)=f(x)+g(x)$, $f g(x)=f(x) \times g(x)$. We use the following conventions to sum the values of a function: If $f$ is a real-valued function whose domain includes the set $S$, then $f(S)=\sum\{f(s) \mid s \in S\}$. Similarly if $f$ has two arguments then $f(S . T)=$ $\sum\{f(s, t) \mid s \in S, t \in T\}$, for $S \times T$ a subset of the domain of $f$.

For any graph $G, V(G)$, and $E(G)$ denote the vertex set and edge set of $G$, respectively. All graphs in this paper are undirected. We regard an edge as being a set of two vertices; hence, a statement like $e \subseteq S$, for $e$ an edge and $S$ a set of vertices, means that both vertices of $e$ are in $S$. We usually denote the edge joining vertices $v$ and $w$ by $v w$. Thus if $e=\nu w$ and $y: E(G) \rightarrow \mathbf{R}$ then $y(e)=y(v)+y(w)$ by our convention for functions. We often identify a subgraph $H$, such as a path or tree, with its set of vertices $V(H)$ or edges $E(H)$. For example $H \subseteq S$ is short for $V(H) \subseteq S$ or $E(H) \subseteq S$, depending on whether $S$ is a set of vertices or edges; the choice will be clear from the context.
A matching on a graph is a set of vertex-disjoint edges. Thus a vertex $v$ is in at most one matched edge $v v^{\prime}$ : a free vertex is in no such edge. A perfect matching has no free vertices. An alternating path (cycle) for a matching is a simple path (cycle) whose edges are alternately matched and unmatched. An augmenting path $P$ is an alternating path joining two distinct free vertices. To augment the matching along $P$ means to enlarge the matching $M$ to $M \oplus P$. thus giving a matching with one more edge ( $M \oplus P$ denotes the symmetric difference of $M$ and $P$ ).

Suppose $c: E \rightarrow \mathbf{R}$ is a function that assigns a numeric cost to each edge; in this paper costs are integers in $[-N \cdots N]$ unless stated otherwise. By our convention the cost $c(S)$ of a set of edges $S$ is the sum of the individual edge costs. A minimum- (maximum-) cost matching is a matching of smallest (largest) possible cost. A minimum (maximum) perfect matching is a perfect matching of smallest (largest) possible cost.
1.1 Edmonds' Minimum Critical Matching Algorithm. It is convenient to work with a variant of the matching problem defined as follows. Let $G$ be a graph and $v$ a fixed vertex. A $v$-matching is a perfect matching on $G-v$. Figure 1 shows an $x$-matching; in all figures of this paper matched edges are drawn wavy and free vertices are drawn square. A minimum (maximum)


Fig. 1. Blossom with base vertex $x$
$v$-matching is a $v$-matching with minimum (maximum) possible cost. $G$ is critical if every vertex $v$ has a $v$-matching. The minimum critical matching problem is, given a critical graph, to find a minimum $v$-matching for each vertex $v$ [11]. It follows from [6] that all the desired matchings can be represented by the blossom tree defined below; we take this blossom tree as a solution to the critical matching problem.
Note that if $G$ is a graph with a perfect matching, a critical graph is obtained by adding a vertex adjacent to every vertex of $V(G)$. Hence, an algorithm for minimum critical matching can be used to find a minimum perfect matching.

Edmonds' algorithm is based on the notion of blossom, which is explained in the next four paragraphs. Let $G$ be a graph with a matching. A blossom forest $F$ is a forest with the following properties. (Figure 2 shows a blossom forest, with just one tree, for Figure 1.) The number of children of any nonleaf node of $F$ is at least three and odd. Each node of $F$ is identified with a subgraph of $G$, as follows: The leaves of $F$ are precisely the vertices of $G$. If $B$ is a nonleaf node, its children can be ordered as $B_{t}, i=1, \ldots, k$ so that

$$
V(B)=\bigcup_{i=1}^{k} V\left(B_{i}\right), \quad E(B)=\bigcup_{i=1}^{k}\left(E\left(B_{i}\right)+e_{i}\right),
$$

where $e_{l}$ is an edge that joins a vertex of $V\left(B_{i}\right)$ to a vertex of $V\left(B_{i+1}\right)$ (interpret $B_{k+1}$ as $B_{1}$ ); furthermore $e_{i}$ is matched precisely when $i$ is even. (Thus, each child of $B$ is incident to two edges $e_{i}$; for $B_{1}$ both edges are unmatched, and for all other children one edge is matched and the other unmatched; there are precisely two possible orderings for the children.) In this


Fig. 2. Blossom tree.
paper, node always refers to an element of $V(F)$ and vertex always refers to an element of $V(G)$.
Each node $B$ of $F$ is a blossom. (Thus, a blossom can also be regarded as a subgraph.) The blossom edges of $B$ are the above edges $e_{i}, i=1, \ldots, k$. Any root, that is, maximal blossom, is a root blossom; all other blossoms are nonroot blossoms. Every vertex is a blossom; a blossom that is not a vertex is a nonleaf blossom.

The subgraph induced by $V(B)$ is denoted $G(B)$. Define functions

$$
\hat{n}(B)=|V(B)|, \quad \hat{m}(B)=|E(G(B))|
$$

We emphasize that a blossom is not defined as an induced subgraph, e.g., $\hat{m}(B)$ is in general larger than $|E(B)|$. A simple induction shows that $\hat{n}(B)$ is odd. The base vertex of $B$ is the unique vertex of $B$ not in a matched edge of $E(G(B))$. The base of a vertex $v$ is $v$; a simple induction shows that the base of a nonleaf blossom $B$ exists and is in the first child blossom $B_{1}$ of $B$.

Any $v$-matching of a critical graph has a blossom forest that consists of one tree $T^{*}$ called a blossom tree. (This can be proved by examining the algorithm of [6].) The root of $T^{*}$ is a blossom having vertex set $V(G)$ and is denoted $G^{*}$. Given $T^{*}$, for any vertex $w$ a $w$-matching of $G$ can be found in time $O(n)$. We now describe a recursive procedure to do this. The procedure is blossom_match $(B, w)$; here, $B$ is a nonleaf node of $T^{*}, w$ is a vertex of $B$, and the procedure constructs a $w$-matching of $B$. To do this, let $B$ have children $B_{i}$ and blossom edges $e_{t}, i=1, \ldots, k$; as above, $e_{t}$ joins $B_{i}$ to $B_{i+1}$. Let $w \in B_{j}$. Match alternate edges of the list $e_{i}$, keeping $e_{j-1}$ and $e_{j}$ unmatched. For $i \neq j$ let $w_{i}$ denote the vertex of $B_{i}$ on a newly matched edge, and let $w_{j}=w$. Complete the procedure by recursively executing blossom $\operatorname{match}\left(B_{i}, w_{i}\right)$ for each nonleaf child $B_{i}$. It is easy to see that blossom$\operatorname{match}\left(T^{*}, w\right)$ constructs the desired $w$-matching in time $O(n)$.

Now we review Edmonds' algorithm for minimum critical matching; further details can be found in [7]. Two functions $y, z$ form (a pair of) dual functions if $y: V(G) \rightarrow \mathbf{R}, z: 2^{V(G)} \rightarrow \mathbf{R}$ and $z(B) \geq 0$ unless $B=V(G)$. Such a pair
determines a dual edge function $y z: E \rightarrow \mathbf{R}$, which for an edge $e$ is defined as

$$
y z(e)=y(e)-z(\{B \mid e \subseteq B\})
$$

(Recall that by convention if $e=v w$ then $y(e)=y(v)+y(w)$. Similarly, the last term denotes $\sum\{z(B) \mid e \subseteq B\}$.) The duals are dominated on edge $e$ if

$$
y z(e) \leq c(e)
$$

they are tight if equality holds.
Edmonds' algorithm maintains a structured matching. This is a matching plus corresponding blossom forest $F$ plus dual functions that collectively satisfy two conditions: (i) $z$ is nonzero only on nonleaf blossoms of $F$. (ii) The duals are dominated on every edge, and tight on every edge that is matched or a blossom edge. A structured $v$-matching is a minimum $v$-matching. (This can be proved by an argument similar to Lemma 2.1(a) below.) Regarding (i), define a weighted blossom as a blossom with a nonzero dual.
An optimum structured matching is a structured $v$-matching for some vertex $v$, whose blossom forest is a blossom tree $T^{*}$. Given $T^{*}$, for any vertex $w$ a minimum $w$-matching is found in $O(n)$ time by the blossom_match procedure. The output of Edmonds' algorithm is an optimum structured matching. Thus, Edmonds' algorithm solves the minimum critical matching problem.

The input to Edmonds' algorithm is a critical graph plus a structured matching. (The structured matching can be the empty matching, a blossom forest of isolated vertices, and dual functions $z=0$ and $y$ small enough to be dominated on every edge.) The algorithm repeatedly does a "search" followed by an "augment step" until some search halts with a $v$-matching ( $v$ arbitrary) and a blossom tree (not forest). (This is a slight difference from the way the algorithm of [7] halts; see below.)
More precisely a search builds a search graph $\mathscr{H}$, defined as follows and illustrated in Figure 3. $V(\mathscr{F})$ is partitioned into root blossoms $B$. $E(\mathscr{P})$ consists of the blossom edges $E(B)$ plus other tight edges. The rest of the description of $\mathscr{F}$ depends on whether or not an augmenting path has been found. First consider $\mathscr{I}$ before an augmenting path has been found. If each root blossom of $\mathscr{H}$ is contracted to a vertex, $\mathscr{y}$ becomes a forest $\mathscr{F}$ ( $\mathscr{y}$ should not be confused with the blossom forest). The roots of $\mathscr{F}$ are precisely the root blossoms of $G$ that contain a free vertex. The path $P(B)$ from a root blossom $B$ to the root of its tree in $\mathscr{F}$ is alternating. Blossom $B$ is outer (inner) if $P(B)$ has even (odd) length. Any descendant, in the blossom forest, of an outer (inner) root blossom is also called outer (inner). (In particular, every free vertex of $G$ is outer.) Any outer vertex $v$ is joined to a free vertex by an even-length alternating path $P(v) \subseteq E(\mathscr{Y})$; if $v$ is in root blossom $B$ then $P(v)$ contains $P(B)$. For instance Figure 3 shows path $P(a)$ for vertex $a$ of Figure 1.

Now consider $\mathscr{G}$ when an augmenting path has been found. In this case, $\mathscr{H}$ contains one or more tight edges $v w$ joining outer vertices in distinct trees of $\mathscr{F}$. Each such edge gives an augmenting path composed of $v w$ plus the above paths $P(v), P(w)$.

The search builds $\mathscr{I}$ using three types of steps. A grow step enlarges $\mathscr{S}$ by adding a tight edge $e$ that was incident to $\mathscr{F}$; the root blossom $B$ at the end of $e$ is also added to $\mathscr{S}$. Grow steps always occur in pairs in Edmonds' algorithm:


Fig. 3 Search graph in Edmonds* algorithm.
first an unmatched edge $e$ is added, along with the above blossom $B$; then, the matched edge incident to $B$ is added. Figure 4 shows a grow step for the unmatched edge $a b$ followed by a grow step for the matched edge $c d$.
A blossom step enlarges ' $y$ ' by adding a tight edge that joins distinct outer root blossoms of $\mathscr{F}$. This step either constructs a new blossom in ${ }^{\mathscr{F}}$, or it discovers that $\mathscr{y}$ contains an augmenting path. In Figure 3 an edge $a e$ would give a blossom step that constructs a new blossom, possibly the one in Figure 1.
An expand step deletes an unweighted root blossom $B$ from the blossom forest, thus making its children into roots; $B$ is also deleted from $y$ and replaced by some of these children (not necessarily all of them) so that $\overline{\mathcal{F}}$ remains a forest. Figure 5 shows an expand step; blossoms $B_{1}, \ldots, B_{5}$ become root blossoms, and $B_{4}$ and $B_{5}$ leave $\mathscr{H}^{\prime}$.

These three steps are repeated as many times as possible, giving a maximal search graph $!f$. If the maximal $f^{\prime}$ does not contain an augmenting path and $G$ is not a blossom, a dual adjustment is done. It starts by computing a quantity $\delta$, as described below. Then, it makes the following changes:

$$
\begin{array}{ll}
y(v) \leftarrow y(v)+\delta, & \text { for each outer vertex } v ; \\
y(v) \leftarrow y(v)-\delta, & \text { for each inner vertex } v ; \\
z(B) \leftarrow z(B)+2 \delta, & \text { for each nonleaf root outer blossom } B ; \\
z(B) \leftarrow z(B)-2 \delta, & \text { for each nonleaf root inner blossom } B .
\end{array}
$$

These assignments do not change the value of $y z(e)$ for $e \in E(\mathscr{F})$, so these edges remain tight. The assignments increase $y z(e)$ only if $e$ joins an outer vertex to a vertex not in $\mathscr{T}$, or if $e$ joins two distinct root outer blossoms. Thus the adjustment maintains condition (ii) above and also allows a new grow,


Frg. 4. Grow steps in Edmonds’ algorithm


FIG. 5. Expand step.
blossom or expand step to be executed, if $\delta$ is chosen as $\min \left\{\delta_{g}, \delta_{b}, \delta_{e}\right\}$, where

$$
\begin{aligned}
& \delta_{g}=\min \{(c-y)(v w) \mid v w \in E(G), v \text { an outer vertex, } w \notin V(\mathscr{S})\} ; \\
& \delta_{b}=\min \left\{\left.\frac{(c-y)(e)}{2} \right\rvert\, e \text { an edge joining two distinct root outer blossoms }\right\} ; \\
& \delta_{e}=\min \left\{\left.\frac{z(B)}{2} \right\rvert\, B \text { a nonleaf root inner blossom }\right\} .
\end{aligned}
$$

Note that $\delta>0$ by the maximality of $\mathscr{G}$. If $\delta=\delta_{g}$, a grow step can be executed after the dual adjustment; similarly, $\delta=\delta_{b}$ gives a blossom step and $\delta=\delta_{e}$ gives an expand step.

After the dual adjustment the search continues to do grow, blossom and expand steps. Eventually the search halts, in one of two ways. Every search but the last finds a weighted augmenting path. This is an augmenting path $P$ whose edges are tight. The augment step enlarges the matching $M$ by one edge to $M \oplus P$. The blossom forest and duals remain valid. Then the algorithm continues with the next search. In the last search the matching is a $v$-matching. The last search eventually absorbs the entire graph $G$ into one blossom. At this time, the algorithm halts with the desired optimum structured matching.

We note two more properties of Edmonds' algorithm for use below. First, any dual adjustment increases the $y$ value of any free vertex by $\delta$. Second, note that a dual adjustment step does divisions by two to calculate $\delta_{b}, \delta_{e}$. If all given costs are even integers, and all $y$ values are initialized to have the same parity, then all quantities computed by the algorithm are integers [24, p. 267, ex. 3]. This integrality property motivates various details of the scaling algorithm. For instance, the scaling algorithm keeps all edge costs even.

It will be helpful to sketch a proof of the above integrality property. Suppose that all given costs are even integers and all $y$ values are initialized to integers of the same parity. The algorithm maintains the invariant that all $y$ values are integral, all $y$ values of free vertices have the same parity, and all $z$ values are even integers. The main step of the proof is to observe that all $y$ values of vertices in $\mathscr{S}$ have the same parity. This follows since $\mathscr{F}$ consists of tight edges and each connected component contains a free vertex. Using this observation it is easy to see that the invariant is preserved.

## 2. The Matching Algorithm: The Scaling Routine

This section gives the overall structure of our matching algorithm. This algorithm solves the minimum critical matching problem in $O(\sqrt{n \alpha(m, n) \log n}$ $m \log (n N)$ ) time. This section describes the main routine of the scaling algorithm for minimum critical matching. The input is a critical graph. (This entails no loss of generality - the algorithm can detect input graphs that are not critical, as indicated below.)

The algorithm works by scaling the costs. Each scale finds a $v$-matching, for some $v$, whose cost is almost minimum in the following sense. A 2-optimum matching is a $v$-matching $M_{v}$, for some vertex $v$, plus a blossom tree $T$, plus dual functions $y, z$ such that $z$ is nonzero only on nonleaves of $T$ and the
following constraints hold:

$$
\begin{array}{ll}
y z(e) \leq c(e), & \text { for } \quad e \in E \\
y z(e) \geq c(e)-2, & \text { for } \quad e \in M \cup \bigcup\{E(B) \mid B \in V(T)\} \tag{1b}
\end{array}
$$

Note that if this definition is satisfied for some vertex $v$, it is satisfied by every vertex $x$ (matching $M_{x}$ is constructed in $O(n)$ time by the blossom_match procedure). Hence, when $M_{x}$ denotes a 2-optimum matching, we understand that $x$ can be chosen arbitrarily.

To motivate this definition, first observe that dropping the -2 term from (1b) gives the dominated and tight conditions used in Edmonds' algorithm. The $\mathbf{- 2}$ term is included so that the algorithm augments along paths of short length. This makes the algorithm efficient. Further motivation is given in [16]. (Actually the bipartite matching algorithm of [16] uses a term of magnitude 1 rather than 2 , and also maintains equality in the constraint for matched edges. Here, we use magnitude 2 because of the aforementioned considerations of integrality. Also equality cannot be maintained on the matched edges, because of details of blossom manipulation.) The following result is the analog of Edmonds' optimality condition.

## Lemma 2.1. Let $M_{x}$ be a 2-optimum matching.

(a) For any vertex $x$, any $x$-matching $X$ has $c(X)>c\left(M_{x}\right)-n$.
(b) If each cost $c(e)$ is a multiple of some integer $k, k \geq n$, then $M_{x}$ is $a$ minimum $x$-matching.

## Proof

(a) Consider any vertex $x$ and let $T$ be the blossom tree. Function $z$ is nonzero only on nonleaves of $T$. For any blossom $B, M_{x}$ contains precisely $\lfloor\hat{n}(B) / 2\rfloor$ edges of $G(B)$, and no matching contains more. Combining these facts with (1a) and (1b) gives

$$
\begin{aligned}
c\left(M_{x}\right) & \leq 2\lfloor n / 2\rfloor+y z\left(M_{x}\right) \\
& \leq(n-1)+y(V(G)-x)-\lfloor\hat{n} / 2\rfloor z(V(T)) \leq c(X)+n-1
\end{aligned}
$$

(Recall that by the conventions of Section $1,\lfloor\hat{n} / 2\rfloor z(V(T)$ ) denotes $\left.\sum\{\lfloor\hat{n}(B) / 2\rfloor z(B) \mid B \in V(T)\}.\right)$
(b) This follows from (a) and the fact that any matching has cost a multiple of $k$.
Now we describe the scaling routine, the main routine of the algorithm. It scales the costs in the following manner. The algorithm always works with even edge costs to preserve integrality. Initially it computes a new cost function $\bar{c}$ equal to $n+1$ times the given cost function. (Clearly the two cost functions have the same minimum $x$-matchings. Also each $\bar{c}(e)$ is even.) It maintains a cost function $c$ equal to $\bar{c}$ in the current scale. More precisely define $k$ $=\lfloor\log (n+1) N \mid+1$, the greatest number of bits in the binary expansion of a $\bar{c}$ cost. For any $s \in[1 \cdots k]$ define a function $b_{s}:[-(n+1) N \cdots(n+$ 1) $N] \rightarrow\{-1,0,1\}$ by taking $b_{s}(i)$ as the $s$ th signed bit in the expansion of $i$ as a $k$-bit binary number. Thus, any new cost value $i=\bar{c}(e)$ equals $\sum_{t=1}^{k-1} b_{t}(i) 2^{k-t}$ (since $b_{k}(i)=0$ ). In the $s$ th scale, the cost $c(e)$ is taken to be $\sum_{t=1}^{s=1} b_{t}(i) 2^{s+1-t}$, so in the $(k-1)$ st scale $c(e)=i$.

The scaling routine computes $\bar{c}$, then initializes $c, y$ and $z$ to the zero function, the matching $M_{x}$ to $\varnothing$, and the blossom tree $T$ to a root $G$ with children $V(G)$. Then, it executes the following loop for index $s$ going from 1 to $k-1$ :

Double Step. Compute new functions $c \leftarrow 2\left(c+b_{s}\right), y \leftarrow 2 y-1, z \leftarrow 2 z$.
Match Step. Call the match routine to find a 2 -optimum matching $M_{x}$, with new duals $y, z$ and new blossom tree $T$.
Lemma 2.1(b) implies that if match works as described in the Match Step, the scaling routine solves the minimum critical matching problem, that is, each final matching $M_{x}$ is a minimum $x$-matching. Each iteration of the loop is called a scale. We give a match routine that runs in $O(\sqrt{n \alpha(m, n) \log n} m)$ time, thereby achieving the desired time bound.
Note that in the first scale the tree $T$ computed in the initialization is not necessarily a blossom tree, since it need not correspond to a blossom structure. We shall see that the algorithm still works correctly, because $z=0$ (see the last paragraph of Section 4).

## 3. The Match Routine

This section describes the overall structure of the routine that finds a 2 -optimum matching in a scale.

On entry to match, $y, z$ are duals computed in the Double Step, and $T$ is the blossom tree of the previous scale (or the initialization, in the first scale). The match routine saves $T$ as the tree $T_{0}$; for the analysis it is convenient to let $y_{0}, z_{0}$ refer to the duals on entry to match.
The scaling routine is similar to the main routine of the bipartite matching algorithm of [16]. In the bipartite algorithm, each scale is similar to the first in that the dual function can be initially taken as zero, and there is no structure on the graph inherited from previous scales. This is not true for general graphs: The function $z_{0}$ can have positive values that cannot be eliminated (see [10]). The match routine is forced to work with blossoms from both the previous scale, in blossom tree $T_{0}$, and the current scale. It is convenient to denote the current blossom forest as $T$ (eventually this forest becomes a blossom tree). An old blossom is a node of $V\left(T_{0}\right)$; a current blossom is a node of $V(T)$. An old blossom $B$ dissolves either when it becomes a current blossom or, if $B \neq G^{*}$, when $z(B)$ becomes zero. Note that current blossoms do not dissolve (in the current scale); hence, we use the term undissolved blossom to refer to an old blossom that has not yet dissolved. A vertex is a current blossom, so only nonleaf blossoms are undissolved. Finally, note that the old matching is implicitly discarded in the Double Step, so "the matching'" refers to the current matching.

The match routine maintains inequalities (1), with $z$ nonzero only on nonleaves of $V(T) \cup V\left(T_{0}\right)$. In (1b) $T$ is the current blossom forest. Observe that both current and old blossoms contribute to the $z$ term in the definition of $y z(e)$. When all old blossoms are dissolved, the matching is 2-optimum and the routine can halt. The reason is that by definition the old blossom $G^{*}$ can dissolve only by becoming a current blossom. When this occurs we have a $v$-matching, a blossom tree, and a function $z$ that is nonzero only on nonleaves of $T$.

Note that after the Double Step, (1a) holds for all edges and (1b) is vacuous. Hence, the Double Step maintains (1) as desired. To help preserve (1a) the match routine also maintains

$$
\begin{equation*}
y \leq y_{0} . \tag{2}
\end{equation*}
$$

In a blossom tree, define the major child $C$ of a node $B$ as a child with largest size $\hat{n}(C)$; a tie for the major child is broken arbitrarily. Hence, any nonleaf has exactly one major child, and any nonmajor child $D$ of $B$ has $\hat{n}(D)<\hat{n}(B) / 2$. A major path is a maximal path in which each node is followed by its major child. The major paths partition the nodes of a blossom tree. (These are essentially the "heavy paths" of [26].) A major path starting at node $R$ is denoted by $P(R)$ and has major path root $R$. Define the rank of any node $B$ as $\lfloor\log \hat{n}(B)\rfloor$. A nonmajor child of a node $B$ has rank less than $B$; a nonmajor child of a node in $P(R)$ has rank less than $R$. In Figure 2 the path from the root to leaf $z$ is the major path $P\left(G^{*}\right)$; the root has rank $[\log 21]=4$.
procedure match.
Initialize the matching to $\varnothing$ and $T$ to the forest having every vertex of $G$ a root. Traverse the major path roots $R$ of $T_{0}$ in postorder. At each root $R$ call a routine path( $R$ ) to dissolve the old blossoms on $P(R)$, while maintaining (1)-(2).

This routine is correct since after match processes root $R=G^{*}$, all old blossoms are dissolved. Thus, the matching is 2 -optimum as observed above. Note that for any major root $R$, on entry to $\operatorname{path}(R)$ all descendants of $R$ have dissolved except those on $P(R)$. Figure 6 illustrates this situation: Suppose the graph of Figure 1 has old blossom tree given by Figure 2. Then, on entry to path $\left(G^{*}\right)$, all blossoms are dissolved except those shown on $P\left(G^{*}\right)$.

Lemma 3.1. If the time for path $(R)$ is $O(\sqrt{\hat{n}(R) \alpha(m, n) \log n} \hat{m}(R))$, then the time for match is $O(\sqrt{n \alpha(m, n) \log n} m)$.

Proof. For any integer $0 \leq r \leq \log n$, consider the major path roots of rank $r$. For any vertex $v \in V(G)$, at most one of these roots $R$ has $v \in$ $V(R)$. Hence any edge (of $E(G)$ ) is in at most one of the subgraphs $G(R)$. Thus for some constant $c$ the time spent on these roots is at most $c \sqrt{2^{r+1} \alpha(m, n) \log n} m$. Summing over all ranks $r$ gives the desired bound.

## 4. Shells and the Path Routine

This section presents the path routine and its main subroutine shell_search. These routines are based on the concept of a shell [10].

If $C$ and $D$ are blossoms with $V(D) \subseteq V(C)$, the shell $G(C, D)$ is the subgraph induced by $V(C)-V(D) . C$ is the outer boundary, $D$ the inner boundary. As a special case, we allow $D=\varnothing$. Extend the function $\hat{n}$ to shells: $\hat{n}(C, D)$ is the number of vertices in a shell $G(C, D)$. A shell is even if $\hat{n}(C, D)$ is even, or equivalently $D \neq \varnothing$; otherwise, it is odd. Figure 6 indicates the even shell $G\left(G^{*}, A\right)$.

We use a number of functions of shells, such as the above $\hat{n}$. We define such functions by using the shell boundaries as arguments, as in the above $\hat{n}(C, D)$. Alternatively, if $X$ denotes a shell, we use $X$ as the argument, e.g., $\hat{n}(X)$. If


FIG. 6. Major path with dissolved blossoms.
$X$ denotes an old blossom, it corresponds to an odd shell, and we write $\hat{n}(X)$ as a shorthand for $\hat{n}(X, \varnothing)$. Which of the two interpretations of $\hat{n}(X)$ is appropriate will always be clear from context.
This paper only refers to shells of $P(R)$, which are shells $G(C, D)$ with $C, D$ on $P(R)$ for the major path root $R$ ( $D$ may be empty). At any time in path $(R)$, if $C$ and $D$ are currently consecutive undissolved blossoms in $P(R)$. then $G(C, D)$ is an undissolved shell (of $P(R)$ ). (An undissolved odd shell has $C$ the currently innermost undissolved blossom of $P(R)$.) The path $(R)$ routine works with undissolved shells. Obviously these shells change as blossoms dissolve.
The $\operatorname{path}(R)$ routine works like the bipartite matching algorithm of [16] in the sense that it finds most augmenting paths quickly and finds the remaining paths at a slower and slower rate. The bipartite algorithm accomplishes this automatically, that is, the algorithm is unchanging, only its performance changes. For general graphs, it seems that some lower-level details of the algorithm must change as the execution progresses. For this reason, we organize the path routine in "phases." More precisely, the phase is defined in terms of a parameter $\rho$ whose value is chosen below (Section 5). Also define $R^{\prime}$ to be the largest undissolved blossom of $P(R)$. (Initially $R^{\prime}=R$ if $R$ is a weighted blossom; $R^{\prime}$ shrinks as the algorithm progresses.) The path routine is a loop. Routine path is in phase 1 during the first $\rho$ iterations of the loop. After that, it is in phase 2 if $R^{\prime}$ has more than one free vertex, and phase 3 otherwise. (Hence, in phase $3, R^{\prime}$ has exactly one free vertex.) It will be apparent that path can go through any sequence of phases that starts with phase 1 and never decreases, that is, $1 ; 1,2 ; 1,2,3$; or 1,3 . We shall also see that an individual iteration consumes more time in a higher phase than in a lower.

The path $(R)$ routine augments the matching along paths of "eligible" edges; it finds these paths by constructing a search graph of eligible edges. Edge $e$ is defined to be eligible if its vertices are in the same undissolved shell of $P(R)$;
furthermore a condition that depends on the phase is satisfied. In phase 1 the condition is that one of these alternatives holds:
(i) $e$ is unmatched and $y z(e)=c(e)$;
(ii) $e$ is matched and $y z(e)=c(e)-2$;
(iii) $e$ is a current blossom edge.

In phase 2 or 3 the condition is $\hat{y z}(e)=c(e)$ or $c(e)-2$. Note that this is always the case if any of (i)-(iii) hold.

Observe a simple consequence of the definition of eligibility: Throughout $\operatorname{path}(R)$, any matched edge or current blossom edge has both ends in the same undissolved shell (recall that the matched or current blossom edges are precisely the edges referred to in inequality (1b)). This observation follows because an edge must be eligible in order to be on an augmenting path and become matched or to enter a search graph and become a blossom edge; an eligible edge has both ends in the same undissolved shell of $P(R)$; the latter property continues to hold even as blossoms subsequently dissolve.

Here is the path routine. It uses a routine shell_search( $S$ ) whose argument is a shell $S$. As above, $R^{\prime}$ denotes the largest undissolved blossom in $P(R)$.
procedure $\operatorname{path}(R)$.
Repeat the following steps until all old blossoms of $P(R)$ are dissolved:
Augment Step. Construct an auxilary graph $H$ from $G\left(R^{\prime}\right)$ by contracting every current root blossom in $R^{\prime}$ and keeping only the eligible edges of $G\left(R^{\prime}\right)$. Find a maximal set $p$ of vertex-disjoint augmenting paths in $H$. For each path of $: \bar{\mu}$, augment along the corresponding path in $G$

Sort Step. Order the undissolved shells of $P(R)$ that contain a free vertex as $S_{2}, i=1, \ldots, k$ so that $\hat{n}\left(S_{t}\right)$ is nonincreasing.

Search Step. For $i=1$ to $k$, if both boundaries of $S_{1}$ are still undissolved call shell__search $\left(S_{1}\right)$ to adjust duals and possibly find an augmenting path of elıgible edges (more details of shell _-search are given below).

The path routine is implemented as follows: It exits immediately if $R$ is a leaf blossom (any vertex is dissolved). Otherwise, the Augment Step finds augmenting paths by doing a depth-first search on $H$. The details of this search are unimportant for the analysis and so are postponed until Section 8. It suffices to note that the Augment Step uses time $O(\hat{m}(R))$.

In the Search Step, the shell_search routine is a search of Edmonds' algorithm modified in three ways: (i) to use eligibility rather than tightness; (ii) to take old blossoms into account; (iii) to change the halting procedure. We discuss each of these in turn.

For (i), note that eligibility plays the role of tightness in Edmonds' algorithm: shell search adds an edge to the search graph $\mathscr{F}^{\prime}$ only when it is eligible. However, a matched edge need not be eligible. (This occurs only in phase 1.) Thus, a grow step may not be done for a matched edge incident to $\mathscr{f}$; also a blossom step may be done when a matched edge is scanned. (For example, in Figure 4, the matched edge $c d$ may be added to $\mathscr{F}$ in a grow step that does not immediately follow the one for $a b$.) This contrasts with Edmonds' algorithm, where a matched edge is always tight, a grow step is always done for a matched
edge incident to $\mathscr{I}$, and a blossom step is done only when an unmatched edge is scanned. These changes to search are straightforward.
We turn to (ii). Consider an undissolved blossom B. To translate B by $\delta$ means to perform the following assignments:

$$
\begin{aligned}
& y(v) \leftarrow y(v)-\delta, \quad \text { for } \quad \text { each } \quad v \in V(B) ; \\
& z(B) \leftarrow z(B)-2 \delta .
\end{aligned}
$$

(For example, a dual adjustment in Edmonds' algorithm translates inner root blossoms.) Observe that translating an undissolved blossom $B$ cannot increase a quantity $y z(e)$, and it maintains inequalities (1) ((1b) holds since as observed above, no matched or current blossom edge has exactly one vertex in the undissolved blossom $B$ ).
Consider an undissolved shell $G(C, D)$ containing a free vertex. The routine shell_search $(C, D)$ executes a search of Edmonds' algorithm on $G(C, D)$, modified to translate $C$ and $D$. More precisely, when search does a dual adjustment, it calculates $\delta=\min \left\{\delta_{g}, \delta_{b}, \delta_{e}, \delta_{d}\right\}$, where the first three quantities correspond to the calculation in Edmonds' algorithm (Section 1.1) and $\delta_{d}=\min \{z(C) / 2, z(D) / 2 \mid D \neq \varnothing\}$. In the dual adjustment, shell_search translates $C$ by $\delta$ and also translates $D$ by $\delta$ (if $D \neq \varnothing$ ). Shell__search does a dissolve step if $\delta=\delta_{d}$, that is, the translation dissolves $C$ or $D$ (or both). Dissolve enlarges the shell to $G\left(C^{\prime}, D^{\prime}\right)$, where $C^{\prime}$ is the smallest undissolved blossom containing $C$ and $D^{\prime}$ is the largest undissolved blossom contained in $D$. (Possibly $C^{\prime}=C$ or $D^{\prime}=D$. but not both. If $C^{\prime}$ does not exist, the search halts, as discussed with (iii) below.) Any free vertex that gets added to the shell is immediately added to $\mathscr{F}$ as an outer vertex. After the dissolve step, the search continues, now working on the enlarged shell $G\left(C^{\prime}, D^{\prime}\right)$.
We now verify that translating $C$ and $D$ ensures that inequalities (2) and (1) are preserved. Translating $C$ ensures (2) is preserved, since search increases a $y$-value by at most $\delta$. Consider (1), for any edge $e$ of $G$. If $e$ is in the shell, then (1) follows easily from Edmonds' algorithm and the fact that the translation does not change $y z(e)$. For other edges $e$, only (1a) applies, so it suffices to show that $y z(e)$ does not increase. If neither end of $e$ is in $C$ or both ends of $e$ are in $D$, then obviously $y z(e)$ does not change. If precisely one end of $e$ is in $C$, then (2) implies that $y z(e)$ does not increase. The remaining case is one end $v$ in $V(D)$ and the other end $w$ in $V(C)-V(D)$. The translations decrease $y(v)$ by $2 \delta$ and $z(C)$ by $2 \delta ; y(w)$ increases by at most $\delta$ in Edmonds' search and decreases by $\delta$ in the translations; thus, $y z(e)$ does not increase.

Finally, we discuss (iii). Recall that a search of Edmonds' algorithm halts either when it finds an augmenting path or (in the last search) when the entire graph is a blossom. Shell_search also halts in these circumstances. (Note that if it finds an augmenting path $P$, then $P$ consists of eligible edges. Shell_search can halt because the current shell is a blossom only for an odd shell.) In addition, shell_search can halt in several other circumstances. We discuss these for the three phases in turn.

In phase 1, each execution of shell_search does at most one dual adjustment, and this adjustment uses $\delta=1$. (In fact, Lemma 5.1 below shows that each phase 1 execution of shell_search does precisely one dual adjustment.) After this adjustment, the search halts. This contrasts with Edmonds' algorithm,
where search chooses $\delta$ large enough so that $\mathscr{H}$ can change. An adjustment of $\delta=1$ may not allow any changes (or any augments in the following Augment Step).

In phase 1, after the dual adjustment there may be unweighted current root blossoms. Specifically, an inner root blossom can become unweighted in the dual adjustment. Shell_search repeatedly removes any unweighted root of the current blossom forest and replaces it by its children, until every nonleaf root is weighted. (This accomplishes what an expand step would have done in Edmonds' algorithm. It is done for the same reason-an unweighted blossom may be hiding an augmenting path.) Because of this rule, in any phase 1 Augment Step after the first, any current root blossom is weighted.

In phase 2 shell_search $(S)$ halts on entry if $S$ is an odd shell that does not contain all the free vertices of $R^{\prime}$. Also it halts when a dissolve step enlarges the shell so that other halting criteria apply. This can happen in three ways. First, dissolve may add to $S$ a new shell $S^{\prime}$ that has already been searched (in the current Search Step) and found to contain an augmenting path. Second, dissolve may enlarge $S$ to an odd shell not containing all free vertices of $R^{\prime}$. The third possibility is implied by the fact that only undissolved shells of $P(R)$ are searched. Consider the undissolved shell $G\left(R^{\prime}, D\right)$, where $D$ is the largest undissolved blossom in $R^{\prime}$. If shell search dissolves $R^{\prime}, G\left(R^{\prime}, D\right)$ is no longer contained in an undissolved shell of $P(R)$. So the search halts. (Note that any augmenting paths that have been created in shell_search will be processed with the major path containing the parent of $R$; this case never occurs when $R=G^{*}$.)

There are no extra halting rules for phase 3 . In phase 3 by definition no augmenting path can be found entirely in $R$. Thus shell_search halts when all blossoms on $P(R)$ are dissolved.

This concludes the statement of path. We now summarize some facts about shell_search that further motivate and justify the phase structure; the details of shell_search are in Section 7.

As already mentioned, the Search Step consumes more time in later phases. The usual implementation of search in Edmonds' algorithm (e.g., [13]) uses a priority queue to find the next dual adjustment quantity $\delta$. In phase 1 search can be implemented without a priority queue, since only one dual adjustment is made. The proper data structures make the time for one Search Step $O(\hat{m}(R))$. In phase 2 the priority queue can be implemented as an array. The phase 2 Search Steps collectively use total time $O(\hat{n}(R) \log \hat{n}(R))$ to scan the array; in addition, each Search Step uses $O(\hat{m}(R) \alpha(m, n))$ time. In phase 3 a standard priority queue is used; the time for phase 3 is (less than) $O(\hat{m}(R) \log \hat{n}(R)$ ).

Next we indicate why the criteria for eligibility change for phases $2-3$. Recall that in search, if a dual adjustment makes an inner blossom $B$ unweighted an expand step is done. As illustrated in Figure 5, expand replaces $B$ in $\mathscr{S}$ by an alternating path of edges of $E(B)$. Expand steps do not occur in phase 1, since a phase 1 search stops right after its dual adjustment. Expands do occur in phases 2-3. Now observe that an edge $e \in E(B)$ that expand places in $\mathscr{H}$ may not satisfy alternatives (i)-(iii) of eligibility. This occurs in the following scenario: A blossom step makes $e$ a blossom edge of $B$ with (ii) holding. Next an augmenting path passes through $B$, changing $e$ from matched to unmatched. A subsequent search makes $B$ an inner blossom. Then an expand step is done for $B$, adding $e$ to $\mathscr{I}$ (as in Figure 5). Now $e$ is an unmatched edges in $\mathscr{I}$,
not in a current blossom, with $y z(e)=c(e)-2$. Thus $e$ does not satisfy any of alternatives (i)-(iii).

To remedy this, phases $2-3$ use the weaker criteria for eligibility. This makes the above edge $e$ eligible. The weaker criteria suffice for these phases. (They would not be adequate for phase 1, however; see Lemma 5.1.)
This concludes the description of path. Showing that path is correct amounts to checking that it accomplishes the goal stated in the match routine: It dissolves all old blossoms on $P(R)$ while maintaining (1)-(2). The discussion above shows that (1)-(2) are preserved. (Note that any edge $e$ in $\mathscr{Y}$ has $y z(e) \geq c(e)-2$; hence, $e$ satisfies (1b) if it gets matched or enters a blossom.) The blossoms on $P(R)$ may all dissolve in phase 1 or 2 . Otherwise they dissolve in phase 3 because of the halting condition. When $R=G^{*}$, since $G$ is critical the entire graph eventually becomes a blossom. This dissolves the old blossom $G^{*}$, and path halts correctly.

This argument applies to the first scale, even though $T_{0}$ need not be a blossom tree. In the first scale $\operatorname{path}(R)$ is trivial except when $R=G^{*}$. In this case path alternates between Augment Steps and shell_search $\left(G^{*}\right)$, and works as desired. Note also that the first scale can detect noncritical input graphs if desired: $G$ is critical if and only if the first scale halts with a blossom $G^{*}$.

## 5. Efficiency: High-Level Analysis

This section analyzes the running time of path. It assumes an inequality that is derived in Section 6, thus completing the high-level analysis.
The efficiency of the algorithm depends on the fact that all quantities are integral. Let us verify this fact. First, observe that an execution of shell_search keeps all values integral if two conditions are met: (i) all costs are even; (ii) the search starts with all $y$ values of free vertices having the same parity. The proof that these two conditions suffice is essentially the same as the proof of the integrality property of Edmonds' algorithm, sketched in Section 1.1. It is modified to account for the fact that the search graph $\mathscr{F}$ consists of eligible edges rather than tight edges. This amounts to the observation that $c(e)-2$ is even if $c(e)$ is.
Thus to prove integrality it suffices to show that the scaling algorithm maintains conditions (i)-(ii). Condition (i) holds because the Double Step makes all costs even. For (ii), we show that throughout the scaling algorithm, any free vertex $v$ has $y(v)$ odd. A Double Step makes $y(v)$ odd. Suppose a call shell_search $(C, D)$ does a dual adjustment by $\delta$. If $v \in G(C, D)$, then $y(v)$ does not change-the dual adjustment increases it by $\delta$ (since $v$ is free) and the corresponding translation of $C$ decreases it by $\delta$. If $v \in D$, then the translations decrease $y(v)$ by $2 \delta$, so it remains odd. These are the only changes to $y(v)$.

In summary, we have shown that all quantities in the scaling algorithm are integral. For instance, each dual adjustment quantity $\delta$ is integral.
Now we present the properties that limit the number of iterations in the three phases. Clearly there is at most one phase 3 iteration. Phases 1 and 2 make progress by either adjusting the duals or augmenting the matching.

A phase 2 iteration need not adjust duals. This is because of the definition of eligibility. However, any phase 2 iteration after the first augments the
matching. To see this, observe that in phase 2 shell_search halts with an even shell only if an augmenting path has been found. The same is true if it halts with an odd shell that contains all free vertices of $R^{\prime}$, when there is more than one free vertex. So if a phase 2 Search Step halts without finding an augmenting path, either all blossoms are dissolved or there is precisely one shell, an odd shell with one free vertex. In the first case path terminates and in the second it starts phase 3.

A phase 1 iteration may not augment the matching, as remarked above. We now show that any phase 1 iteration does a dual adjustment. Note that in bipartite matching the analogous statement is true because there are no augmenting paths of eligible edges when a search starts [16]. This is not true for general graphs -an augmenting path of eligible edges may exist when shell_search is called. We show that the duals get adjusted nonetheless.

Lemma 5.1. In phase 1 after the first Search Step, any execution of shell_search $(S)$ adjusts duals by $\delta=1$.

Proof. By integrality it suffices to show that shell_search adjusts duals by some positive amount. Suppose for the sake of contradiction that shell_search $(S)$ halts before doing a dual adjustment, because it finds an augmenting path of eligible edges $P$. Let $r^{\prime \prime}$ denote the preceding Augment Step and let $H$ denote the auxiliary graph in $\mathscr{\sim}$.
First observe that $P$ corresponds to a path $P$ in $H$. This follows if we show that the edges of $P$ were eligible at the start of $\mathscr{\alpha}$ (for then each edge of $P$ is either in $H$ or in a blossom contracted in $H$ ). An Augment Step does not create an eligible edge. Also shell_search $(S)$ can do grow and blossom steps, but no expand steps (as observed above, at the start of shell_search $(S)$ every current root blossom is weighted). It follows that shell_search( $S$ ) does not create an eligible edge. This gives the desired conclusion.

Next observe that $\bar{P}$ is a simple path in $H$. This follows because if $B$ is a blossom contracted in $H$, shell_search $(S)$ did not expand $B$, so $E(P) \cap$ $E(G(B)$ ) is a subpath (possibly empty) of $P$.
We conclude that $\bar{P}$ is an augmenting path in $H$ at the end of $s /$. Now the maximality criterion in the Augment Step implies that in $H, \bar{P}$ contains a vertex $v$ of some path $Q \in \mathscr{\mathscr { F }}$. Thus $\bar{P}$ contains the matched edge $e$ incident to $v$. But $e$ became ineligible in the augment of $Q$. (This depends on the definition of eligibility in phase 1.) This is the desired contradiction.

We can now do the high-level timing analysis for $\operatorname{path}(R)$, where $R$ is any major path root. For convenience, write $\hat{n}=\hat{n}(R)$ and $\hat{m}=\hat{m}(R)$. Recall that $\rho$ is the number of iterations of the loop of path in phase 1 . Let $R_{1}$ be the largest undissolved blossom in $P(R)$ at the end of phase 1 . Let $F_{1}$ denote the set of free vertices of $R_{1}$. The number of phase 2 iterations is at most $1+\left(\left|F_{1}\right| / 2\right)$ since every phase 2 iteration after the first augments the matching. Assume for the moment that this product inequality holds:

$$
(\rho-\log \hat{n})\left(\left|F_{1}\right|-1\right) \leq 5 \hat{n} \log \hat{n}
$$

Thus, if $\rho \geq 2 \log \hat{n}$, then the number of phase 2 iterations is at most $(5 \hat{n} \log \hat{n}) / \rho+3 / 2$.

Recall the time bounds for the various phases, as already mentioned and presented in detail in Section 7: $O(\hat{m})$ for one iteration in phase 1 ,
$O(\hat{m} \alpha(m, n))$ for one iteration in phase 2 plus $O(\hat{n} \log \hat{n})$ total extra time, and $O(\hat{m} \log \hat{n})$ total time for phase 3. Then the total time for $\operatorname{path}(R)$ is $O(\rho \hat{m}+((\hat{n} \log \hat{n}) / \rho) \hat{m} \alpha(m, n)+\hat{m} \log \hat{n})$. If $\sqrt{\hat{n}} \geq 2 \log \hat{n}$, then take

$$
\rho=\sqrt{\hat{n} \alpha(m, n) \log \hat{n}} .
$$

This gives time $O(\sqrt{\hat{n} \alpha(m, n) \log \hat{n}} \hat{m})$ for $\operatorname{path}(R)$. If $\sqrt{\hat{n}}<2 \log \hat{n}$ then the time is $O(1)$ so the previous bound holds again. Now Lemma 3.1 gives the desired time bound for the entire algorithm.

To complete the timing analysis, we need only prove the above product inequality. We now show this inequality follows from the "witness inequality" defined below. To state the witness inequality, we first introduce two quantities that are fundamental in the next section. For a vertex $v$ and an old blossom $B$, at any time in path define

$$
\begin{aligned}
\hat{\delta}(B)= & \text { the total of all translations of } B \\
\hat{\delta}(v, B)= & \text { the total of all translations of } B \\
& \text { made when } v \text { is in an undissolved shell with outer boundary } B .
\end{aligned}
$$

Since translating $B$ by one decreases $z(B)$ by two, $\hat{\delta}(B) \leq z_{0}(B) / 2$; equality holds if $B$ dissolves before becoming a current blossom. The quantity $\hat{\delta}(v, B)$ counts all translations of $B$ "witnessed by" $v$. It is positive only if $v \in B$. There can be more than one inner boundary of shells contributing to $\hat{\delta}(v, B)$. We often write expressions like $\hat{\delta}(F, P)$, where $F$ is a set of vertices and $P$ a set of blossoms (for instance the blossoms in a path of the blossom tree); recall that by our conventions for functions this denotes $\sum\{\hat{\delta}(v, B) \mid v \in F, B \in P\}$.

Now choose any time in $\operatorname{path}(R)$. Let $\omega$ be a free vertex in the innermost possible blossom of $P(R)$. Let $F$ denote the set of free vertices of $P(R)$. The witness inequality is

$$
\hat{\delta}(F-\omega, P(R)) \leq 5 \hat{n} \log \hat{n} .
$$

(This inequality is one reason why Section 6 analyzes the algorithm in terms of $\hat{\delta}$ rather than the quantities $y, z$ directly involved in the algorithm. Intuitively, $\hat{\delta}(\nu, \cdot)$ is directly related to progress made by the algorithm, since the translations witnessed by a free vertex $v$ correspond to searches for an augmenting path involving $v$. On the other hand $y(v)$ does not change when progress is being made by searching for an augmenting path from $v$ - recall Section 4, modification (ii) of shell_search.)

To derive the product inequality, consider any vertex $v \in F_{1}$. In the Sort Step of any iteration of path, let $S(v)$ denote the undissolved shell of $P(R)$ that contains $v$. We show that the Search Step adjusts the duals in $S(v)$ in all but at most $\log \hat{n}$ phase 1 iterations of path. All but possibly one phase 1 execution of shell_search $(S(\nu))$ adjusts duals by $\delta=1$ (Lemma 5.1). The Search Step does not execute shell_search $(S(v))$ only if a boundary of $S(v)$ dissolves before $S(v)$ is examined. In this case the ordering of the Sort Step implies that the quantity $\hat{n}(S(v))$ doubles. This can happen only $\log (\hat{n} / 2)$ times, since a shell has at least two vertices. Thus, the Search Step adjusts duals in $S(v)$ in all but at most $1+\log (\hat{n} / 2)=\log \hat{n}$ iterations, as desired.

We conclude that in the $\rho$ iterations of phase $1, v$ witnesses at least $\rho-\log \hat{n}$ translations, that is, $\hat{\delta}(v, \mathrm{P}(\mathrm{R})) \geq \rho-\log \hat{n}$, and $\hat{\delta}(F-\omega, P(R))$
$\geq(\rho-\log \hat{n})\left(\left|F_{1}\right|-1\right)$. This together with the witness inequality obviously implies the product inequality.

## 6. The Witness Inequality

This section derives the witness inequality, thereby completing the high-level efficiency analysis for path. The section ends by proving a related inequality needed for the implementation of shell_search.

We start with terminology. The derivation centers around the old blossom tree $T_{0}$. We use an interval notation for paths in $T_{0}$ : If node $C$ is an ancestor of $D,[C, D]$ denotes the path from $C$ to $D$ with both endpoints included; $[C, D)$ is the same path with $D$ excluded, etc. For an odd shell $G(C, D)$ of $P(R)$, any interval ending with $D$, e.g., $[C, D)$, is interpreted as if $D$ were the last node of $P(R)$ (notice that $D=\varnothing$ for an odd shell). Recall that $G^{*}$ is the root of the old blossom tree $T_{0}$, so $\left[G^{*}, C\right]$ is the path from the root to $C$.

If $B$ is a node of a tree, $f^{\prime}(B)$ denotes the set of its nonmajor children and $\mathscr{D}(B)$ denotes the set of its descendants (including $B)$. These functions can also be applied to sets of nodes, e.g., if $P$ is a path in a tree, $\mathscr{F}^{\prime}(P)$ denotes the set of all descendants of nonmajor children of nodes of $P$. If $P$ is an interval, we omit the enclosing parentheses in these notations, so $H^{\prime}[C, D], F[C, D)$, etc. have the obvious meanings.

Edge $e$ crosses a set of vertices $B$ if precisely one of its ends is in $B$, that is, $|e \cap B|=1$. (In this notation $B$ is usually a blossom.) The crossing function $\gamma: 2^{V(G)} \rightarrow \mathbf{Z}$ of a matching $M$ is defined by $\gamma(B)=\mid\{e \mid e \in M$ crosses $B\} \mid$. For example if $B$ is a blossom of $M$ then $\gamma(B) \leq 1$. A blossom $B$ is uncrossed if $\gamma(B)=0$; otherwise it is crossed. A shell $G(C, D)$ (even or odd) is uncrossed if all its boundaries are uncrossed; otherwise it is crossed. During the algorithm an old blossom is uncrossed if it is undissolved; it is crossed if all its vertices have become matched; from when it dissolves until all its vertices become matched, it can alternate arbitrarily between crossed and uncrossed.

The first step in the derivation is to summarize the changes in duals $y, z$ caused by scaling and shell_searches. This leads to an inequality similar to the witness inequality, Lemma 6.1. To state it, fix a time in the execution of $\operatorname{path}(R)$. Let $M$ be the current matching. Let $\gamma$ be the crossing function for $M$. Choose a free vertex $\omega$ in the innermost possible blossom of $P(R)$. Let $M_{0}$ be the $\omega$-matching on $R$ given by the 2 -optimum matching of the previous scale. Let $\gamma_{0}$ be the crossing function for $M_{0}$. Thus an old blossom $B$ has $\gamma_{0}(B)=$ if $\omega \in B$ then 0 else 1. (In the first scale $R=G^{*}, \omega$ can be any vertex and $M_{0}$ any $\omega$-matching.)

Let $G(C, D)$ be an uncrossed shell of $P(R)$. (Bear in mind that $C$ or $D$ may be currently dissolved or undissolved, and $D$ can be $\varnothing$ for an odd shell.) Let $F_{\omega}$ be the set of free vertices of $G(C, D)-\omega$. (Possibly $F_{\omega}=\varnothing$.) Recall that the set of old blossoms is $V\left(T_{0}\right)$ and the set of current blossoms is $V(T)$. In the following lemma all time-dependent quantities $\left(\gamma, \hat{\delta}, F_{\omega}\right)$ are evaluated at the chosen time in the execution of $\operatorname{path}(R)$.

Lemma 6.1. At any time in path $(R)$ an uncrossed shell $G(C, D)$ of $P(R)$ satisfies

$$
\left(\gamma-\gamma_{0}\right) \hat{\delta}((C, D) \cup \mathscr{D} \mathscr{N}[C, D))+\hat{\delta}\left(F_{\omega}, V\left(T_{0}\right)\right) \leq 5 \hat{n}(C, D)
$$

Proof. We start with some terminology. We frequently use our convention of identifying a subgraph with its vertices or edges, e.g., $M \cap G(C, D)$ abbreviates $M \cap E(G(C, D))$. Define

$$
\begin{aligned}
M^{\prime} & =M \cap G(C, D), & M_{0}^{\prime} & =M_{0} \cap G(C, D), \\
d & =c\left(M_{0}^{\prime}\right)-c\left(M^{\prime}\right), & \mu(B) & =\left|M_{0}^{\prime} \cap G(B)\right|-\left|M^{\prime} \cap G(B)\right| .
\end{aligned}
$$

In the last definition $B$ is a blossom, old or current. Say that an old blossom and a shell "intersect" if they have a common vertex. Thus, the old blossoms that intersect $G(C, D)$ are those in $\left[G^{*}, D\right) \cup \mathscr{D} \mathscr{N}[C, D)$. The argument is based on estimating $d$ in two ways.

Observe that neither $M$ nor $M_{0}$ crosses $C$ or $D$. For $M$, this holds by hypothesis. Since $M$ does not cross $C$ and $\hat{n}(C)$ is odd, $\omega \in C$. (This depends on the fact that $G(C, D)$ is a shell of $P(R)$.) Hence, $M_{0}$ does not cross $C$. Similarly, $M_{0}$ does not cross $D$ if $D \neq \varnothing$.

First estimate $d$ using the initial duals $y_{0}, z_{0}$. Conditions (1) of the previous scale and the Double Step of the scaling routine imply

$$
\begin{array}{ll}
y_{0} z_{0}(e) \leq c(e), & \text { for } \quad e \in M \\
y_{0} z_{0}(e) \geq c(e)-8, & \text { for } \quad e \in M_{0}
\end{array}
$$

(This holds for the first scale $s=1$, since $y_{0} z_{0}(e)=-2$ and $|c(e)| \leq 2$.) Adding the $M$ inequalities and subtracting the $M_{0}$ inequalities for the edges of $M^{\prime} \cup M_{0}^{\prime}$ gives

$$
-y_{0}\left(F_{\omega}\right)+\mu z_{0}\left(V\left(T_{0}\right)\right) \leq-d+8\left|M_{0}^{\prime}\right| .
$$

This inequality depends on the fact that neither matching crosses $C$ or $D$. Recall that by the conventions of Section 1, $\mu z_{0}\left(V\left(T_{0}\right)\right)$ denotes $\sum\left\{\mu(B) z_{0}(B) \mid B \in V\left(T_{0}\right)\right\}$.

Next estimate $d$ in a similar way using the current duals $y, z$. Since (1) holds for the current duals, adding (1a) for $M_{0}^{\prime}$ and subtracting (1b) for $M^{\prime}$ gives

$$
y\left(F_{\omega}\right)-\mu z\left(V\left(T_{0}\right) \cup V(T)\right) \leq d+2\left|M^{\prime}\right|
$$

This also depends on the fact that neither matching crosses $C$ or $D$.
Next we bound the terms involving $V(T)$ and $V\left(T_{0}\right)$ in the two $d$ estimates. A current blossom $B \in V(T)$ has $\mu(B) \leq 0$. This follows since $\left|M^{\prime} \cap G(B)\right|$ is as large as possible; this in turn follows since $M$ induces a $u$-matching on $B$ for some $u \in B$, and no edge of $M$ crosses $C$ or $D$. (Note however that an edge of $B$ can cross $C$ or $D$.) Since $\mu$ is nonpositive on $V(T)$, the $V(T)$ term in the second $d$ estimate can be dropped.

We turn to the $V\left(T_{0}\right)$ terms. First note that an argument similar to the above shows that $\mu$ is nonnegative on old blossoms. This fact will be used below.

Clearly, $\mu$ vanishes on blossoms not intersecting $G(C, D)$, so we can restrict attention to old blossoms $B$ intersecting $G(C, D)$. Define

$$
f(B)=\left|F_{\omega} \cap V(B)\right| ;
$$

in addition. define $\gamma^{\prime}$ and $\gamma_{0}^{\prime}$ as the crossing functions of $M^{\prime}$ and $M_{0}^{\prime}$, respectively. We show the following inequality in order to bound the $V\left(T_{0}\right)$ terms:

$$
\mu\left(z_{0}-z\right)(B) \geq\left(f+\gamma^{\prime}-\gamma_{0}^{\prime}\right) \hat{\delta}(B) .
$$

To prove this, we first prove the equation $2 \mu(B)=\left(f+\gamma^{\prime}-\gamma_{0}^{\prime}\right)(B)$ : By definition $2 \mu(B)$ is how many more vertices of $V(B) \cap G(C, D)$ that $M_{0}$ matches on edges of $G(B) \cap G(C, D)$ compared to $M$. A vertex of $V(B) \cap$ $G(C, D)$ is not matched on an edge of $G(B) \cap G(C, D)$ if it is free or it is matched on an edge crossing $V(B) \cap G(C, D)$. There are $f(B)$ vertices of the first type (note that $\omega$ is free in both matchings). There are $\left(\gamma^{\prime}-\gamma_{0}^{\prime}\right)(B)$ vertices of the second type, since no edge of either matching crosses $C$ or $D$. This gives the desired equation.

Now we show that an old blossom $B$ has $\left(z_{0}-z\right)(B) \geq 2 \hat{\delta}(B)$ or $\mu(B)=0$. Note that this relation, together with the above equation and the nonnegativity of $\mu(B)$, implies the desired inequality for the $V\left(T_{0}\right)$ terms. The relation follows by considering three cases. If $B$ has never become a current blossom then $\left(z_{0}-z\right)(B)=2 \hat{\delta}(B)$. If $B$ has become a current blossom, but is not now a current blossom, then $\left(z_{0}-z\right)(B)=z_{0}(B) \geq 2 \hat{\delta}(B)$. If $B$ is a current blossom, then $\mu(B)=0$. The desired relation follows, as does the inequality for $V\left(T_{0}\right)$ terms.

Next we deduce

$$
\left(y-y_{0}\right)\left(F_{\omega}\right)+\left(f+\gamma^{\prime}-\gamma_{0}^{\prime}\right) \hat{\delta}\left(\left[G^{*}, D\right) \cup \operatorname{sit}[C, D)\right) \leq 5 \hat{n}(C, D)
$$

This follows by adding the two $d$ estimates and replacing the $\mu\left(z_{0}-z\right)$ terms using the above observations. In addition, note that $\left|M^{\prime}\right| \leq\left|M_{0}^{\prime}\right| \leq$ $\hat{n}(C, D) / 2$.

A free vertex $v$ has $y(v)=y_{0}(v)-(\hat{\delta}-\hat{\delta}(v, \cdot))\left[G^{*}, v\right)$. (To show this, observe that if $v$ is in a blossom that gets translated by $\delta$, then $y(v)$ decreases by $\delta$; if $\mathcal{v}$ is in a shell that gets duals adjusted by $\delta$, then $y(v)$ increases by $\delta$; these are the only times that $y(v)$ changes.) Summing these equations for all $v \in F_{\omega}$ implies

$$
\left(y-y_{0}\right)\left(F_{\omega}\right)+f \hat{\delta}\left(\left[G^{*}, D\right) \cup(C, D)\right)-\hat{\delta}\left(F_{\omega}, V\left(T_{0}\right)\right)=0
$$

Finally, subtract the last equation from the preceding inequality. This gives the lemma if we use two observations to simplify the expression ( $\gamma^{\prime}-$ $\left.\gamma_{0}^{\prime}\right) \hat{\delta}\left(\left[G^{*}, D\right) \cup \mathscr{D} \neq[C, D)\right)$ : The terms $\left(\gamma^{\prime}-\gamma_{0}^{\prime}\right) \hat{\delta}\left(\left[G^{*}, C\right]\right)$ vanish, since a blossom $B \in\left[G^{*}, C\right]$ has $\gamma^{\prime}(B)=\gamma_{0}^{\prime}(B)=0$. The remaining terms are for blossoms $B$ contained in $G(C, D)$. Hence the function $\gamma^{\prime}-\gamma_{0}^{\prime}$ simplifies to $\gamma-\gamma_{0}$.

Rewrite the lemma as

$$
\begin{align*}
& \left(\gamma-\gamma_{0}\right) \hat{\delta}((C, D))+\hat{\delta}\left(F_{\omega}, P(R)\right) \\
& \quad \leq 5 \hat{n}(C, D)+\left(\left(\gamma_{0}-\gamma\right) \hat{\delta}-\hat{\delta}\left(F_{\omega}, \cdot\right)\right) \underline{Z}, t[C, D) \tag{3}
\end{align*}
$$

Let us interpret this inequality and survey the rest of the derivation. Each scale after the first starts off with an "error" of $O(n)$, in the sense that the 2-optimum matching of the previous scale can cost $O(n)$ more than that of the current scale. If the graph is bipartite, this is the only source of error [16]. For general graphs there is a second type of error when $\operatorname{path}(R)$ begins. It comes from changes in the duals made by previous calls to path for descendants of $R$. Specifically, the error from scaling corresponds to the first term on the right-hand side of (3) and the error from descendants corresponds to the second
term. Loosely speaking the second term is bounded by the total amount of all dual adjustments in phases $1-2$. We will show this in a precise sense, in the process of proving that the second term is $O(\hat{n}(C, D) \log \hat{n}(R))$. Then (3) will imply that $\hat{\delta}\left(F_{\omega}, P(R)\right)$ is $O(\hat{n}(R) \log \hat{n}(R))$, the desired witness inequality.

We begin with a convenient notation that corresponds to "total dual adjustment." First let us clarify a phrase used throughout this section, a search of shell $G(A, B)$. This refers to the operation of shell_search when it is working on shell $G(A, B)$. Notice that it is possible for this search to occur during an execution of shell_search $(S)$ where $S \neq G(A, B)$ : If $S$ is a proper subshell of $G(A, B)$ then after one or more dissolve steps the search can be working on $G(A, B)$.

For an even shell $G(C, D)$ of $P(R)$ define

$$
\begin{aligned}
\Delta(C, D)= & \text { the total of all translations in searches of shells } \\
& G(A, B) \text { for } A, B \in[C, D] .
\end{aligned}
$$

Equivalently, $\Delta(C, D)$ is twice the total of all dual adjustments made in searches of shells $G(A, B)$, for $A, B \in[C, D] ; \Delta$ is evaluated after the last translation of a shell in $G(C, D)$. (Note that $\Delta(C, D)$ differs slightly from 'the total of all translations in executions shell_search $(A, B)$ for $A, B \in$ $[C, D]$.' If an execution shell_search $(C, D)$ dissolves $C$ or $D$ and proceeds to adjust duals by some positive amount, $\Delta(C, D)$ does not count the corresponding translations whereas the alternative definition would.)

The main step in the derivation discussed above is to show that any even shell $G(C, D)$ of $P(R)$ has

$$
\begin{equation*}
\Delta(C, D) \leq 5 \hat{n}(C, D)\lfloor\log \hat{n}(R)\rfloor . \tag{4}
\end{equation*}
$$

We prove (4) by induction on $\hat{n}(R)$. The base case, $R$ a leaf blossom, is vacuous. The inductive step is done in Lemmas 6.2-6.3 and the following paragraph. The lemmas inductively assume (4) for shells in major paths of descendants of $R$.

We start by bounding the error from descendants. (This essentially gives the witness inequality.) In the following lemma, all notation is as in Lemma 6.1.
Lemma 6.2. At any time in path $(R)$ an uncrossed shell $G(C, D)$ of $P(R)$ satisfies

$$
\left(\gamma-\gamma_{0}\right) \hat{\delta}((C, D))+\hat{\delta}\left(F_{\omega}, P(R)\right) \leq 5 \hat{n}(C, D)\lfloor\log \hat{n}(R)\rfloor .
$$

Proof. By (3) we need only show

$$
\begin{equation*}
\left(\left(\gamma_{0}-\gamma\right) \hat{\delta}-\hat{\delta}\left(F_{\omega}, \cdot\right)\right) \mathscr{F} \cdot[C, D) \leq 5 \hat{n}(C, D)(\lfloor\log \hat{n}(R)]-1) \tag{5}
\end{equation*}
$$

The left-hand side of (5) is made up of contributions coming from translations made prior to the call path $(R)$. More precisely consider a shell $G(A, B)$ of $P(S)$, where $S$ is a major path root descending from $\mathcal{N}[C, D)$. Let $f=\mid F_{\omega}$ $\cap G(A, B) \mid$. Suppose that a search of $G(A, B)$ does a dual adjustment of $\delta$. The corresponding translations contribute $\left(\left(\gamma_{0}-\gamma\right)(\{A, B\})-f\right) \delta$ to the left-hand side. The first step is to show that this contribution is positive only for uncrossed even shells containing no free vertices.

The number of vertices of $G(A, B)-\omega$ not matched on edges of $G(A, B)$ has the same parity in the old matching $M_{0}$ as in the current matching $M$.
(Recall that $\omega$ is free in both matchings; also note that $G(A, B)$ can be odd or even.) This number is $\gamma_{0}(\{A, B\})$ for $M_{0}$, and $\gamma(\{A, B\})+f$ for $M$. Thus $\left(\gamma_{0}-\gamma\right)(\{A, B\})-f$ is even. Since $\gamma_{0}(\{A, B\}) \leq 2$, the contribution is positive only if $\gamma_{0}(\{A, B\})=2$ and $\gamma(A)=\gamma(B)=f=0$, that is, $G(A, B)$ is an uncrossed even shell with no free vertices, as claimed.

Let $G(A, B)$ be an uncrossed even shell with no free vertices. Any blossom descending from $\psi[A, B)$ is crossed (since it has no free vertices). Thus, any uncrossed even shell with no free vertices that intersects $G(A, B)$ is a shell of $P(S)$. Among these shells is a maximal one. Let $G(A, B)$ denote such a maximal shell.

The maximal shells $G(A, B)$ are vertex disjoint. Ignoring the negative terms on the left-hand side of (5) shows that $G(A, B)$ contributes at most $\Delta(A, B)$. The inductive assumption (4) implies $\Delta(A, B) \leq 5 \hat{n}(A, B)|\log \hat{n}(S)|$, where $S$ is a descendant of a nonmajor child of $P(R)$, that is, $\hat{n}(S)<\hat{n}(R) / 2$. Thus $\Delta(A, B) \leq 5 \hat{n}(A, B)(|\log \hat{n}(R)|-1)$. Summing over all the disjoint shells $G(A, B)$ gives (5).

Now we estimate $\Delta$ for even shells. Call $G(C, D)$ active if during the course of the algorithm there is a search of shell $G(C, D)$ that adjusts duals by some positive amount. For an active shell $G(C, D)$ consider the last dual adjustment. Let $M$ be the matching at the time of this last dual adjustment and let $\gamma$ be the crossing function of $M$. Note there is a natural correspondence between intervals $[A, B)$ in the blossom tree and shells $G(A, B)$. Call an interval even, uncrossed, active, etc. if the corresponding shell is.

Lemma 6.3. An active interval $[C, D)$ of $P(R)$ can be partitioned into uncrossed intervals $\left[C, C^{\prime}\right),\left[C^{\prime}, D^{\prime}\right)$ and $\left[D^{\prime}, D\right)$ such that if $F$ is the set of free vertices of $M$ in $\left[C^{\prime}, D^{\prime}\right)$,

$$
\Delta(C, D) \leq \gamma \hat{\delta}\left(\left(C^{\prime}, D^{\prime}\right)\right)+\hat{\delta}(F,[C, D))+\Delta\left(C, C^{\prime}\right)+\Delta\left(D^{\prime}, D\right)
$$

Proof. Choose any free vertex $v$ in $G(C, D)$ and let $G\left(C^{\prime}, D^{\prime}\right)$ be the minimal uncrossed shell containing $\nu$. (Thus $C^{\prime}$ is the innermost blossom of $P(R)$ containing $v$ with $\gamma\left(C^{\prime}\right)=0$ and $D^{\prime}$ is the outermost blossom of $P(R)$ not containing $v$ with $\gamma\left(D^{\prime}\right)=0$.) Clearly $\left[C^{\prime}, D^{\prime}\right) \subseteq[C, D)$ since neither $C$ nor $D$ is crossed at the time of the last dual adjustment. Possibly $C=C^{\prime}$ or $D=D^{\prime}$.

Consider any shell $G(A, B)$ that gets searched, with $A, B \in[C, D]$. If $A, B \in\left[C, C^{\prime}\right]$ then the translations for this shell are counted by the term $\Delta\left(C, C^{\prime}\right)$ on the right-hand side; similarly if $A, B \in\left[D^{\prime}, D\right]$. The remaining possibility is that $G(A, B)$ intersects $G\left(C^{\prime}, D^{\prime}\right)$. Let $G\left(A^{\prime}, B^{\prime}\right)$ denote the intersection, that is, $A^{\prime}=$ if $A \in\left(C^{\prime}, D^{\prime}\right)$ then $A$ else $C^{\prime}, B^{\prime}=$ if $B \in$ $\left(C^{\prime}, D^{\prime}\right)$ then $B$ else $D^{\prime}$. Let $f$ denote the number of free vertices in $G\left(A^{\prime}, B^{\prime}\right)$. Since $G\left(A^{\prime}, B^{\prime}\right)$ is an even shell, $\gamma\left(\left\{A^{\prime}, B^{\prime}\right\}\right)+f$ is even. This quantity is positive, since if $\gamma\left(A^{\prime}\right)=\gamma\left(B^{\prime}\right)=0$ then the minimality of $G\left(C^{\prime}, D^{\prime}\right)$ implies that $A^{\prime}=C^{\prime}$ and $B^{\prime}=D^{\prime}$; thus $f>0$. We conclude that $\gamma\left(\left\{A^{\prime}, B^{\prime}\right\}\right)+f \geq 2$.

In a search of $G(A, B)$, a dual adjustment of $\delta$ contributes $2 \delta$ to the left-hand side of the inequality of the lemma. To show that it contributes at least this much to the right, observe that $\gamma\left(\left\{A^{\prime}, B^{\prime}\right\}\right) \delta$ is counted by the first term on the right, $f \delta$ by the second. Furthermore $\gamma\left(A^{\prime}\right)>0$ only if $A^{\prime}=A$, by the
above formula for $A^{\prime}$. Hence the term $\gamma \hat{\delta}\left(A^{\prime}\right)$ on the right only counts translations of $A$. The same holds for $B^{\prime}$, so our method of counting is valid.

Now we complete the proof of the inductive assertion (4). First observe that the active intervals are nested, that is, if two active intervals intersect then one contains the other. Thus any even interval $[A, B)$ of $P(R)$ has a unique set of maximal active subintervals [ $C, D$ ) such that $\Delta(A, B)$ equals the sum of all values $\Delta(C, D)$. We prove (4) for any even interval $[C, D)$ by induction on $\hat{n}(C, D)$. If $[C, D)$ is not active, apply the inductive assertion to each maximal active subinterval to get (4) for [ $C, D$ ). Suppose $[C, D$ ) is active. Apply Lemma 6.3. The first two terms on the right, $\gamma \hat{\delta}\left(\left(C^{\prime}, D^{\prime}\right)\right)+\hat{\delta}(F,[C, D))$, are bounded above by $5 \hat{n}\left(C^{\prime}, D^{\prime}\right)\lfloor\log \hat{n}(R)]$, by Lemma 6.2 applied to the uncrossed shell $\left[C^{\prime}, D^{\prime}\right.$ ). (Since $D$ is uncrossed, $\omega \in D$ and $\gamma_{0}$ vanishes on $(C, D)$.) Thus $\Delta(C, D) \leq 5 \hat{n}\left(C^{\prime}, D^{\prime}\right)\lfloor\log \hat{n}(R)]+\Delta\left(C, C^{\prime}\right)+\Delta\left(D^{\prime}, D\right)$. The desired conclusion follows from the inductive hypothesis for $\left[C, C^{\prime}\right.$ ) and [ $D^{\prime}, D$ ). Thus we have proved (4), and in the process shown that Lemmas 6.2-6.3 hold with no assumptions.

Now we derive the witness inequality. Let $R$ be a major path root. At any time in the execution of $\operatorname{path}(R)$ let $F$ be the set of free vertices of $R$, and let $\omega$ be a free vertex in the innermost possible blossom of $R$. The choice of time in path $(R)$ implies that $\gamma(R)=0$ and $\omega$ exists. Note that a blossom $B \in P(R)$ has $\gamma(B) \geq \gamma_{0}(B)$ (since $\gamma_{0}(B)=$ if $\omega \in B$ then 0 else 1 , and $\omega \notin B$ implies $\gamma(B)>0$ ). Thus Lemma 6.2 applied to the uncrossed shell $R$ gives the witness inequality,

$$
\hat{\delta}(F-\omega, P(R)) \leq 5 \hat{n}(R) \log \hat{n}(R) .
$$

We conclude with a related inequality used to implement the priority queue for shell_search (in Section 7).
Corollary 6.1. For a major path root $R$, the total dual adjustment in all phase 2 shell_searches of path $(R)$ is at most $5 \hat{n}(R) \log \hat{n}(R)$.
Proof. Write the total dual adjustment in phase 2 as $d_{1}+d_{2}$, where $d_{1}$ is the total adjustment when there are free vertices in $R^{\prime}$ that are not in the (undissolved) odd shell, and $d_{2}$ is the remainder, that is, the total adjustment when all free vertices in $R^{\prime}$ are in the odd shell. We show that each $d_{i} \leq(5 / 2) \hat{n}(R) \log \hat{n}(R)$.
The dual adjustments counted in $d_{1}$ all occur in a search of an even shell. If $D$ is the smallest blossom in $P(R)$, (4) shows that $\Delta(R, D) \leq$ $5 \hat{n}(R, D)\lfloor\log \hat{n}(R)\rfloor$. Thus the definition of $\Delta$ implies that $d_{1} \leq$ $(5 / 2) \hat{n}(R) \log \hat{n}(R)$.
For $d_{2}$, consider the odd shell $G\left(R^{\prime}, \varnothing\right)$ immediately after the last dual adjustment counted in $d_{2}$. The definition of phase 2 implies that at this time the odd shell has at least three free vertices $\nu$. Each such vertex has witnessed every dual adjustment of an odd shell in phase 2 , that is, $\hat{\delta}(\nu, P(R)) \geq d_{2}$. Thus the witness inequality implies that $d_{2} \leq(5 / 2) \hat{n}(R) \log \hat{n}(R)$.

## 7. The Search Step

This section gives the data structures and details of the Sort and Search Steps.
We begin with a data structure needed for the Search Step. At the start of the match routine, the old blossom tree $T_{0}$ is ordered so that every major child is a
rightmost child. The vertices of $G$, which are the leaves of $T_{0}$, are numbered from 1 to $n$ in left-to-right order. In the following discussion we identify each vertex with its number. Each node $B$ of $T_{0}$ stores $l o(B)$, its lowest-numbered leaf descendant. The given graph $G$ is represented by adjacency lists, two lists for each vertex $v$. One list for $v$ contains the edges $\{v w \mid w<v\}$ ordered by decreasing $w$. The other list contains the edges $\{v w \mid w>v\}$ ordered by increasing $w$.

This data structure is constructed (once in each scale) in time $O(m)$ using a bucket sort. The main property of the vertex order is that in any execution of $\operatorname{path}(R)$, the vertices of an undissolved shell $G(C, D)$ (even or odd) constitute the interval $[l o(C) \ldots l o(D))$. Hence for any vertex $v$ in an undissolved shell $G(C, D)$, the edges incident to $v$ in $G(C, D)$ can be found by scanning the appropriate part of $v$ 's two adjacency lists (assuming the values $l o(C), l o(D)$ are known). The time for this scan is $O(1)$ plus time proportional to the number of edges found in $G(C, D)$.

Now consider an execution of path( $R$ ). As in Section 5 let $\hat{n}=\hat{n}(R)$ and $\hat{m}=\hat{m}(R)$. The undissolved blossoms of $P(R)$ are stored in a doubly-linked list $\|$; the order of blossoms in $\mathscr{U}$ is the same as in $P(R)$.

The Sort Step can be done in $O(\hat{n})$ time using a bucket sort.
We turn to the Search Step. During the course of a Search Step a vertex goes through three states: it starts out asleep; while it is in the shell being searched it is active; when its shell has been searched it is dead. (Vertices outside of the largest undissolved shell $R^{\prime}$ are always dead.) In greater detail, when a boundary of a shell $G(C, D)$ dissolves, the new larger shell is composed of the vertices of $G(C, D)$, currently active, and new vertices, which were either asleep or dead. If they were asleep all vertices of the new shell are active; if they were dead all vertices of the new shell are dead. (If $C$ and $D$ dissolve simultaneously, with one boundary introducing asleep vertices and the other dead vertices, the new shell is dead.) Thus a vertex is active in the search for only one augmenting path-once dead, it remains dead for the rest of the Search Step. All this follows from the statement of the Search Step and the halting criterion.

Consider the time in the Search Step for dissolving shell boundaries. (This is important in phases 2 and 3.) Suppose shell $G(C, D)$ is being searched and $C$ dissolves ( $D$ dissolving is similar). Let $B$ be the blossom preceding $C$ in list ${ }^{\prime} \ell$, that is, the smallest undissolved blossom containing $C$. $C$ is deleted from ' 4 . Suppose the new vertices (those in $B-C$ ) become active. The edges in the new shell $G(B, D)$ are found by scanning the adjacency lists of the new vertices (the interval for the new shell is $[\operatorname{lo}(B) \ldots l o(D))$ ). Since a vertex becomes active only once, the total time for scanning edges in dissolve steps in one Search Step is $O(\hat{m})$. (Some additional processing that is done when a blossom dissolves is discussed below: it concerns dual values.)

Next consider the time spent in the Search Step manipulating the priority queue and doing related processing to find the next dual adjustment quantity $\delta$ (as described in Section 1.1). We consider phases 1,3 , and 2 in that order.

In phase 1 the priority queue is not needed, since only one dual adjustment with $\delta=1$ is made. Blossom steps are implemented in linear time using the incremental tree set-merging algorithm of [15]. This makes the time for one Search Step in phase $1 O(\hat{m})$.

The total time for phase 3 is $O(\hat{m} \log \hat{n})$. This can be achieved by imple-
menting the priority queue as a balanced tree [17]. Gabow, et al. [13] give an even better bound but this is not needed here.

For phase 2, Corollary 6.1 shows that the total dual adjustment is at most $5 \hat{n} \log \hat{n}$. The priority queue is implemented using two arrays, $P[0 \ldots \hat{n}-1]$ and $Q[0 \ldots 5 \log \hat{n}]$. Each array entry points to a list of priority queue entries. At any time, an edge $e$ that will become eligible after the shell_search has done a total dual adjustment of $d$ units is placed in the list for $Q(\lfloor d / \hat{n} \mid)$. When the total dual adjustment is in the range $[i \hat{n} \ldots(i+1) \hat{n})$, the edges in list $Q(i)$ are stored in the lists for $P$, the above edge $e$ being placed in the list for $P(d \bmod \hat{n})$. The edges in the list for $P(j)$ are eligible when the total dual adjustment reaches $i \hat{n}+j$. After a shell is searched, any entries remaining in this priority queue are removed, in linear time. This makes the total overhead for the queue in all phase 2 searches $O(\hat{n} \log \hat{n})$. (See [16] for a more detailed discussion of the implementation of such a queue.)

To implement expand steps the list-splitting algorithm of [10] is used. This makes the time for one Search Step in phase $2 O(\hat{m} \alpha(m, n))$ (since a vertex is active in only one search).
The last aspect of the Search Step discussed here is maintaining the duals $y, z$. Most details are the same as in an efficient implementation of Edmonds' algorithm (see [13] and [17]; although the main concern of these papers is implementing the priority queue discussed above, the details needed here are also given). The main technique is using offsets to facilitate the adjustment of dual values. In addition we use offsets in connection with old blossoms and their translations. We show how the algorithm translates a blossom in $O(1)$ time, and also how it calculates $y z(e)$ in $O(1)$ time. This plus the details in [13] and [17] give the desired time bound for our algorithm.
We start by describing the data structure. The algorithm stores two values for each old blossom $B, z_{1}(B)$ and $t(B)$, plus a value $y^{\prime}(v)$ for each vertex $v$. Intuitively $z_{1}$ keeps track of old $z$ values, $y^{\prime}$ keeps track of $y$ values and $t$ keeps track of translations. More precisely the algorithm maintains these invariants:

$$
\begin{aligned}
y(v) & =y^{\prime}(v)-t\left(\left[G^{*}, v\right]\right) ; & & \\
z\left(\left[G^{*}, B\right]\right) & =z_{1}(B)-2 t\left(\left[G^{*}, B\right]\right), & & \text { for } B \text { undissolved; } \\
t(B) & =0, & & \text { for } B \text { dissolved. }
\end{aligned}
$$

Initially each old blossom $B$ has $z_{1}(B)=z_{0}\left(\left[G^{*}, B\right]\right)$ and $t(B)=0$, and each vertex $v$ has $y^{\prime}(v)=y(v)$, so the invariant holds.

Consider a search of shell $G(C, D)$. To calculate $y z(e)$, write $y z(e)=y(e)$ $-z\left(\left[G^{*}, C\right]\right)-z(\{B \mid B$ is a current blossom containing $e\})$. The last term is calculated as in Edmonds' algorithm, so we concentrate on $y(e)-z\left(\left[G^{*}, C\right]\right)$. This equals $y^{\prime}(e)-z_{1}(C)$, since $C$ is the smallest undissolved blossom containing either end of $e$. Hence $y z(e)$ can be calculated in $O(1)$ time as claimed.
Next suppose the shell_search does a dual adjustment of $\delta$. This necessitates translating blossoms $C$ and $D$ by $\delta$. The algorithm does this by increasing $t(C)$ and $t(D)$ by $\delta$. This has the same effect as a translation, maintaining the above two invariants. Hence the dual adjustment is done correctly in $O(1)$ time.

It remains only to discuss how dual values are processed when a boundary $C$ or $D$ dissolves in the shell_search. In general consider consecutive undissolved blossoms $Z \subseteq Y \subseteq X$ of $P(R)$ and suppose blossom $Y$ has dissolved.

The invariants can be preserved by working on $G(X, Y)$ using time $O(\hat{n}(X, Y))$ or working on $G(Y, Z)$ using time $O(\hat{n}(Y, Z))$. For the first, assign $y^{\prime}(v) \leftarrow y^{\prime}(v)+t(Y)$ for each vertex $v \in X-Y$, then $t(X) \leftarrow t(X)$ $+t(Y), \quad z_{1}(X) \leftarrow z_{1}(X)+2 t(Y)$ and finally $t(Y) \leftarrow 0$. For the second, assign $y^{\prime}(v) \leftarrow y^{\prime}(v)-t(Y)$ for each vertex $v \in Y-Z$, then $t(Z) \leftarrow t(Z)+$ $t(Y)$ and $t(Y) \leftarrow 0$.

The algorithm adjusts duals by working on the vertices whose state changes. This implies that the time for all dissolves in a Search Step is $O(\hat{n})$.

Note that this discussion applies to both even shells and the odd shell. Observe how an odd shell $G(C, \varnothing)$ is processed: $C$ dissolves by becoming either an unweighted or current blossom. If $C=R^{\prime}$ the new odd shell is dead. In this case $t\left(R^{\prime}\right)$ becomes zero and, after the invariant has been restored, each vertex $\mathcal{v} \in R$ has $y^{\prime}(v)=y(v)$. In particular each scale halts with $y^{\prime}=y$.

## 8. The Augment Step

This section shows that the Augment Step can be done in linear time. This amounts to solving the following problem in linear time: Given an arbitrary graph with a matching $M$, find a maximal set $\mathscr{\mathscr { P }}$ of vertex-disjoint augmenting paths. We present an algorithm based on depth-first search and the properties of blossoms [6]. This section uses the following notation: $F$ denotes the set of free vertices of $M$. For a vertex $v \notin F, v^{\prime}$ denotes the vertex matched to $v$.

The algorithm grows a search graph similar to the search graph $\mathscr{F}$ in Edmonds' weighted matching algorithm (Section 1.1) with three changes: First, the requirement that an edge of $\mathscr{S}$ be tight is dropped (tightness is irrelevant since there are no edge costs). Second, an inner blossom is always a vertex, never a nonleaf blossom. (This comes about because the search starts with a graph that has no blossoms. As a consequence the algorithm has no expand steps-only grow and blossom steps.) Third, the free vertices are added to the search graph one at a time rather than simultaneously. A free vertex $f$ is either outer (if some search starts from $f$ ) or inner (if a search ends by finding an augmenting path to $f$ ). The contracted subgraph $\mathscr{F}$ (of Section 1.1) is always a forest: an augmenting path corresponds to a path in $\mathscr{S}$ joining a free outer vertex to a free inner vertex.

The final and most important difference from Edmonds' weighted matching algorithm is that the search is done depth-first. Figure 7 gives the recursive depth-first search procedure find_ap. Figure 8 shows a search graph constructed by find_ap. Here the input to find_ap is vertex 1, and vertices are labelled in the order they become outer. We explain the search algorithm by first presenting a simplified version and then discussing find_ap.

It is convenient to use the terminology of ordered trees. Thus the children of a vertex are ordered from left to right. Given two vertices in an ordered tree, one vertex is either an ancestor of or to the left of the other.

The goal of the simplified depth-first search is to find one augmenting path. The search maintains a search tree $\mathscr{T}$ (an ordered tree) in a graph $\bar{G}$ that is a contraction of the given graph $G$. It will be seen that $\bar{G}$ is the result of contracting blossoms in $G$. In discussing the simplified depth-first search we write $G$-vertex or $\bar{G}$-vertex to specify to which graph a vertex belongs. A path from a $\bar{G}$-vertex $v$ in $\mathscr{T}$ to the root is alternating; $v$ is outer (inner) if the length of the path is even (odd). The search maintains an active vertex, an outer vertex on the rightmost path of $\mathscr{T}$, from which edges are scanned.

```
procedure find_ap(x) \(\{x\) is an outer vertex \(\}\)
for each edge \(x y \notin M\) do \{examine an edge\}
    if \(y \notin V\left(y^{\prime}\right)\) then
        if \(y\) is free then begin \{an augmenting path has been found \}
            add \(x y\) to \(I\), and add path \(y P(x)\) to,\(\vec{y}\)
            terminate all active recursive calls to find__ap
            end
    else begin \{two grow steps\}
            add \(x y, y y^{\prime}\) to \(\mathscr{I}^{\prime}\), by setting \(l\left(y^{*}\right) \leftarrow x\)
            find_ap( \(y^{\prime}\) )
            end
    else if \(b(y)\) became outer strictly after \(b(x)\) then begin \{blossom step \}
    let \(u_{i}, i=1, \ldots, k\) be the inner vertices in \(P(y, b(x))\), ordered so that \(u_{i}\) precedes \(u_{t-1}\)
    for \(i \leftarrow 1\) to \(k\) do begin \{update \(f\) \}
            \(l\left(u_{t}\right) \leftarrow(y, x)\)
            for each vertex \(v\) with \(b(v) \in\left\{u_{i}, u_{i}^{\prime}\right\}\) do \(b(v) \leftarrow b(x)\)
            end
        for \(i \leftarrow 1\) to \(k\) do find_ap \(\left(u_{t}\right)\)
    end
```

Fig. 7. Depth-first search procedure find_ap


Fig. 8. Example graph for find_ap.

Initially $\bar{G}=G$ and $\bar{T}$ consists of a root which is a free vertex $f$ of $G ; f$ is also the active vertex.

The search works by repeatedly scanning an edge incident to the active vertex. The active vertex changes as the search progresses. At any point if all edges incident to the current active vertex $x$ have been scanned from $x$ then the search backs up-the grandparent of $x$ becomes the new active vertex (recall the parent of $x$ is inner, the grandparent is outer). If the search attempts to back up from the root of $\mathscr{T}$ then it halts-it will be seen that in this case there is no augmenting path from the initially chosen free vertex $f$.

The procedure to scan an edge incident to the active vertex $x$ is as follows. An edge $x y$ that has not been scanned from $x$ is chosen arbitrarily. First suppose $y \notin \mathscr{T}$. (It will be seen that only vertices in $\mathscr{T}$ get contracted; hence $y$ is a $G$-vertex.) If $y$ is free then the search halts (an augmenting path from $f$ to
$y$ has been found). Otherwise two grow steps are done: Recall that $y^{\prime}$ denotes the vertex matched to $y . \mathscr{T}$ is extended by adding edges $x y$ and $y y^{\prime}$, and the active vertex becomes $y^{\prime}(y$ is made the rightmost child of $x$, so the new active vertex is on the rightmost path of $\mathscr{T})$. The search continues by scanning an edge incident to the new active vertex.

Now suppose $y \in \mathscr{J}$. If $y$ is an outer descendant of $x$ then a blossom step is done: The path from $y$ to $x$, plus edge $x y$, forms a blossom. It is shrunk to a new vertex $B$ giving a new contracted graph $\bar{G}$. The active vertex becomes $B$ (clearly $B$ is on the rightmost path). The search continues by scanning an edge incident to the new active vertex. (Note that when $y \in \mathscr{T}$ nothing is done if $y$ is not an outer descendant of $x$, even though in some cases a new blossom could be shrunk.)

This completes the description of the simplified algorithm. It finds an augmenting path leading to the free vertex $f$ if one exists. We shall not prove this completely (it follows from arguments similar to Lemma 8.1 below). Instead we prove two properties, designated (i) and (ii), needed to analyze find_ap. Property (i) is that if $v$ and $w$ are adjacent outer vertices and edge $v w$ has been scanned from both vertices then $v$ and $w$ are in the same blossom. (Although it is not logically necessary for our development, we remark that property (i) essentially proves that the simplified algorithm works correctly when it halts without finding an augmenting path: The property implies no more blossom steps are possible, which in turn implies that no augmenting path exists.)

We first show that the algorithm maintains the invariant that no edge goes from an outer vertex to the right. It is easy to see that at any time any outer vertex not on the rightmost path has been completely scanned. Thus, in a grow step, no outer vertex to the left of the rightmost path is adjacent to $y$ or $y^{\prime}$. Thus grow steps preserve the invariant. A blossom step preserves the invariant, since in general contracting an edge in a tree does not move a vertex to the left or right of another. We conclude that the invariant always holds.

We make a second observation. Consider edges $x y, y y^{\prime}$ added to $\mathscr{T}$ in two grow steps that make $y$ inner and $y^{\prime}$ outer. Suppose that after $y^{\prime}$ is made active the search eventually backs up from $y^{\prime}$. Then if $y^{\prime}$ is ever contained in the active vertex again, edges $x y$ and $y y^{\prime}$ are contracted. To prove this consider the first time $y^{\prime}$ or one of its descendants again enters the active vertex. This occurs in a blossom step that contracts a path from the current active vertex, a nondescendant of $y^{\prime}$, to a descendant of $y^{\prime}$. Clearly edges $x y$ and $y y^{\prime}$ are on this path and get contracted.

Now we prove property (i). Suppose that at any time an edge $v w$ joins two distinct outer vertices of $\bar{G}$. The invariant implies that one vertex, say $v$, is an ancestor of the other, $w$. We show that if edge $v w$ gets scanned from vertex $v$ (or from some $\bar{G}$-vertex containing $v$ ) then $v$ and $w$ are vertices in the same blossom. (Clearly this implies (i).) Consider the time when $v w$ is scanned from $v$, and assume $w$ is not in the same blossom as $v$. Vertex $w$ is in $\mathscr{T}$ when $v w$ is scanned (otherwise a grow step would make $w$ a child of $v$; but then $v w$ never joins two distinct outer blossoms). Hence $w$ is inner or outer at this time.

Suppose $w$ is inner. Let $a$ be the first ancestor of $w$ on the path from the active vertex $v$ to the root. Observe that the search has backed up from every outer vertex on the path $P$ from $w^{\prime}$ to $a$, except $a$. By the time $w$ becomes outer, all vertices on $P$ are in the same blossom as $a$ (this follows from the second observation above). But since $a$ is an ancestor of $v, v$ is never a proper
ancestor of the outer vertex $w$. This contradiction shows that $w$ cannot be inner.
We conclude that $w$ is outer when $v w$ is scanned from $v$. Furthermore $w$ is a descendant of $v$ (since edge contractions cannot change a proper ancestor into a proper descendant). Hence a blossom step is done, placing $v$ and $w$ in the same blossom, as claimed. This proves (i).
Now we extend the simplified algorithm to record information about $G$. Call a $G$-vertex outer if it is contained in an outer $\bar{G}$-vertex. The search implicitly constructs an even-length alternating path from each outer $G$-vertex $x$ to the free vertex $f$; let $P(x)$ denote this path. We extend the algorithm to make the paths $P(x)$ explicit, using the following data structure ([8]).
In this discussion interpret a path $P$ as an ordered list of vertices. For example the first vertex of $P(x)$ is $x$. Let $P^{r}$ denote the reverse path of $P$; if $Q$ is also a path let $P Q$ denote the concatenation of the two paths. (For this to be a path the last vertex of $P$ must be adjacent to the first vertex of $Q$.) For vertex $y \in P(x)$, let $P(x, y)$ denote the subpath of $P(x)$ from $x$ to $y$.
Each outer vertex $x$ has a label $l(x)$ that defines path $P(x)$ as follows. A label is either a singleton label $l(x)=y$, where $y$ is an outer vertex, or a pair label $l(x)=(y, z)$, where $y$ and $z$ are outer vertices and the pair is ordered. If $x$ has a singleton label $l(x)=y$ then $P(x)=x x^{\prime} P(y)$. (As a degenerate case a free vertex $x$ has singleton label $l(x)=\varnothing$ and we define $P(x)=x$.) If $x$ has a pair label $l(x)=(y, z)$ then necessarily $x \in P(y)$, and $P(x)=P(y, x)^{r} P(z)$.

The algorithm assigns labels as follows. In the grow steps, $y^{\prime}$ gets the singleton label $x$. In the blossom step, each inner vertex $u$ in the path from $y$ to $x$ gets the pair label $(y, x)$.

Property (ii) is that the labels define $P(x)$ as an even-length alternating path from $x$ to $f$. Recall that by definition an alternating path is simple. This is important for correctness of the algorithm, since it is well-known that augmenting a matching along a nonsimple augmenting path can produce a set that is not a matching. Property (ii) is easy to prove by induction. The inductive assertion also includes the fact that if $x$ is in blossom $B$ of $\bar{G}$, then $P(x) \cap, y^{y}$ is precisely the tree path from $B$ to the root of $\mathscr{T}$.

This completes the discussion of the simplified depth-first search, which finds one augmenting path. The final version of the algorithm finds a maximal set of augmenting paths. To do this we must be more precise about selecting the active vertex. In the final algorithm the active vertex is a $G$-vertex. The final algorithm maintains the following property: For the active vertex $x, P(x)$ contains any outer $G$-vertex that has not been completely scanned. (This is analogous to ordinary depth-first search of a directed or undirected graph, where the search path leading to the vertex currently being scanned contains all vertices that are not completely scanned [1].) In order to maintain this property the blossom step works as follows: First it makes the inner vertices on the path $P$ from $y$ to $x$ outer; then it makes these vertices active in the reverse order of their occurrence in $P$.

The final algorithm uses one other data structure, to represent the blossom structure: For each vertex $x, b(x)$ denotes the base of the root blossom containing $x$.

Incorporating these changes into the simplified algorithm gives the algorithm find_ap stated in Figure 7. Observe that for any free vertex $f$, find_ap $(f)$ is a correct implementation of the simplified algorithm searching from $f$. Here we
assume that find_ap( $f$ ) begins with $\mathscr{\mathscr { P }}$ empty, each $b(v)$ initialized to $v$ and the search graph $\mathscr{F}$ initialized to root $f$. (Note that $\mathscr{F}$ in find_ap corresponds to $T$ in the simplified algorithm; in its final usage below, $y$ will contain more than one connected component, and will be analogous to the search graph $\mathscr{F}$ of Section 1.1.) Correctness depends on two observations. First, the recursion correctly implements the notion of the active vertex in $\bar{G}$. (The recursive calls in a blossom step are consistent with the simplified algorithm.) Second, the test preceding a blossom step, that $b(y)$ became outer strictly after $b(x)$, is equivalent to the simplified algorithm's test that $y$ is in an outer descendant of $x$ (since if $b(y)$ is outer then so is $y$, and edge $x y$ implies $x$ and $y$ are related in the search tree of $\bar{G})$.

Since find_ap implements the simplified algorithm it satisfies properties (i)-(ii). In particular property (i) states that if $v$ and $w$ are adjacent outer vertices and edge $v w$ has been scanned from both vertices then $b(v)=b(w)$.

We further observe the following property (iii): When find_ap halts, every outer vertex not in $\mathscr{P}$ has been completely scanned. To prove this it suffices to show that every time find_ap(x) examines an edge every unscanned outer vertex is in $P(x) \cup \mathscr{P}$. (Say that find_ap "examines an edge" each time control passes to the line so labelled in Figure 7. This includes the last time, when no unscanned edges $x y$ exist.) This invariant follows easily because of the order of activating vertices $u_{i}$ in a blossom step.

The final depth-first algorithm uses a main routine find_ap_set. This routine initializes the search graph $\mathscr{S}$ to empty and each $b(v)$ to $v$. Then it examines each vertex $f \in F$ in turn. If $f$ is not in a path of $\mathscr{P}$ when it is examined, the routine adds $f$ to $S$ (by assigning $l(f) \leftarrow \varnothing$ ) and calls the recursive procedure find_ap $(f)$. Note that when find_ap discovers an augmenting path, it immediately terminates itself and all currently active recursive calls. Also note that vertices are never removed from $\mathscr{S}$-when find_ap_set calls find_ap, $\mathscr{Y}$ is the same as when the last call terminated. Thus it is possible that a scanned edge $x y$ has $y$ in a previous search tree. Examining find_ap shows that in this case $y$ is ignored. It is easy to see that properties (i)-(iii) hold for find_ap_set as well as find_ap.

Now we can show that find_ap_set performs as desired.
Lemma 8.1. When find_ap_set halts, $\mathscr{P}$ is a maximal set of vertex-disjoint augmenting paths.

Proof. By property (ii), ip consists of vertex-disjoint augmenting paths. We need only show that when find_ap_set halts, $\mathscr{y}$ is maximal, that is, any augmenting path contains a vertex of $\mathscr{P}$.

When find_ap_set halts, consider an alternating path with vertices $x_{1}$, $i=0, \ldots, k$ that starts at a free vertex $x_{0}$ and is vertex-disjoint from $5 \mathscr{y}$. We show by induction that every $x_{2}$, is outer and $b\left(x_{2 \jmath}\right)=x_{h}$ for some $h \leq 2 j$. (The argument will also show that the path is not augmenting.)

The base case $j=0$ holds because find_ap $\left(x_{0}\right)$ was called. For the inductive step assume that $x_{2,}$ is outer. Since $x_{2,} \notin \mathscr{P}$ it has been completely scanned, by property (iii). Thus since $x_{2 j+1} \notin \mathscr{\mathscr { F }}, x_{2 j+1}$ is not free, that is, $x_{2 j+1}$ is matched and either inner or outer. If $x_{2 j+1}$ is inner then $x_{2 j+2}$ is outer; obviously $b\left(x_{2 J+2}\right)=x_{2 J+2}$ so the inductive assertion holds. (Note that the vertices $x_{2 \jmath}$ and $x_{2 \jmath+1}$ need not be in the same search tree of find_ap.) If $x_{2 j+1}$ is outer then it has been completely scanned. Thus $b\left(x_{2 \jmath+1}\right)=b\left(x_{2 \jmath}\right)$,
from property (i). Since $b\left(x_{2 \jmath}\right) \neq x_{2 \jmath+1}$ by the inductive hypothesis, $b\left(x_{2 j+1}\right)$ $\neq x_{2 j+1}$. Hence $x_{2 j+2}$ is outer and $b\left(x_{2 j+2}\right)=b\left(x_{2 j}\right)$. This completes the induction.
The inductive argument has also shown that no vertex $x_{2 j+1}$ is free. This shows there is no augmenting path disjoint from $\mathscr{P}$.

The time for find_ap_set is $O(m)$. To see this, first note that the values of $b$ can be updated and accessed in total time $O(m)$ using the incremental tree set-merging algorithm of [15]. Next note that in a blossom step the vertices $u_{i}$ are found using the following observation: The vertices $u_{t}^{\prime}$ are the predecessors of $b(x)$ in the sequence $(b l)^{j} b(y), j=0, \ldots$ (this follows since if $v$ is a blossom base then in the simplified algorithm its grandparent is $b(l(v)))$. This implies that in all blossom steps the total time to find all vertices $u_{t}$ is $O(n)$. It is obvious that the rest of the time for find ${ }_{-} a p_{-} s e t$ is $O(m)$.
After find_ap_set the Augment Step augments along each path of $\mathscr{P}$. This takes total time $O(n)$. To see this, note that it is easy to give a recursive routine that finds the edges in a path $P(x)$ in time proportional to their number; after finding an augmenting path it can be augmented. Alternatively, Gabow [8] gives a one-pass procedure.

## 9. Analysis Completed: Size of Numbers

This section completes the efficiency analysis. We have implicitly assumed that all arithmetic operations use $O(1)$ time. To justify this assumption, we show that all numerical values calculated by the algorithm have magnitude $O\left(n^{2} N \log (n N)\right)$. Since the input values require a word size of at least $\max \{\log N, \log n\}$ bits this implies that at worst quadruple-word integers are needed. Thus an aritmetic operation uses $O(1)$ time.

Lemma 9.1. At any time in the scaling routine a $y$ or $z$ value has magnitude $O\left(n^{2} N \log (n N)\right)$.

Proof. The result is proved in three steps. First we prove it for $y$ values. Define $N_{s}$ as the largest magnitude of a cost in scale $s$ : it is easy to see that $N_{s} \leq 2^{s+1}-2$. Let $Y_{s}$ denote the largest magnitude of a $y$ value in scale $s \geq 1$, and set $Y_{0}=0$. It suffices to verify the recurrence

$$
Y_{s} \leq 2 Y_{s-1}+1+10 n \log n+2 n N_{s} .
$$

This implies $Y_{s} \leq\left(2^{s}-1\right)(1+10 n \log n)+n s 2^{s+2}$. Hence in the last scale $Y_{s}=O\left(n^{2} N \log (n N)\right)$. This implies the desired bound for $y$.
To obtain the recurrence begin by observing that the match routine never increases a $y$ value: $y$ values change only in dual adjustments or translations and, if a dual adjustment increases $y(v)$ by $\delta$, the accompanying translation decreases $y(v)$ by $\delta$. Hence it suffices to examine the $y$ values at the end of the scale.
Let $\omega$ be the vertex that is free at the end of the scale. We show that at the end of the scale

$$
y(\omega) \geq y_{0}(\omega)-10 n \log n .
$$

In path $(R)$, suppose some shell_search $(C, D)$ adjusts duals by $\delta$. Then $y(\omega)$ does not change if $\omega \in V(C)-V(D)$ and it decreases by $2 \delta$ if $\omega \in V(D)$.

Thus the total decrease in $y(\omega)$ is at most the total translation of all even shells, which is at most $5 \hat{n}(R) \log \hat{n}(R)$ by (4). Summing over all major path roots $R$ containing $\omega$ gives a geometric progression with ratio $1 / 2$. Thus the total decrease is at most $10 n \log n$ as desired.

Now consider any vertex $x$, with matching $M_{x}$ at the end of the scale. As in Lemma 2.1, $c\left(M_{x}\right) \leq 2\lfloor n / 2\rfloor+y(V(G))-y(x)-\lfloor\hat{n} / 2\rfloor z(V(T))$, and $c\left(M_{x}\right) \geq y(V(G))-y(x)-\lfloor\hat{n} / 2\rfloor z(V(T))$. Thus any vertex $x$ has $y(x) \geq$ $y(\omega)-2\lfloor n / 2\rfloor+c\left(M_{\omega}\right)-c\left(M_{x}\right)$. Since $c\left(M_{x}\right) \in\left[-n N_{s} / 2 \ldots n N_{s} / 2\right]$, we deduce

$$
y(x) \geq y(\omega)-2 n N_{s}
$$

Combining this with the above inequality for $y(\omega)$ shows that any vertex $x$ has $y(x) \geq y_{0}(\omega)-10 n \log n-2 n N_{5}$. The Double Step shows that in scale $s$, $y_{0}(\omega)$ has magnitude at most $2 Y_{s-1}+1$. Together these imply the desired recurrence for $Y_{s}$.

It remains to analyze the magnitude of $z$ values. Consider first the value $z\left(G^{*}\right)$. This value is nonpositive-it can decrease in Double Steps and translations of shell_searches, but it never increases. In some scale $s$, let $e$ be a matched edge involved in the last blossom step (this step forms the blossom $G^{*}$ and ends the scale). Since $G^{*}$ is the only blossom containing $e$, (1a) implies that $z\left(G^{*}\right) \geq y(e)-N_{s}$. Thus $z\left(G^{*}\right)$ satisfies the lemma.

Finally consider any $z$ for nonroot blossoms. These values are nonnegative. Consider a nonroot blossom $B$. Without loss of generality $B$ is a current blossom, so it contains a matched edge $e$. (1b) for $e$ implies that in scale $s$, $y(e)+N_{s}+2-z\left(G^{*}\right) \geq z(\{C \mid e \subseteq C\})-z\left(G^{*}\right) \geq z(B)$. Thus $z(B)$ satisfies the lemma.

All other quantities computed in the algorithm are easily related to $y$ and $z$. This includes the quantities of Section $7, z_{1}, t$ and $y^{\prime}$. (Note that a $t$ value is bounded by a value $z\left(\left[G^{*}, B\right]\right)-z\left(G^{*}\right)$; the last paragraph of the proof shows that this satisfies the bound of the lemma.)

This completes the analysis of the scaling routine.
Theorem 9.1. The minimum critical matching problem can be solved in $O(\sqrt{n \alpha(m, n) \log n} m \log (n N))$ time and $O(m)$ space.

It is interesting that the proofs of the above lemma and Lemma 2.1 use the dual objective function $y(V(G))-\lfloor\hat{n} / 2\rfloor z(V(T))$. (This is the objective function of the linear programming dual of the matching problem [7].) It is tempting to analyze the matching algorithm using this dual objective function (as done in [10]). Here are some easily-proved facts: The dual objective does not decrease in path. In the entire execution of path the dual objective increases by $O(n)$ (from the Double Step). A dual adjustment of $\hat{o}$ in the search of a shell containing $f$ free vertices increases the dual objective by at least $(f-2) \delta$. These facts give a good bound on the time spent in shell_searches of shells with at least three free vertices. Unfortunately it is possible for an even shell to contain only two free vertices. Such shells do not seem amenable to an easy analysis. Hence the attractiveness of this approach remains unclear.

## 10. Other Matching Problems

This section gives applications of the minimum critical matching algorithm.
Theorem 10.1. A minimum perfect matching can be found in $O(\sqrt{n \alpha(m, n) \log n} m \log (n N))$ time and $O(m)$ space. The same bounds
apply to minimum-cost matching and minimum-cost maximum-cardinality matching.

Proof. The application to minimum perfect matching has been noted in Section 1.1. For minimum-cost matching, observe that a minimum-cost matching on $G$ corresponds to a minimum perfect matching on the graph formed by taking two copies of $G$ and joining each pair of copies of the same vertex by a zero-cost edge. Minimum-cost maximum-cardinality matching uses the same construction except that each added edge costs $n N$.

As observed in [10], cost-scaling algorithms can be used as approximation algorithms when input numbers are real, rational or very large integers. We illustrate this with the problem of finding an approximately-minimum perfect matching.
Consider a given cost function $c$ with nonnegative real values. Fix an integer $a$. We define a new cost function $c^{\prime}$ with values in $\left[0 \ldots n^{2+a}\right]$ such that if $M$ $\left(M^{\prime}\right)$ is a minimum perfect matching for $c\left(c^{\prime}\right)$, then $c\left(M^{\prime}\right) \leq\left(1+n^{-a}\right) c(M)$. To define $c^{\prime}$ let $B$ be the bottleneck cost of a minimum bottleneck matching; that is, let $B$ be the minimum value such that there is a perfect matching $A$ on the edges costing at most $B$. Assume $B>0$; otherwise the problem is trivial. Delete all edges costing more than $c(A)$; clearly these edges are not in a minimum-cost matching. Define cost function $c^{\prime}=\left\lfloor n^{1+a} c / B\right\rfloor$. Note that if $e$ is an edge that is not deleted then $c(e) \leq n B / 2$. Hence $c^{\prime}(e) \leq n^{2+a}$ as desired. Furthermore $c^{\prime}\left(M^{\prime}\right) \leq c^{\prime}(M)$ implies $n^{1+a} c\left(M^{\prime}\right) / B<n^{1+a} c(M) / B$ $+n / 2$. Since $c(M) \geq B, \quad c\left(M^{\prime}\right)<c(M)+B / n^{a} \leq\left(1+n^{-a}\right) c(M)$ as desired.

Theorem 10.2. Given arbitrary nonnegative edge costs and a positive integer $a$, a perfect matching costing at most $\left(1+n^{-a}\right)$ times minimum can be found in $O(a \sqrt{n \alpha(m, n) \log n} m \log n)$ time and $O(m)$ space.

Proof. First a bottleneck matching is found in time $O(\sqrt{n \log n} m)$ [14]. Then $c^{\prime}$ is computed in time $O(m)$. Finally the scaling algorithm is executed with costs $c^{\prime}$. This algorithm runs in the time of the theorem.
This leads to an efficient implementation of Christofides' approximation algorithm for a travelling salesman tour [4]. Recall that this approximation algorithm works as follows. Given are $n$ cities and the distance between each pair of cities. We assume that the distances satisfy the triangle inequality. The algorithm constructs a tour by finding a minimum spanning tree $T$, finding a minimum perfect matching $M$ on the odd-degree vertices of $T$, and reducing the Eulerian graph $T \cup M$ to a tour.

Recall the accuracy analysis of this algorithm: Let $H$ denote a minimumlength tour of the given cities. Let $c(e)$ denote the length of an edge $e$ joining two cities. The approximation algorithm gives a tour of length at most $c(T)+$ $c(M)$. It is easy to see that $c(T) \leq(1-1 / n) c(H)$ and $2 c(M) \leq c(H)$. This implies that $c(T)+c(M) \leq(3 / 2) c(H)$. Hence the algorithm gives a tour at most $3 / 2$ times optimum.

The running time of this algorithm is $O\left(n^{3}\right)$, the time to find the matching. We improve this by making one change: Instead of $M$ use a perfect matching that is at most $(1+1 / n)$ times minimum. It is easy to see that the resulting tour is at most $3 / 2-1 /(2 n) \leq 3 / 2$ times optimum. We find the approximatelyminimum matching using the algorithm of Theorem 10.2 .

Theorem 10.3. Christofides' approximation algorithm for a travelling salesman tour on $n$ cities, where distances obey the triangle inequality, can be implemented to take $O\left(n^{2.5}(\log n)^{15}\right)$ time and $O\left(n^{2}\right)$ space.

A number of applications of matching require not just an optimum matching but the output of Edmonds' algorithm, an optimum structured matching (recall the definition from Section 1.1). One example is updating a weighted matching: Suppose we have an optimum structured matching and the graph changes at one vertex $v$ (that is, edges incident to $v$ are added or deleted, and costs of edges incident to $v$ change). A new optimum structured matching can be found in the time for one search of Edmonds' algorithm [2, 5, 10, 29]. Another example is the single-source shortest path problem on undirected graphs with no negative cycles [ 21 , pp. 220-222]. We now give an algorithm to find an optimum structured matching.

The algorithm starts by executing the scaling routine with one change: The new cost function $\bar{c}$ is $(2 n+2) c$ (in Section $2, \bar{c}=(n+1) c$ ). Change the number of scales correspondingly to $k=\lfloor\log (n+1) N\rfloor+2$. Suppose the scaling routine halts with matching $M_{x}$, blossom tree $T$ and dual functions $y_{0}, z_{0}$. Our structured matching has the same matching and blossom tree. The dual function $y$ is defined by

$$
\begin{aligned}
& Y=y_{0}(V(G))-\lfloor\hat{n} / 2\rfloor z_{0}(V(T)) \\
& y=\left\lfloor\frac{y_{0}-Y}{2 n+2}\right\rfloor
\end{aligned}
$$

To define $z$, for each blossom $B$ choose (arbitrarily) a blossom edge $e_{B}$ of $B$. For a blossom $B$ with parent $A$,

$$
\begin{aligned}
z(B) & =(y-c)\left(e_{B}\right)-(y-c)\left(e_{A}\right) \\
z\left(G^{*}\right) & =(y-c)\left(e_{G^{*}}\right)
\end{aligned}
$$

To prove that the algorithm is correct, define dual functions $\bar{y}=(2 n+2) y$, $\bar{z}=(2 n+2) z$. It suffices to show that changing the duals to $\bar{y}, \bar{z}$ gives an optimum structured matching for the cost function used by the scaling routine, $\bar{c}$. This amounts to showing the following: (i) the duals are tight on every blossom edge; (ii) the duals are dominated on every edge; (iii) $\bar{z}(B) \geq 0$ unless $B=G^{*}$. The following proof depends on the fact that all values of $\bar{y}, \bar{z}$ and $\bar{c}$ are multiples of $2 n+2$.

We start by showing that for any vertex $v$.

$$
\begin{equation*}
y_{0}(v)-\bar{y}(v)-Y \in[0 \ldots n) \tag{6}
\end{equation*}
$$

First observe that the quantity $y_{0}(v)-Y+\bar{c}\left(M_{v}\right) \in[0 \ldots n)$, by (1a)-(1b) applied to the edges of $M_{v}$. Furthermore $\bar{c}\left(M_{\nu}\right)$ is a multiple of $2 n+2$. Since by definition $\bar{y}(v)=(2 n+2)\left\lfloor\left(y_{0}(v)-Y\right) /(2 n+2)\right\rfloor$, (6) follows.

To prove (i)-(ii) we first show that any edge $e$ satisfies the relation

$$
\left(\bar{y}-y_{0}\right) e-\left(\bar{z}-z_{0}\right)\left(\left[G^{*}, B\right]\right) \in[-2 n+2 \ldots 2 n]
$$

To derive this, first observe the tightness equation for any edge $e_{B}, \bar{y} \bar{z}\left(e_{B}\right)=$ $\bar{c}\left(e_{B}\right)$ (this follows from the definition of $z$ ). From (1a)-(1b), $y_{0} z_{0}\left(e_{B}\right)-$ $\bar{c}\left(e_{B}\right) \in[-2 \ldots 0]$. Subtract the previous tightness equation and use (6) on the
ends of $e_{B}$ to get $\left(\bar{z}-z_{0}\right)\left(\left[G^{*}, B\right]\right)+2 Y \in[-2 n . .0]$. The desired relation for $e$ follows by using this last relation plus (6) on the ends of $e$.

To prove (i), consider any blossom edge $e$ of $B$. Since $y_{0} z_{0}(e)-\bar{c}(e) \in$ [ $-2 \ldots 0$ ], the above relation for $e$ implies that $\bar{y} \bar{z}(e)-\bar{c}(e) \in[-2 n . .2 n]$. Since the left-hand expression is divisible by $2 n+2$, it equals zero, that is, the duals are tight on $e$. For (ii), similarly consider an edge $e$ such that $B$ is the smallest blossom containing it. Since $y_{0} z_{0}(e) \leq \bar{c}(e)$ it is easy to see that $\bar{y} \bar{z}(e)-\bar{c}(e) \leq 2 n$. Since the left-hand expression is divisible by $2 n+$ $2, \bar{y} \bar{z}(e) \leq \bar{c}(e)$, that is, the duals are dominated on $e$.

Lastly consider (iii). From (1a)-(1b), $\left(y_{0}-\bar{c}\right) e_{B}-\left(y_{0}-\bar{c}\right) e_{A}-z_{0}(B) \in$ [-2..2]. Combining this with the definition of $\bar{z}$ shows that $\left(\bar{z}-z_{0}\right)(B)+$ $\left(y_{0}-\bar{y}\right) e_{B}-\left(y_{0}-\bar{y}\right) e_{A} \in[-2 \ldots 2]$. Applying (6) to the ends of $e_{B}$ and $e_{A}$ implies that $\left(\bar{z}-z_{0}\right)(B) \in[-2 n \ldots 2 n]$. This implies that $\bar{z}(B) \geq-2 n$, since $z_{0}(B) \geq 0$. This gives the desired relation $\bar{z}(B) \geq 0$, since $\bar{z}(B)$ is divisible by $2 n+2$.

Theorem 10.4. An optimum structured matching can be found in the time and space bounds of Theorem 9.1.

The minimum critical matching algorithm can be modified to find a maxi-mum-cardinality matching on an arbitrary graph $G$. The cardinality matching algorithm works as follows. Define all edge costs to be zero. Execute the scaling routine by simply omitting the Double Step and doing the Match Step once, that is, call the match routine once to find the desired matching. Define $\rho$, the number of phase 1 iterations of path, to be $\lceil\sqrt{n}\rceil$. The remaining details of phase 1 are unchanged. (The algorithm works with one shell, $G^{*}$, which is even or odd depending on the number of vertices of G.) After phase 1 the algorithm is simpler than before. Instead of phases $2-3$ it abandons all edge costs and dual variables and does the following: Repeatedly call find_ap_set to find a maximal set of disjoint augmenting paths $\mathscr{P}$, and augment along these paths. find_ap_set operates on the graph $G$, unmodified. The algorithm halts when find_ap_set does not find an augmenting path.

The analysis of this algorithm is a special case of critical matching. We sketch it for completeness. First recall that the old blossom tree $T$ has root $G^{*}$ with children $V(G)$. The following version of Lemma 6.1 holds: At any time in the execution of $\operatorname{path}\left(G^{*}\right)$ let $M$ be the current matching. Let $M_{0}$ be a maximum-cardinality matching whose free vertices are all free in $M$. Let $F_{\omega}$ be the set of vertices that are free in $M$ but not $M_{0}$.

Lemma 10.1. At any time in path $\left(G^{*}\right)$ during phase 1 of the cardinality matching algorithm, $\hat{\delta}\left(F_{\omega}, G^{*}\right) \leq n$.
Proof. The proof is a special case of Lemma 6.1. As in Lemma 6.1, define $\mu(B)=\left|M_{0} \cap G(B)\right|-|M \cap G(B)|$. For the current duals $y, z$, adding (1a) for $M_{0}$ and subtracting (1b) for $M$, and recalling that all edge costs are zero, gives

$$
y\left(F_{\omega}\right)-\mu z\left(\left\{G^{*}\right\} \cup V(T)\right) \leq 2|M| .
$$

Note that $y$ is zero for any free vertex and $\mu$ is nonpositive for a current blossom (as in Lemma 6.1). Hence $-\mu z\left(G^{*}\right) \leq 2|M|$. Now the relations $z\left(G^{*}\right)=-2 \hat{\delta}\left(G^{*}\right), \mu\left(G^{*}\right)=\left|F_{\omega}\right| / 2$, and $|M| \leq n / 2$ imply the lemma.

At the end of phase 1 any free vertex $v$ has $\hat{\delta}\left(v, G^{*}\right)=\rho$. Thus $\left|F_{\omega}\right| \rho \leq n$. This implies that phase 1 ends with $O(\sqrt{n})$ more free vertices than a maxi-mum-cardinality matching. Thus find ap_set is executed $O(\sqrt{n})$ times.

As in the weighted algorithm, the time for a phase 1 iteration and the time for one execution of find_ap_set are both $O(m)$. Thus the total time is $O(\sqrt{n} m)$.

Theorem 10.5. A maximum cardinality matching can be found in $O(\sqrt{n} m)$ time and $O(m)$ space.

This bound is the same as the algorithm of Micali and Vazirani [23]. Note that our algorithm is not the same as theirs. For instance it operates with inner blossoms effectively "shrunk' during phase 1. Also our depth-first search seems to involve less overhead than the "double depth-first search" of [23].

In practice a different organization after phase 1 is probably faster: The algorithm initializes the search graph $\mathscr{J}^{\prime}$ to empty. Then it calls find_ap $(f)$ for each free vertex $f$. If find_ap( $f$ ) finds an augmenting path then the matching is augmented; also the vertices that it added to $\mathscr{F}^{\prime}$ are removed from $\mathscr{f}$. If find_ap(f) does not find an augmenting path then nothing is done-the vertices it added to $\mathscr{\prime}$ remain in $\mathscr{\mathscr { \prime }}$. These vertices will be effectively ignored in subsequent searches. This is correct because these vertices cannot be in augmenting paths (arguing as in Lemma 8.1; alternatively see [6] and [8].

## 11. Concluding Remarks

The matching algorithm generalizes to degree-constrained subgraph problems. Consider a graph with two functions $l, u: V \rightarrow \mathbf{Z}$. A degree-constrained subgraph ( $D C S$ ) is a subgraph in which each vertex $v$ has degree in the range $[l(v) \ldots u(v)]$. In a perfect $D C S$ each degree is exactly $u(v)$. The size of a perfect DCS is denoted $U=u(V)$. The weighted degree-constrained subgraph problem is to find a minimum-cost maximum-cardinality DCS or a minimum-cost DCS. A degree-constrained subgraph problem on a graph of $n$ vertices and $m$ edges can be reduced in linear time to a matching problem on a graph of $O(m)$ vertices and edges [11]. Thus our algorithm immediately implies a bound of $O(\sqrt{m \alpha(m, m) \log m} m \log (m N))$ to solve the weighted DCS problem. The same bound applies to the problem of finding a minimum cost flow on a $0-1$ bidirected network [21]. A more careful implementation of our ideas gives a bound of $O(\sqrt{U \alpha(m, n) \log U} m \log (n N))$ for the weighted DCS problem; this bound holds for multigraphs as well. Details will be given in a forthcoming paper.

Pravin Vaidya has recently investigated the matching problem for points in the plane. If distance is measured by the $L_{1}, L_{2}$ or $L_{\infty}$ norm, a minimum perfect matching on a set of $2 n$ points can be found in time $n^{2.5} \log ^{\circ(1)} n$ and space $O(n \log n)$ [28]. Furthermore it appears that applying our algorithm reduces the time by a factor of about $\sqrt{n}$ [28].

## REFERENCES

1. Aho. A. V.. Hopcroft, J. E., and Ullman, J. D. The Design and Analysis of Computer Algorithms, Addison-Wesley, Reading, Mass., 1974.
2. Ball, M. O., and Derios, U. An analysis of alternative strategies for implementing matching algorithms. Networks 13, 4 (1983), 517-549.
3. Bertsekas. D. P. Distrıbuted asynchronous relaxation methods for linear network flow problems, LIDS Report P-1606. M.I.T., Cambridge, Mass., 1986; preliminary version. In Proceeding of the 25th Conference on Decision and Control. IEEE, New York, December 1986.
4. Christofides, N. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical Report, Graduate School of Industrial Administration, Carnegie-Mellon Univ.. Pittsburgh, Pa., 1976.
5. Cunningham, W. H., and Marsh, A. B. III. A primal algorithm for optimum matching. Math. Prog. Stud. 8 (1978), 50-72.
6. Edmonds, J. Paths, trees and flowers. Canad. J. Math. 17 (1965), 449-467
7. Edmonds. J. Maximum matching and a polyhedron with 0,1-vertices. J. Res. Nat. Bur. Standards 69B (1965), 125-130.
7a. Edmonds, J., and Karp, R. M. Theoretical improvements in algorithmic efficiency for network flow problems. J. ACM 19. 0 (1972), 248-264.
8. Gabow, H. N. An efficient implementation of Edmonds' algorithm for maximum matching on graphs. J. ACM 23, 2 (1976), 221-234.
9. Gabow, H. N. Scaling algorithms for network problems. J. Comp. and System Sci., 31, 2 (1985), 148-168.
10. Gabow, H. N. A scaling algorithm for weighted matching on general graphs. In Proceedings of the 26th Annual Symposium on Foundations of Computer Science. IEEE. New York, 1985, pp. 90-100.
11. Gabow, H. N. Duality and parallel algorithms for graph matching, manuscript.
12. Gabow, H. N. Data structures for weighted matching and nearest common ancestors with linking. In Proceedings of the 1st Annual ACM-SIAM Symposium on Discrete Algorithms. ACM, New York. 1990, pp. 434-443.
13. Gabow, H. N., Galil. Z., and Spencer. T. H. Efficient implementation of graph algorithms using contraction. J. ACM 36, 3 (1989), 540-572.
14. Gabow. H. N., and Tarjan, R. E. A linear-time algorithm for a special case of disjoint set union. J. Comp. and System Sci.. 30, 2 (1985), 209-221.
15. Gabow, H. N.. and Tarjan, R. E. Algorithms for two bottleneck optimization problems. J. Algorithms, 9,3 (1988), 411-417.
16. Gabow. H. N., and Tarian, R. E. Faster scaling algorithms for network problems. SIAMJ. Comput., 18, 5(1989), 1013-1036.
17. Galil, Z., Micali, S., and Gabow, H. N. An $O(E V \log V)$ algorithm for finding a maximal weighted matching in general graphs. SIAM J. Comput., 15, 1 (1986), 120-130.
18. Goldberg, A. V. Efficient graph algorithms for sequential and parallel computers. Ph. D. Dissertation. Dept. of Electrical Eng. and Comp. Scı., MIT. Technical Rep. MIT /LCS/TR-374, Cambridge. Mass.. 1987.
19. Goldberg, A. V., and Tarjan, R E. Finding minimum-cost circulations by successive approximation. Math. of Oper. Res., 15. 3 (1990), 430-466.
20. Hopcroht. J.. and Karp. R. An $n^{5 / 2}$ algorithm for maximum matchings in bipartite graphs, SLAM J. Comput., 2, 4 (1973), 225-231.
21. Lawler, E. L. Combinatorial Optımization: Networks and Matroids. Holt, Rinehart and Winston, New York, 1976.
22. Lovász. L., and Plummer, M. D. Matching Theory North-Holland Mathematical Studies 121, Annals of Discrete Mathematics 29, North-Holland, New York, 1986.
23. Micali, S., and Vatirani, V V. An $O(\sqrt{|V|} \cdot|E|)$ algorithm for finding maximum matchng in general graphs. In Proceedings of the 21st Annual Symposium on the Foundations of Computer Science. IEEE. New York, 1980, pp. 17-27.
24. Papadimitriou, C. H.. and Steiglitz, K. Combinatorial Optimization: Algorithms and Complexity. Prentice-Hall, Inc, Englewood Cliffs, N.J., 1982.
25. Tardos. E. A strongly polynomial minimum cost circulation algorithm. Combinatorica 5, 3 (1985), 247-255.
26. Tarjan, R. E. Applications of path compression on balanced trees. J. ACM 26, 4 (1979), 690-715.
27. Tarjan, R. E. Data Structures and Network Algorithms. SIAM Monograph. Philadelpha, Pa., 1983.
28. Vaidya. P. M. Geometry helps in matching In Proceeding of the 20 th Annual ACM Symposium on Theory of Computing. ACM, New York, 1988, pp 422-425.
29. Weber, G. M. Sensitıvity analysis of optımal matchıngs, Networks II (1981), 41-56.

RECEIVED APRIL 1989; REVISED MAY 1990; ACCEPTED MAY 1990

