MATROIDS WITH PARITY CONDITIONS:
A NEW CLASS OF COMBINATORIAL OPTIMIZATION PROBLEMS

by

Eugene L. Lawler

Memorandum No. ERL-M334

November 22, 1971

ELECTRONICS RESEARCH LABORATORY

College of Engineering
University of California, Berkeley
94720

MATROIDS WITH PARITY CONDITIONS:
A NEW CLASS OF COMBINATORIAL OPTIMIZATION PROBLEMS

Eugene L. Lawler

Department of Electrical Engineering and Computer Sciences
and the Electronics Research Laboratory
University of California, Berkeley, California 94720

ABSTRACT

Let $M = (E, \mathcal{J})$ be a given matroid, and let π be a given partition of E which pairs the elements. I.e. each block of π contains exactly two elements e and $\bar{e}$; we call e the mate of $\bar{e}$ and vice versa. A set $A \subseteq E$ is said to be a parity set if, for each element e, $e \in A$ if and only if $\bar{e} \in A$. The matroid parity problem is to find an independent parity set with a maximum number of elements.

The matroid intersection problem and the nonbipartite matching problem are specializations of the matroid parity problem. A min-max duality theorem for the parity problem generalizes duality theorems for matroid intersection and for matching. An "augmentation" algorithm for the parity problem combines features of algorithms for matroid intersection and for matching. The algorithm is computationally efficient, provided there exists an efficient subroutine to test arbitrary subsets of elements for independence in the given matroid.

Research sponsored by the Air Force Office of Scientific Research, Grant AFOSR-71-2076.

1. PROBLEM DEFINITION

Let $M = (E, \mathcal{I})$ be a given matroid, and let π be a given partition of E which pairs the elements. I.e. each block of π contains exactly two elements e and $\bar{e}$; we call e the <u>mate</u> of $\bar{e}$, and vice versa. A set $A \subseteq E$ is said to be a <u>parity set</u> if, for each element e, $e \in A$ if and only if $\bar{e} \in A$. The <u>matroid parity problem</u> is to find an independent parity set with a maximum number of elements. (The "weighted" parity problem, as opposed to the "cardinality" problem, will be the subject of a separate paper.)

The matroid intersection problem [2]-[4] and the nonbipartite matching problem [1] are specializations of the matroid parity problem. A min-max duality theorem for the parity problem generalizes duality theorems for matroid intersection and for matching. An "augmentation" algorithm for the parity problem combines features of algorithms for matroid intersection and for matching. The algorithm is shown to be computationally efficient, provided there exists an efficient subroutine to test arbitrary subsets of elements for independence in the given matroid.

Before proceeding, we present some necessary definitions.

A <u>matroid</u> $M = (E, \mathcal{I})$ is a combinatorial structure in which E is a finite set of <u>elements</u> and $\mathcal{I}$ is a nonempty family of subsets of E (called <u>independent sets</u>) satisfying the axioms:

(1.1) If $I \in \mathcal{I}$ and $I' \subseteq I$, then $I' \in \mathcal{I}$.

(1.2) If I_p and I_{p+1} are sets in $\mathcal{I}$ containing respectively p and p+1 elements, then there exists an element $e \in I_{p+1} - I_p$ such that $I_p + e \in \mathcal{J}$.

(We use I + e and I - e to denote $I \cup \{e\}$ and $I - \{e\}$ respectively. We also denote the symmetric difference of two sets by "$\oplus$", and the number of elements in I by $|I|$.)

A maximal independent set is a base. For a given subset $A \subseteq E$, we call the cardinality of a maximal independent subset of A the rank of A, denoted r(A). (All maximal independent subsets of A have the same cardinality.) The span of A, denoted sp(A), is the unique maximal superset of A such that r(sp(A)) = r(A).

A set which is not independent (i.e. not in the family $\mathcal{J}$) is said to be dependent. A minimal dependent set is called a circuit. It is a basic theorem of matroid theory that if I is independent and I + e is dependent, then I + e contains precisely one circuit. If a subroutine exists for testing for independence, then the unique circuit in I + e can be discovered by removing one element at a time from I + e and testing for independence. If the removal of an element produces independence, the element is returned to the set. The subset remaining at the end is the unique circuit.

Intuitively, the span of A is the set which contains all the elements of A, together with all elements e which form circuits with subsets of A. Clearly, if A is independent and a subroutine exists for testing for independence, it is possible to compute the span of A by testing A + e, for all $e \notin A$. It is assumed in the statement of the algorithm that subroutines are available to test a given set for independence.

2. PROBLEM REDUCTIONS

Suppose π is an arbitrary partition of E, i.e. each of its blocks does not necessarily contain exactly two elements. An apparent generalization of the matroid parity problem is obtained by asking for an independent set I which contains a maximum number of elements, subject to the condition that I contains an even number of elements from each block of π. In fact, however, any parity problem of this type can be reduced to a parity problem in which the elements are paired.

Let $M = (E,\mathcal{I})$ and π, an arbitrary partition of E, be given. Let P be the family of all pairs $\{e_i,e_j\}$ such that e_i and e_j are contained in the same block of π. Create as many copies of each element e_i as there are pairs in P in which e_i is contained. Denote these new elements $e_i^{(k)}$, where k indexes the pairs in P and $e_i^{(k)}$ belongs to the k^{th} subset. Let E^* contain all such elements $e_i^{(k)}$, where $|E^*| = 2|P| \leq |E|^2 - |E|$.

Let $M^* = (E^*,\mathcal{I}^*)$ be the matroid obtained from M by letting all copies $e_i^{(k)}$ of e_i be "parallel" elements. I.e. if $\{e_{i(1)}, e_{i(2)}, \ldots, e_{i(p)}\}$ is an independent set of M, and $e_{i(1)}^{(k(1))}, e_{i(2)}^{(k(2))}, \ldots, e_{i(p)}^{(k(p))}$ are elements of E^*, then $\{e_{i(1)}^{(k(1))}, e_{i(2)}^{(k(2))}, \ldots, e_{i(p)}^{(k(p))}\}$ is an independent set of M^*, whereas $\{e_i^{(k)}, e_i^{(k')}\}$ is a circuit of M^* for any i, $k \neq k'$. Let π^* be the partition of E^*, each block of which contains exactly two elements $e_i^{(k)}, e_j^{(k)}$ (corresponding the k^{th} pair $\{e_i,e_j\}$ contained in P).

For each independent set I of M satisfying the parity conditions given by π, there corresponds at least one independent set I^* of M^* (with $|I| = |I^*|$) satisfying the parity conditions given by π^*. For each independent set I^* of M^* satisfying the parity conditions given by π^* there corresponds a unique independent set I of M satisfying the parity conditions given by π.

Now consider the reduction of the matching problem to a matroid parity problem.

Let $G = (N,A)$ be a graph with node set N and arc set A. We say that a subset $S \subseteq A$ is a <u>matching</u> in G if no two arcs in S are incident to the same node. The (cardinality) <u>matching problem</u> is to find a matching with a maximum number of arcs.

Replace each of the m arcs of G by a pair of arcs e and $\bar{e}$, with a new node between them, thereby obtaining the subdivision graph G'. Let E be the set of 2m arcs of G', and let $\mathcal{I}$ contain all subsets $I \subseteq E$, such that no two arcs of I are incident to the same node of G', unless it is one of the nodes created by subdivision. Then $M = (E,\mathcal{I})$ is a matroid, and each independent parity set of M is identified with a matching in G, and conversely. Therefore the matching problem is a matroid parity problem.

We can also characterize matroid parity problems which can be reduced to matching problems, as follows. Let E be a finite set and τ be an arbitrary partition of its elements into blocks, $B_1, B_2, \ldots, B_k$. Let

$$\mathcal{I} = \{I \subseteq E \,\big|\, |I \cap B_i| \leq 1,\ i = 1, 2, \ldots, k\}.$$

Then $M = (E,\mathcal{I})$ is a <u>partition matroid</u> (defined by the partition τ). A matroid parity problem is equivalent to a matching problem if and only if the given matroid is a partition matroid.

Finally, we illustrate the reduction of the matroid intersection problem to a matroid parity problem.

Let $M_1 = (E, \mathcal{I}_1)$ and $M_2 = (E, \mathcal{I}_2)$ be two matroids over the same set of elements E. We say that $I \subseteq E$ is an <u>intersection</u> of the two matroids if I is an independent set of both M_1 and M_2. The (cardinality) <u>matroid intersection problem</u> is to find an intersection of the two matroids with a maximum number of elements.

Replace each element e of M_2 by a distinct element $\bar{e} \notin E$, so that M_2 becomes a matroid over a new set $\bar{E}$, disjoint from E. Next take the sum of M_1 and M_2. The sum $M_1 + M_2 = (E \cup \bar{E}, \mathcal{I})$, where

$$\mathcal{I} = \{I_1 \cup I_2 | I_1 \in \mathcal{I}_1, \ I_2 \in \mathcal{I}_2\},$$

is a matroid [2]. Each independent parity set of M is identified with an intersection of M_1 and M_2, and conversely. Therefore, the matroid intersection problem is a matroid parity problem.

We can also characterize matroid parity problems which can be reduced to intersection problems, as follows. Let $M = (E, \mathcal{I})$ be an arbitrary matroid and let π be a partition of E which pairs the elements. If there exists a partition τ with exactly two blocks, such that (1) e and $\bar{e}$ do not belong to the same block of τ, for any e, and (2) each circuit of M is entirely contained within a single block of τ, then we say that M is <u>separable with respect to π</u>.

A matroid parity problem is equivalent to a matroid intersection problem if the partition π is such that M is separable with respect to π.

If in addition M is a partition matroid then the parity problem is equivalent to a bipartite matching problem. Separability is a matroid analog of graphic bipartiteness.

3. AUGMENTING SEQUENCES

Let $S = (e_1, \bar{e}_1, e_2, \bar{e}_2, \ldots, e_m, \bar{e}_m)$ be a sequence of distinct elements. Let $S_i = \{e_1, \bar{e}_1, e_2, \bar{e}_2, \ldots, e_{i-1}, \bar{e}_{i-1}, e_i\}$, and $\bar{S}_i = S_i + \bar{e}_i$. Let I be an independent parity set. S is said to be an alternating sequence with respect to I if

(3.1) $$e_i, \bar{e}_i \in \begin{cases} E - I, & \text{for } i \text{ odd}, \\ I, & \text{for } i \text{ even}, \end{cases}$$

(3.2) $I + e_1$ is independent,

(3.3) $sp(I \oplus S_i) = sp(I + e_1)$, for all i,

Note that, for all even i, $I \oplus S_i$ is necessarily independent and $|I \oplus S_i| = |I| + 1$. It follows that, for all odd $i < m$, $I \oplus \bar{S}_i$ is dependent.

An alternating sequence S is said to be an augmenting sequence with respect to I if, in addition,

(3.4) m is odd

(3.5) $I \oplus \bar{S}_m$ is independent.

(Note that $|I \oplus \bar{S}_m| = |I| + 2$ and that $sp(I \oplus \bar{S}_m) \supseteq sp(I)$.)

Theorem 3.1

Let I_{2k} and I_{2k+2} be independent parity sets with 2k and 2k + 2 elements respectively. Then there exists an augmenting sequence $S \subseteq I_{2k} \oplus I_{2k+2}$ with respect to I_{2k}.

Proof:

Proof is by induction on $|I_{2k} \oplus I_{2k+2}|$. If $|I_{2k} \oplus I_{2k+2}| = 2$, it is clear that $I_{2k} \oplus I_{2k+2} = \{e_1, \bar{e}_1\}$ yields an augmenting sequence. Suppose $|I_{2k} \oplus I_{2k+2}| = 2p > 2$. By matroid axiom (1.2), there must exist an element $e_1 \in I_{2k+2} - I_{2k}$ such that $I_{2k} + e_1$ is independent. If $I_{2k} + e_1 + \bar{e}_1$ is independent then $S = (e_1, \bar{e}_1)$ is an augmenting sequence. If $I_{2k} + e_1 + \bar{e}_1$ is dependent, then there exists a circuit $C \subseteq I_{2k} + e_1 + \bar{e}_1$, with $C - I_{2k+2} \neq \phi$. Let e_2 be any element belonging to $C - I_{2k+2}$. Then $(e_1, \bar{e}_1, e_2, \bar{e}_2)$ is an alternating sequence, and $I'_{2k} = I_{2k} + e + \bar{e}_1 - e_2 - \bar{e}_2$ is independent. By inductive assumption, there exists an augmenting sequence with respect to I'_{2k}, since $I'_{2k} \oplus I_{2k+2} = 2p - 4$. Let $S' = (a_1, \bar{a}_1, a_2, \bar{a}_2, \ldots, a_m, \bar{a}_m)$ be such a sequence. There are two cases to consider:

Case 1: There is an element $a_p \in S'$ such that $I + a_p \in \mathcal{I}$. Let a_p be the element of S' with largest index for which this is the case. Then $S = (a_p, \overline{a_p}, a_{p+1}, \overline{a_{p+1}}, \ldots, a_m, \overline{a_m})$ is an augmenting sequence with respect to I_{2k}.

Case 2: There is no such element a_p. Then $S = (e_1, \overline{e_1}, e_2, \overline{e_2}, a_1, \overline{a_1}, a_2, a_2, \ldots, a_m, \overline{a_m})$ is an augmenting sequence with respect to I_{2k}.

Corollary 3.2

An independent parity set I contains a maximum number of elements if and only if there exists no augmenting sequence with respect to I.

Corollary 3.3

For any independent parity set I there exists an independent parity set I^* with a maximum number of elements, such that $sp(I) \subseteq sp(I^*)$.

4. BLOSSOMS

Let I be an independent parity set, B a parity set, and S an alternating sequence. B is said to be a blossom with respect to I, and S, an alternating sequence, is its stem, if

(4.1) $r(B) = |I \cap B| + 1$ (from which it follows that B is nonempty and $r(B)$ is odd).

(4.2) S is disjoint from B and contains an even number of element pairs (S may be empty).

(4.3) for each element $e_i \in sp(I) - sp(I - B)$, $e_i \notin I$, there exists an alternating sequence of the form $(S;B_i;e_i,\bar{e}_i)$, where $B_i \subseteq B$.

Given B, it follows from (4.1) that I is such that $|I \cap B|$ is maximal. Now suppose there is an augmenting sequence of which S is a prefix and in which e_i, $\bar{e}_i$ are contained, in that order, where $e_i \in sp(I) - sp(I - B)$ and $e_i \notin I$. From (4.2) and (4.3) it follows that there exists such an augmenting sequence for which $|I \cap B|$ remains maximal after augmentation.

Blossoms do not exist for matroid parity problems which are separable with respect to the given parity conditions. The existence of blossoms distinguishes nonseparable matroid parity problems from matroid intersection problems, just as the existence of blossoms distinguishes nonbipartite matching problems from bipartite matching problems.

The terms "blossom" and "stem" were originated by Edmonds [1] for matching problems, where the terms are suggested by the appearance of these structures in a drawing of the graph. It is also intuitively appealing to make drawings of stems and blossoms for the matroid parity problem.

Let elements be represented by arcs or lines. An element e contained in I will be drawn as a wavy line, one not in I as a straight line. Thus a typical alternating sequence is drawn as shown in Figure 1(a). We shall attempt to construct augmenting sequences by extending the length of alternating sequences that have already been constructed. Since a given alternating sequence can serve as the prefix of many other alternating sequences, the result of this construction is a tree structure, as shown in Figure 1(b).

Suppose a computation is undertaken in which alternating sequences are extended, in an effort to construct an augmenting sequence. And suppose in the course of this computation two alternating sequences P and Q are found, where

$$P = (p_1,\bar{p}_1,p_2,\bar{p}_2,\ \ldots,\ p_k,\bar{p}_k,e_i,\bar{e}_i),$$

$$Q = (q_1,\bar{q}_1,q_2,\bar{q}_2,\ \ldots,\ q_\ell,\bar{q}_\ell,\bar{e}_i,e_i),$$

and where the following properties are satisfied:

(4.4) $\quad sp(I + p_1) = sp(I + q_1)$

(4.5) $\quad$ except for e_i, $\bar{e}_i$, any element pair contained in both P and Q appears in the same order in each of them.

When such sequences P and Q are found, a blossom has been detected, and it can be constructed as follows.

Because of the way in which the augmenting sequences are constructed (cf. the tree structure of Figure 1(b)), we may assume that $T = P \cap Q - \{e_i, \bar{e}_i\}$ is a prefix of both sequences. I.e. there is an integer m such that

$$P = (T; p_{m+1}, \bar{p}_{m+1}, \ldots, p_k, \bar{p}_k, e_i, \bar{e}_i)$$

$$Q = (T; q_{m+1}, \bar{q}_{m+1}, \ldots, q_\ell, \bar{q}_\ell, \bar{e}_i, e_i)$$

If T is empty (i.e. P and Q are disjoint except for e_i, $\bar{e}_i$), the stem of the blossom is empty, and we say that the blossom is rooted. For each element $e_j \in (P \cup Q) - \{p_1, q_1\}$, $e_j \notin I$, find the unique circuit $C_j \subseteq I + e_j$. Let $par(C_j)$ denote the smallest parity set containing C_j. Then the rooted blossom B is the union of $P \cup Q = (P \oplus Q) \cup \{e_i, \bar{e}_i\}$ and $par(C_j)$, for all C_j.

If T is nonempty and contains an even number of element pairs, then T is the stem of the blossom and we say that the blossom is short. For each element $e_j \in (P \oplus Q) \cup \{e_i, \bar{e}_i\}$, $e_j \notin I$, find the unique circuit

$C_j \subseteq I + e_j$. Let B' be the union of $(P \oplus Q) \cup \{e_i, \bar{e}_i\}$ and $\text{par}(C_j)$, for all C_j. Then the short blossom $B = B' - \{\bar{p}_m\}$, where $\{p_m, \bar{p}_m\} = \{q_m, \bar{q}_m\}$ is the final element pair of the stem.

If T contains an odd number of element pairs, then the stem of the blossom contains all but the final pair $\{p_m, \bar{p}_m\} = \{q_m, \bar{q}_m\}$ of T, and we say that the blossom is <u>tall</u>. The final element pair of T is called the <u>base</u> of the tall blossom. For each element $e_j \in (P \oplus Q) \cup \{e_i, \bar{e}_i, \bar{p}_m\}$, $e_j \notin I$, find the unique circuit $C_j \subseteq I + e_j$. The tall blossom B is the union of $(P \oplus Q) \cup \{e_i, \bar{e}_i, p_m, \bar{p}_m\}$ and $\text{par}(C_j)$, for all C_j.

Blossoms can have a rather complex structure, as indicated schematically in Figues 2 - 4. However, in each case it is possible to verify that the sets B, as constructed above, satisfy conditions (4.1) - (4.3) of the definition of a blossom. For example, to show that (4.1) is satisfied in the case of a rooted blossom, let $b' = B - \{p_1, q_1\}$. It follows from the construction of the circuits C_j that $r(B') = |I \cap B'|$. Hence $r(B) = |I \cap B| + 1$, because $I + p_1$ and $I + q_1$ are independent and $sp(I + p_1) = sp(I + q_1)$, by assumption. In the case of a short blossom, let $B' = B + \bar{p}_m$, where $\{p_m, \bar{p}_m\} = \{q_m, \bar{q}_m\}$ is the final element pair of the stem. Then $r(B') = |I \cap B'| = r(B) = |I \cap B| + 1$. In the case of a tall blossom, let $B' = B - p_m$, where $\{p_m, \bar{p}_m\}$ is the base of the blossom. Then $r(B') = |I \cap B'| = r(B) - 1 = |I \cap B|$.

When blossoms are detected in the course of the construction of alternating sequences, we "shrink" them, thereby obtaining a (smaller) problem on a matroid over the elements E - B. We then proceed with the construction of alternating sequences, as though the new matroid were separable with respect to the parity conditions. That is, until another blossom is detected and shrunk, etc.

As implied by condition (4.3), the shrinking of blossoms should be such that for every alternating sequence passing through the blossom before shrinking, there is a corresponding alternating sequence after shrinking. For example, in place of the sequence $(e_1, \bar{e}_1, \ldots, e_7, \bar{e}_7)$ in Figure 4, there shou d be the sequence $(e_1, \bar{e}_1, e_2, \bar{e}_2, e_7, \bar{e}_7)$ after shrinking.

Let e_m, $\bar{e}_m$ be the final element pair of the stem of the blossom B, if the stem is nonempty. Let

$$D = (sp(I) - sp(I - B)) - I + \bar{e}_m,$$

if the stem is nonempty, and

$$D = (sp(I) - sp(I - B)) - I,$$

otherwise. Let M(D) be the partition matroid which has as its bases all singleton subsets of D. We <u>shrink</u> the blossom B by replacing the matroid M by

$$M_B = M \text{ ctr } B + M(D),$$

where M ctr B denotes the <u>contraction</u> of B in M. I.e. if $M = (E, \mathcal{I})$, then $M \text{ ctr } B = (E - B, \mathcal{I}')$, where

$$\mathcal{I}' = \{I' | r(I' \cup B) = r(I') + r(B)\}$$

The effect of shrinking a blossom is indicated in Figure 5. In this figure,

$$D = \{\bar{e}_m, d_1, d_2, d_3\}.$$

Theorem 4.1

There exists an augmenting sequence with respect to I in M if and only if there exists an augmenting sequence with respect to I - B in M_B.

We outline a proof of the theorem as follows. First show that for any sequence S' in M there is a corresponding sequence in M_B. If S' is disjoint from B in M, then S' exists in M_B. If S' is not disjoint from B and has the stem of B as a prefix, then S' - B is a sequence in M_B. If S' is not disjoint from B and does not have the stem S of B as a prefix, then there is a sequence in M_B which does have S as a prefix. Examples of these latter two cases are illustrated in the upper and lower drawings in Figure 6 respectively.

The converse, i.e. that for any sequence S' in M_B there is a corresponding sequence in M, is perhaps a bit simpler to prove. If $\bar{e}_m \notin S'$, then S' is a sequence in M. If $\bar{e}_m \in S'$, then there is a sequence passing through B (by condition (4.3)), in M.

5. LABELLING PROCEDURE

The construction of augmenting sequences can be carried out by a labelling procedure similar to that employed for the matroid intersection problem [3], [4]. We first consider the adaptation of that procedure to the parity problem, without considering the complications introduced by the existence of blossoms.

An element is labelled when it is found to be contained in an alternating sequence. The label identifies the previous element pair in the sequence. Thus, if e_i is given the label (j), the sequence contains e_j, $\bar{e}_j$, e_i, $\bar{e}_i$, in that order; if it is given the label ($\bar{j}$), the sequence contains $\bar{e}_j, e_j, e_i, \bar{e}_i$. In addition, we append a minus or a plus to a label to remind us whether or not the element in question belongs to I. Thus, the label $(j)^-$ on e_i means that $e_i \in I$; $(j)^+$ means that $e_i \notin I$.

The labelling procedure begins with the application of the label $(\phi)^+$ to an element e_i in E - sp(I) (and also to all other elements in $sp(I + e_1) - sp(I)$). The label $(\phi)^+$ indicates that the element has no predecessor in an alternating sequence. Additional elements are labelled (i.e. alternating sequences are extended) by "scanning" existing labels. A "+" label on e_i is scanned by first determining if $I + e_1 + \bar{e}_i$ is independent. In this case an augmenting sequence has been discovered, with e_i, $\bar{e}_i$ as the final pair of elements. If $I + e_1 + \bar{e}_i$ is dependent, the unique circuit $C \subseteq I + e_1 + \bar{e}_i$ is found, and the label $(i)^-$ is given to each unlabelled element in C. A "-" label on e_i is scanned by giving each unlabelled element e_j in $sp(I) - sp(I - \bar{e}_i)$ the label i^+.

The labelling procedure terminates when no further elements can be labelled or when an augmenting sequence is discovered, as described above. The complete augmenting sequence can be obtained by "backtracking". I.e. if the label of e_i is $(j)^+$, the second-to-last element pair in the sequence is e_j, $\bar{e}_j$. If the label of e_j is $(\bar{k})^-$, the third-to-last element pair is $\bar{e}_k$, e_k, etc. The initial element of the sequence, of course, has the label $(\phi)^+$.

It is not difficult to verify that the above rules are sufficient to construct augmenting sequences in the case that no blossoms are encountered.

A blossom is detected whenever both elements e_i and $\bar{e}_i$ of a pair receive labels. The sequences P and Q (terminating in e_i, $\bar{e}_i$ and $\bar{e}_i$, e_i respectively), described in the previous section, can be found by backtracking from e_i and $\bar{e}_i$, respectively. One can then construct the blossom B, its stem S, and, if the blossom is tall, its base, according to the rules indicated in the previous section.

There is little difficulty in applying the same labelling procedure to M_B = M ctr B + M(D) and I - B after the blossom B is shrunk. In particular, a subroutine for independence testing in M can be adapted to independence testing in M_B, for the sets required by the algorithm. However, there is a problem involved in constructing an augmenting sequence in M from an augmenting sequence in M_B. I.e. one must know how to find a "path" through the blossom B.

We cope with this problem by giving each element e_i two labels, or rather, a label with two components. The first component indicates the predecessor of e_i in an alternating sequence in M and the second component indicates the predecessor of e_i in an alternating sequence in M_B. The first component is used for backtracking to construct an augmenting sequence in M, whereas the second component is used for backtracking to construct a blossom in M_B.

For example, in the lower drawing in Figure 2, the label given to e_3, after the shrinking of the blossom, is $(2,\phi)^+$; in Figure 3 the label given to e_5 will be $(4,2)^+$ (or $(\bar{4},2)^+$); in Figure 4 the label given to e_7 will be $(6,2)^+$. The second component of a label may be revised with the detection and shrinking of further blossoms.

6. CARDINALITY PARITY ALGORITHM

The complete matroid parity algorithm is summarized as follows.

Step 0 (Start)

Let I_0 be any independent parity set (possibly the empty set) in the matroid $M_0 = (E_0, \mathcal{J}_0)$. Set $b = 0$. No elements are labelled. (All references to independence, circuits, spans, etc. are with respect to M_b.)

Step 1 (Labelling)

1.0 If there are no unlabelled elements in $E_b - sp(I_b)$, go to Step 4 (I_0 contains a maximum number of elements). Otherwise, find an unlabelled element e_i in $E_b - sp(I_b)$ and give it, and all other elements in $sp(I_b + e_i) - sp(I_b)$, the label $(\phi,\phi)^+$. Set $I' = I_b + e_i$.

1.1 If all labels have been scanned, go to Step 1.0. Otherwise, find an element e_i with an unscanned label. If $\bar{e}_i$ is also labelled, go to Step 3 (a blossom has been detected). Otherwise, if the label of e_i is a "+" label, go to Step 1.2; if it is a "-" label, go to Step 1.3.

1.2 Scan the "+" label of e_i as follows. If $I' + \bar{e}_i$ is independent, go to Step 2 (an augmenting sequence has been discovered). Otherwise, identify the unique circuit C in $I' + \bar{e}_i$ and give each unlabelled element in $C - \bar{e}_i$ the label $(i,i)^-$. Return to Step 1.1.

1.3 Scan the "-" label of e_i as follows. Give each unlabelled element in $sp(I') - sp(I' - \bar{e}_i) - \bar{e}_i$ the label $(i,i)^+$. Return to Step 1.1.

Step 2 (Augmentation)

An augmenting sequence with respect to I_0 has been discovered, with e_i, $\bar{e}_i$ (found in Step 1.2) as the final element pair. The preceeding elements in the sequence are discovered by "backtracking," using first components of labels. I.e. if the first component of the label of e_i is j, then the second-to-last element pair is $e_j, \bar{e}_j$. If the first component of the label of e_j is $\bar{k}$, then the third-to-last element pair is $\bar{e}_k$, e_k, etc. The initial element in the sequence has ϕ as the first component of its label. Augment I_0 by adding to I_0 all elements in the sequence with "+" labels and removing from I_0 all elements with "-" labels. Set b = 0. Remove all labels from elements. Return to Step 1.0.

Step 3 (Blossoming)

3.0 A blossom has been detected, containing the pair e_i, $\bar{e}_i$ found in Step 1.1. Backtrack from e_i and from $\bar{e}_i$, using second components of labels, thereby obtaining sequences P and Q, where

$$P = (p_1, \bar{p}_1, p_2, \bar{p}_2, \ldots, p_k, \bar{p}_k, e_i, \bar{e}_i),$$

$$Q = (q_1, \bar{q}_1, q_2, \bar{q}_2, \ldots, q_e, \bar{q}_e, \bar{e}_i, e_i).$$

Find $T = P \cap Q - \{e_i, \bar{e}_i\}$,

$$= (p_1,\bar{p}_1,p_2,\bar{p}_2, \ldots, p_m,\bar{p}_m),$$

$$= (q_1,\bar{q}_1,q_2,\bar{q}_2, \ldots, q_m,\bar{q}_m).$$

If T is empty, go to Step 3.1.0. If T is nonempty and contains an even number of element pairs, go to Step 3.1.1. If T contains an odd number of element pairs, go to Step 3.1.2.

3.1.0 (The blossom is rooted.) For each element $e_j \in (P \cup Q) - I'$, identify the unique circuit $C_j \subseteq I' + e_j$. Set $B_b = \bigcup_j \text{par}(C_j)$. Set $D_b = (\text{sp}(I') - \text{sp}(I' - B_b)) - B_b$. The labels of all elements in B_b are considered to be unscanned. Go to Step 3.2.

3.1.1 (The blossom is short.) For each element $e_j \in (P \oplus Q) \cup \{e_i,\bar{e}_i\} - I_b$, find the unique circuit $C_j \subseteq I' + e_j$. Set $B_b = \bigcup_j \text{par}(C_j) - \{p_m,\bar{p}_m\}$. Set $D_b = (\text{sp}(I_b) - \text{sp}(I_b - B_b)) - B_b + \bar{p}_m$. The labels of all elements in B_b are considered to be unscanned. Go to Step 3.2.

3.1.2 (The blossom is tall.) For each element $e_j \in (P \oplus Q) \cup \{e_i,\bar{e}_i,\bar{p}_m\} - I_b$, find the unique circuit $C_j \subseteq I' + e_j$. Set $B_b = \bigcup_j \text{par}(C_j)$. Set $D_b = (\text{sp}(I_b) - \text{sp}(I_b - B_b)) - B_b + \bar{p}_{m-1}$, if $m > 1$, and $D_b = (\text{sp}(I') - \text{sp}(I'-B_b)) - B_b$, if $m = 1$. The labels of all elements in B_b are considered to be unscanned. Go to Step 3.2.

3.2 If all labels of elements in B_b have been scanned, go to Step 3.4. Otherwise, find an element $e_i \in B_b$ with an unscanned label. If the label of e_i is a "+" label, go to Step 3.3; if it is a "-" label, go to Step 3.5.

3.3 Scan the "+" label of e_i as follows. If $I' + \bar{e}_i$ is not independent, identify the unique circuit C in $I' + \bar{e}_i$. Give each unlabelled element in $(C - \bar{e}_i) - T$ the label $(i,i)^-$. (Each element so labelled is in B_b; the second component of the label will never be used.) Return to Step 3.2.

3.4 Scan the "-" label of e_i as follows. Let $e_m' = p_m$ and $e_{(m-1)'} = P_{m-1}$. $((m-1)' = \phi$ if $m = 1.)$ Give each unlabelled element in $(sp(I') - sp(I' - \bar{e}_i) - \bar{e}_i) - T$ the label $(i,\phi)^+$ if B_b is rooted, $(i,m')^+$ if B_b is tall. For each labelled element in $(sp(I') - sp(I' - \bar{e}_i) - T$ set the second component of the label to ϕ if B_b is rooted, m' if B_b is short, and $(m - 1)'$ if B_b is tall. (If such a label has been scanned, it is still considered to have been scanned.) Return to Step 3.2.

3.5 Set $M_{b+1} = M_b$ ctr $B_b + M(D_b)$, $I_{b+1} = I_b - B_b$, and $E_{b+1} = E_b - B_b$. If B_b is rooted, or tall with $m = 1$, set $I' = (I' - B_b) + e_j$, where e_j is any element in D_b; otherwise set $I' = I' - B_b$. Set $b = b + 1$. Return to Step 1.1.

<u>Step 4</u> (Hungarian Labelling)

The labelling is "Hungarian." I_0 contains a maximum number of elements. (A dual solution can be constructed from the labelling and from the blossoms, as described in Section 7.) The computation is completed.

7. MIN-MAX DUALITY THEOREM

A duality theorem can be proved for the matroid parity problem. This theorem is of the same character as the min cut-max flow theorem for network flows, the König-Egervary theorem for bipartite matchings, Edmonds' odd-set covering theorem for nonbipartite matchings, etc. In fact, these other theorems can be considered to be corollaries of the matroid parity duality theorem.

Let $A_1, A_2, \ldots, A_m$ be subsets of E. We say that the collection

$A_1, A_2, \ldots, A_m$ is a parity covering of E if

$$E = \bigcup_{i=1}^{m} \mathrm{par}(A_i),$$

and we define the rank of the covering to be

$$\sum_{i=1}^{m} r^*(A_i),$$

where, for i=1,2,...,m, $r^*(A_i)$ is the largest even integer not greater than $r(A_i) + r(\tilde{A}_i)$, and where $\tilde{A}_i \subseteq A_i$ is the subset of elements of A_i whose mates are not contained in A_i, i.e.

$$\tilde{A}_i = \{e_i \mid e_i \in A_i,\ \bar{e}_i \notin A_i\}.$$

Lemma 7.1

The rank of a covering cannot be strictly less than the cardinality of an independent parity set.

Proof:

Let I be an independent parity set and $A_1, A_2, \ldots, A_m$ be a covering. Then it must certainly be the case that

$$|I \cap A_i| \leq r(A_i)$$

and

$$|I \cap \tilde{A}_i| = |I \cap (\mathrm{par}(A_i) - A_i)| \leq r(\tilde{A}_i).$$

Thus, for i = 1,2,...,m,

$$|I \cap \mathrm{par}(A_i)| \leq r^*(A_i).$$

Since $A_1, A_2, \ldots, A_m$ is a covering,

$$|I| \leq \sum_{i=1}^{m} |I \cap \mathrm{par}\ (A_i)| \leq \sum_{i=1}^{m} r^*(A_i)$$

and the lemma is proved.

The duality theorem asserts that there exists a covering and an independent parity set that satisfy the inequality of Lemma 7.1 with strict equality. This can be proved constructively, using the "Hungarian" labelling that exists at the end of the matroid parity algorithm. However, in order to do this, we must first show how to transform B_0, B_1, ..., B_{b-1}, which are blossoms for the matroids M_0, M_1, ..., M_{b-1}, into a set of blossoms for M_0. (This construction is also essential for the computation for the weighted parity problem.)

Let "~" be a relation on the indices of blossoms, where $i \sim j$ if either $i = j$ or $D_i \cap B_j \neq \phi$. (Note that $i \sim j$ implies $i \leq j$. Hence the transitive closure of "~" is a partial ordering.) Let

$$B_0^* = B_0$$

$$B_j^* = \bigcup_{i \sim j} B_i^* \qquad (j = 1, 2, \ldots, b - 1).$$

Lemma 7.2

Each of the sets B_j^* is a blossom of M_0. Moreover, the maximal blossoms B_j^* are disjoint.

We omit the proof of Lemma 7.2.

Theorem 7.3 (Matroid Parity Duality Theorem)

The minimum rank of a parity covering is equal to the maximum cardinality of an independent parity set.

Proof:

Consider the "Hungarian" labelling that exists at the end of the algorithm. Let $L \subseteq I_0$ be the set of elements of I_0 which are labelled, but do not belong to any of the blossoms $B_0, B_1, \ldots, B_{b-1}$. Let $\mathbf{U}$ be the (parity) set of all elements e_i such that neither e_i nor its mate $\bar{e}_i$ is labelled. Set $A_1 = sp(L) \cup \mathbf{U}$. Let $A_2, A_3, \ldots, A_m$ be the blossoms which are maximal within the collection $B_0^*, B_1^*, \ldots, B_{b-1}^*$, provided $b \geq 1$.

It can be verified that $par(A_1)$ contains all elements not in blossoms. Thus, $A_1, A_2, \ldots, A_m$ is a parity covering of E. Moreover, from the construction, the sets $par(A_1), par(A_2), \ldots, par(A_m)$ are disjoint, and, for $i = 1,2,\ldots, m$

$$|I_0 \cap A_i| = r^*(A_i).$$

Thus,

$$|I_0| = \sum_{i=1}^{m} r^*(A_i),$$

and the theorem is period.

We now indicate how the duality theorems for graphical matchings and matroid intersections follow as corollaries of Theorem 7.3.

Following Edmonds [1] terminology for matchings in graphs, a set of $2r + 1$ nodes is said to have (odd-set) rank one if $r = 0$ and rank r otherwise. If $r = 0$, the set is a singleton set and is said to cover all the arcs incident to that single node. If $r \geq 1$, the set covers all arcs which are incident to two nodes in the set. A collection of odd sets is an odd-set covering of the arcs, if every arc is covered by at least one set in the collection.

Corollary 7.4 (Odd-Set Covering Theorem [1])

For a given graph G, the minimum rank of an odd-set covering of the arcs of G is equal to the maximum cardinality of a matching in G.

Proof:

Reduce the matching problem to a matroid parity problem by the construction of Section 2, and apply Theorem 7.3. A minimal rank covering $A_1, A_2, \ldots, A_m$ constructed as in the proof of 7.3 yields an odd-set covering as follows. (Note: References to "nodes" in the following are to the "original" nodes of G, not nodes created by subdivision in the course of the problem reduction.) Take as single nodes in the odd-set covering all nodes incident to arcs in $L \subseteq A_1$, plus any single one of the nodes to which the arcs in $U \subseteq A_1$ are incident, if U is nonempty. Take as an odd set the remaining nodes to which arcs in U are incident. Take as an odd-set the nodes to which arcs in each of the blossoms $A_2, \ldots, A_m$ are incident. The resulting odd-set covering has (odd-set) rank equal to one half the (parity) rank of the parity covering. (Recall that a maximum cardinality independent parity set has twice as many elements, i.e. arcs of the subdivision graph, as G has arcs in a maximum cardinality matching.)

Let $M_1 = (E, \mathcal{I}_1)$, $M_2 = (E, \mathcal{I}_2)$ be two matroids, and let $sp^i(\ )$, $r^i(\ )$ indicate span and rank in matroid i, i = 1, 2. We say that E_1, E_2 provide a covering of E if $E_1 \cup E_2 = E$, and that the (intersection) <u>rank</u> of the covering is $r^1(E_1) + r^2(E_2)$.

<u>Corollary 7.4</u> (Matroid Intersection Duality Theorem [2]-[4])

Given two matroids M_1, M_2 over E, the minimum rank of a covering of E is equal to the maximum cardinality of an intersection.

Proof:

Reduce the matroid intersection problem to a parity problem by the construction of Section 2 and apply Theorem 7.3. Because the resulting matroid is separable with respect to the parity conditions, there exists a minimal rank covering in the form of a single set A_1 (no blossoms exist). We may assume, without loss of generality, that the labelled element e_i at the root of each Hungarian tree is identified with an element $e_i \in E\text{-}sp^1(I_0)$. Let $L' \subseteq E$ and $U' \subseteq E$ be the elements of E identified with elements of L and U. Then $E_1 = sp^1(U')$ and $E_2 = sp^2(L')$ is the desired intersection covering with the appropriate rank. (The rank of the parity covering being twice the rank of the intersection covering.)

8. COMPLEXITY OF THE PARITY ALGORITHM

We seek to establish an upper bound on the number of computational steps required by the algorithm, as a function of n, the number of element pairs (i.e. $|E| = 2n$). We assume that the number of computational steps required for independence testing in M is c(n).

We assert that, for the sets that must be tested in each of the matroids M_i formed by shrinking, the number of steps required for independence testing is $O(nc(n))$. (In nearly all cases independence testing is only $O(c(n))$.)

First, we notice that at most n augmentations are possible, if the computation is begun with the empty set in Step 0. We therefore proceed to investigate the amount of computation required for a single augmentation.

For a single augmentation, Steps 1.0 - 1.3 are each performed $O(n)$ times, Step 2 once, and Step 3 $O(n)$ times.

Step 1.0 requires the computation of sp(I) (the first time the step is performed) and of $sp(I + e_i)$. Suppose that spans are computed by successive tests for independence, as described in Section 1, and that the subroutine for independence testing is $n\ c(n)$ in length. Then each execution of Step 1.0 is $O(n^2c(n))$ in length.

Step 1.1 is, at worst, $O(n)$ in length.

Step 1.2 requires the testing of $I + e_i$ for independence, which is $n\ c(n)$ in length, and possibly identifying a circuit in $I + e_i$. If the circuit computation is carried out as suggested in Section 1, it is $O(n^2 c(n))$ in length, and this establishes an overall bound for the step.

Step 1.3 requires the computation of $sp(I - \bar{e}_i)$, which establishes a bound of $O(n^2 c(n))$ for the step.

Step 2 is $O(n)$ in length.

Step 3.0 requires backtracking which is $O(n)$ and various operations on two sequences which may also be assumed to be $O(n)$.

In Steps 3.1.0 - 3.1.2, each circuit C_j requires a computation which is $O(n^2 c(n))$ in length, and each such circuit C_j is computed at most once per augmentation. Thus, we may assume that the computation attributable to Steps 3.1.0 - 3.1.2 is $O(n^3\ c(n))$ per augmentation.

Steps 3.2 - 3.5 require scanning of labels of elements in the blossom, regardless of whether or not they have been scanned before. However, at worst, this implies only a single rescanning of any given label. Thus, at worst, we need only repeat the amount of computation performed in Step 1.

(Admittedly, there is a redundant computation of circuits in the algorithm. The circuit $C_j \subseteq I' + e_j$ is computed once when e_j is found to belong to a blossom and possibly two other times as part of scanning. operations. Possibly this redundancy could be eliminated, with some additional complication of the algorithm, i.e. additional housekeeping operations. However, this would at best improve the computational bound only by a linear scale factor.)

When we multiply the bound for each step by the number of times it is performed per augmentation, we obtain for Step 1.0, $O(n^3 c(n))$, Step 1.1 $(O(n^2)$, Step 1.2, $O(n^3 c(n))$, Step 1.3, $O(n^3 c(n))$, Step 2, $O(n)$, and Step 3, $O(n^3 c(n))$. The maximum of these bounds is $O(n^3 c(n))$, and this provides a bound on the length of computation per augmentation. Since there are $O(n)$ augmentations a bound on the total computation is $O(n^4 c(n))$, a factor of n greater than that required for matroid intersection [4].

Thus, if $c(n)$, the amount of computation required for independence testing, is a polynomial function of n, then the overall amount of computation required by the algorithm is itself a polynomial function of n.

9. ACKNOWLEDGEMENT

The author wishes to acknowledge the careful reading of the manuscript by Arnon Rosenthal, and his help in correcting errors.

REFERENCES

[1] J. Edmonds, "Paths, Trees, and Flowers," Can. J. Math.17, 449-467 (1965).

[2] J. Edmonds, "Submodular Functions, Matroids, and Certain Polyhedra," Combinatorial Structures and their Applications, Proceedings of the Calgary International Conference, Gordon and Breach, N. Y. (1970).

[3] E. L. Lawler, "Optimal Matroid Intersections," Combinatorial Structures and Their Applications, Proceedings of the Calgary International Conference, Gordon and Breach, N. Y. (1970), Abstract only, p. 233.

[4] E. L. Lawler, "Matroid Intersection Algorithms," submitted for publication.

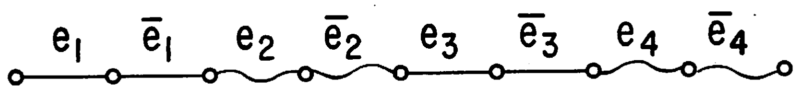

(a)

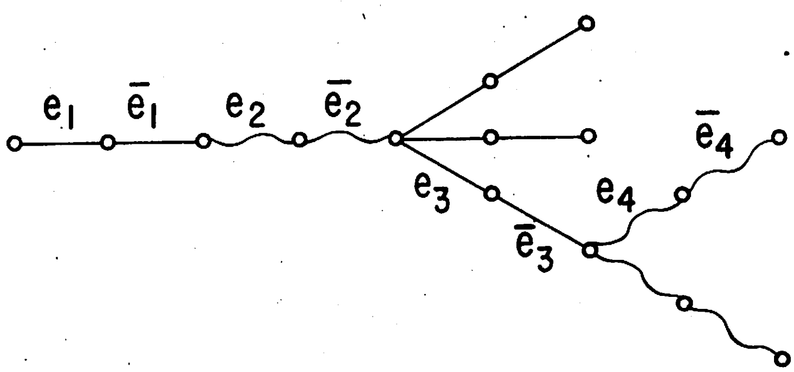

(b)

Fig. 1. Typical Augmenting Sequence and its Representation

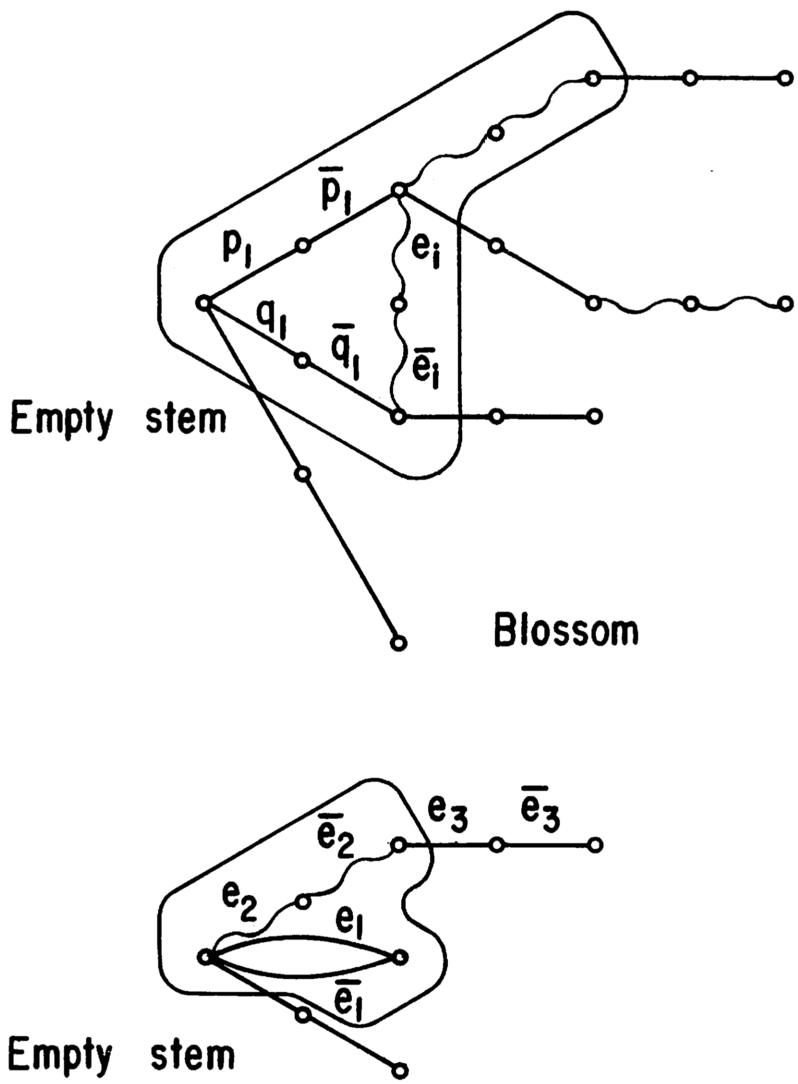

Fig. 2. Rooted Blossoms

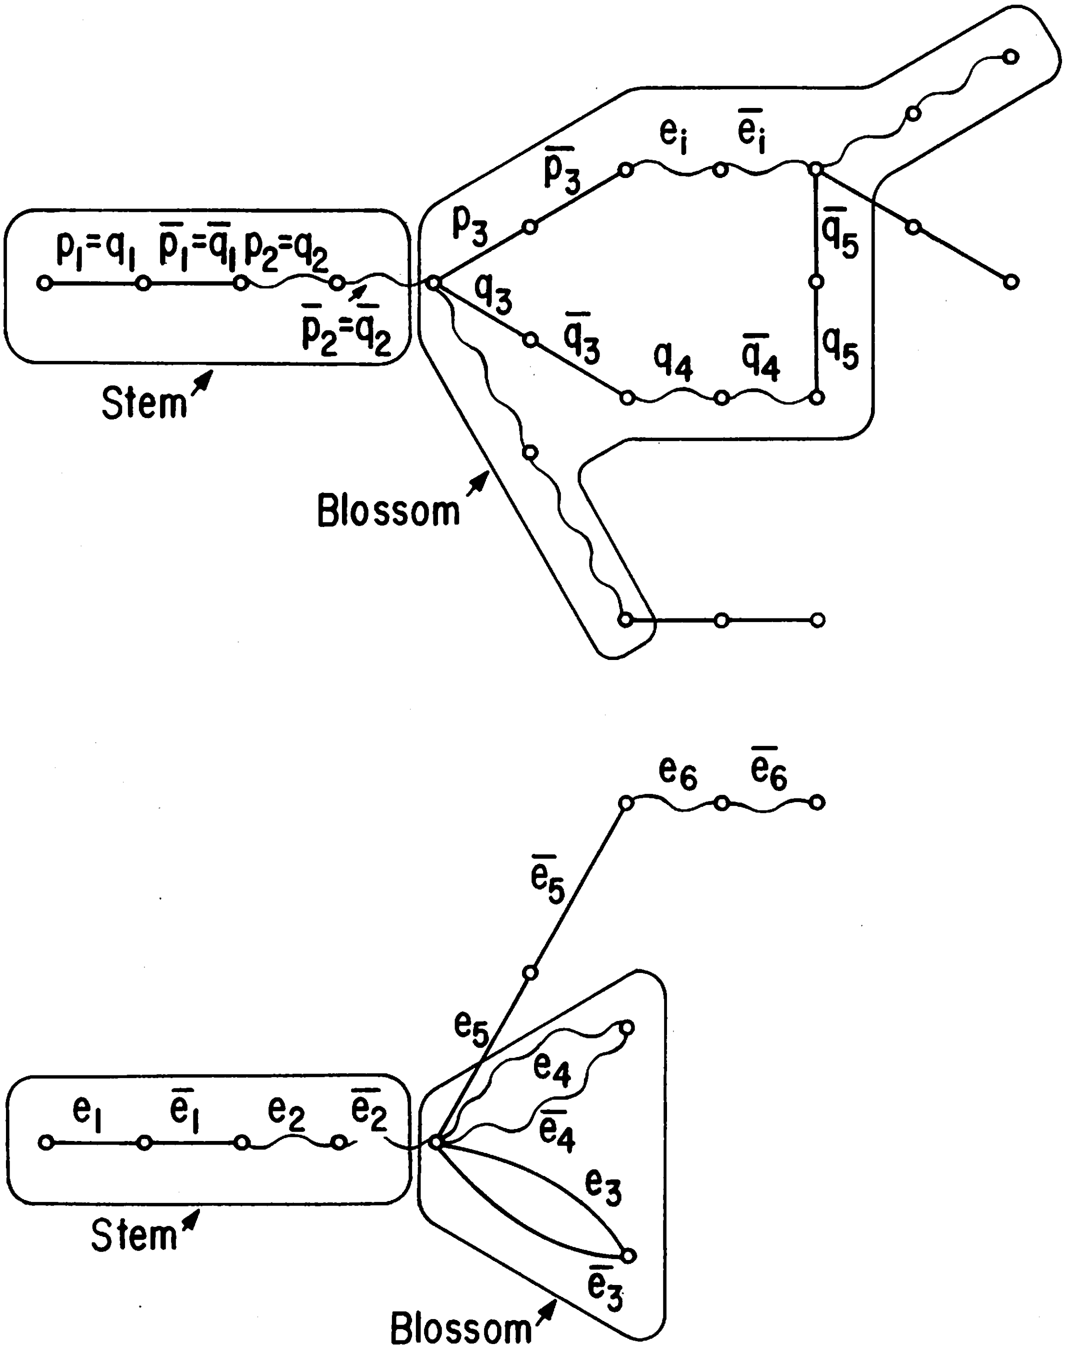

Fig. 3. Short Blossoms

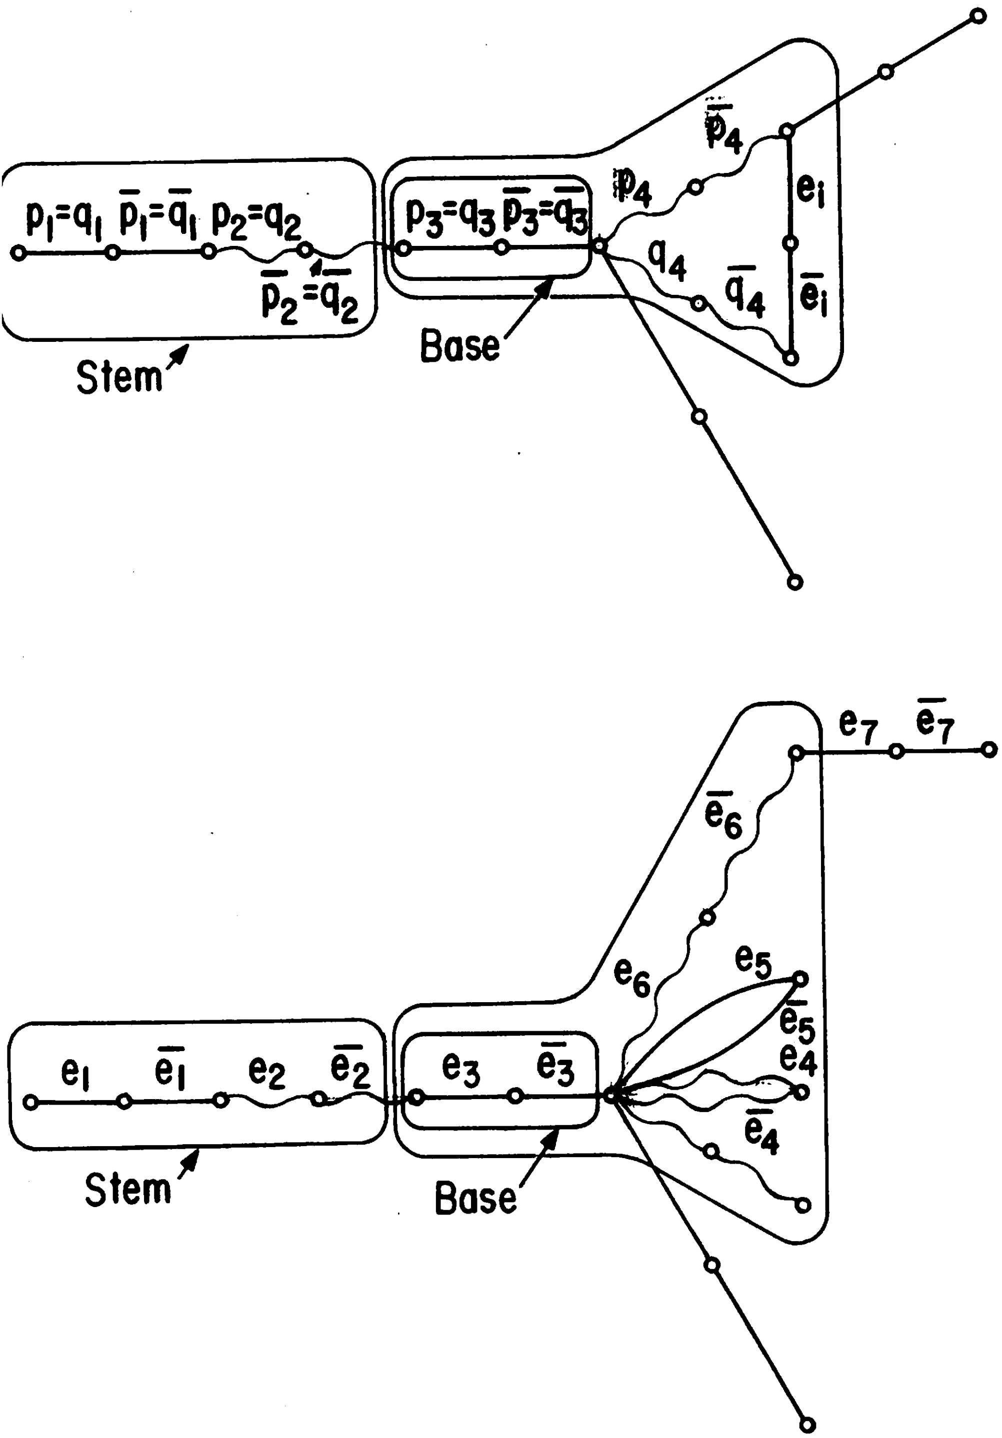

Fig. 4. Tall Blossoms

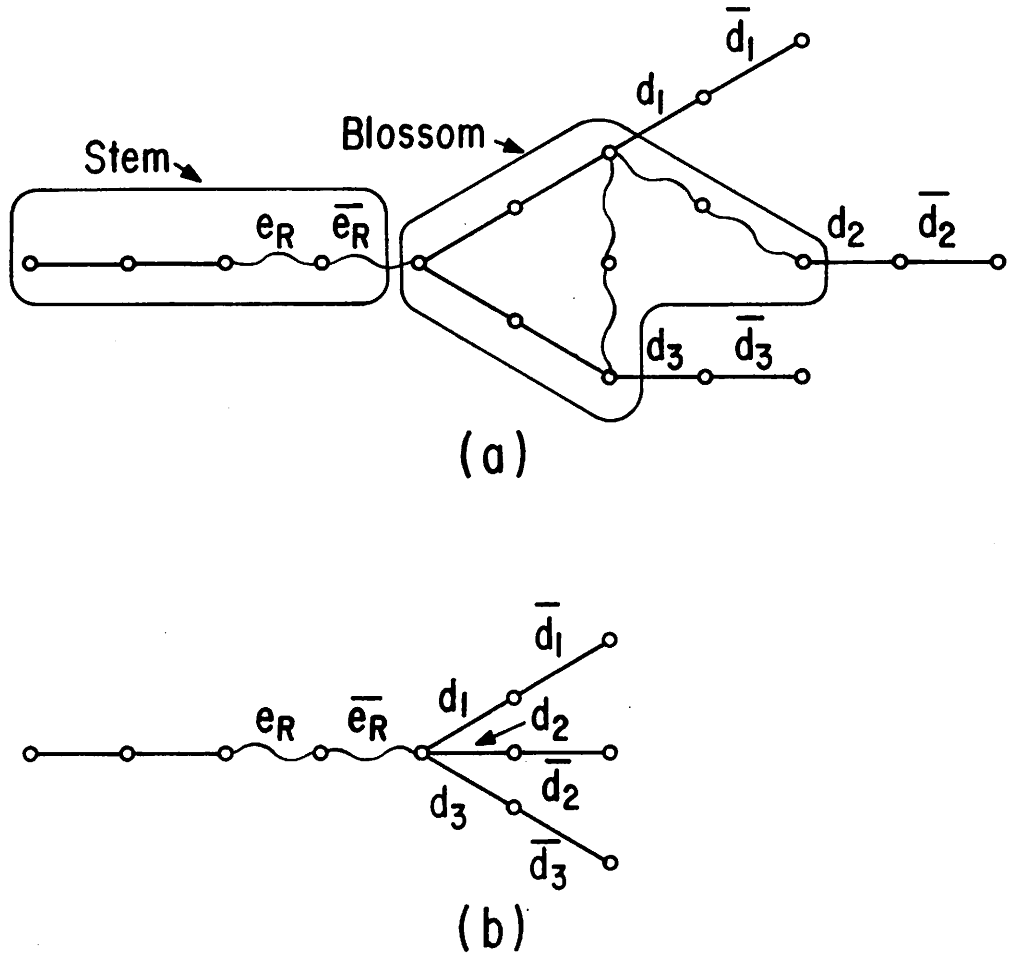

Fig. 5. Shrinking of Blossom

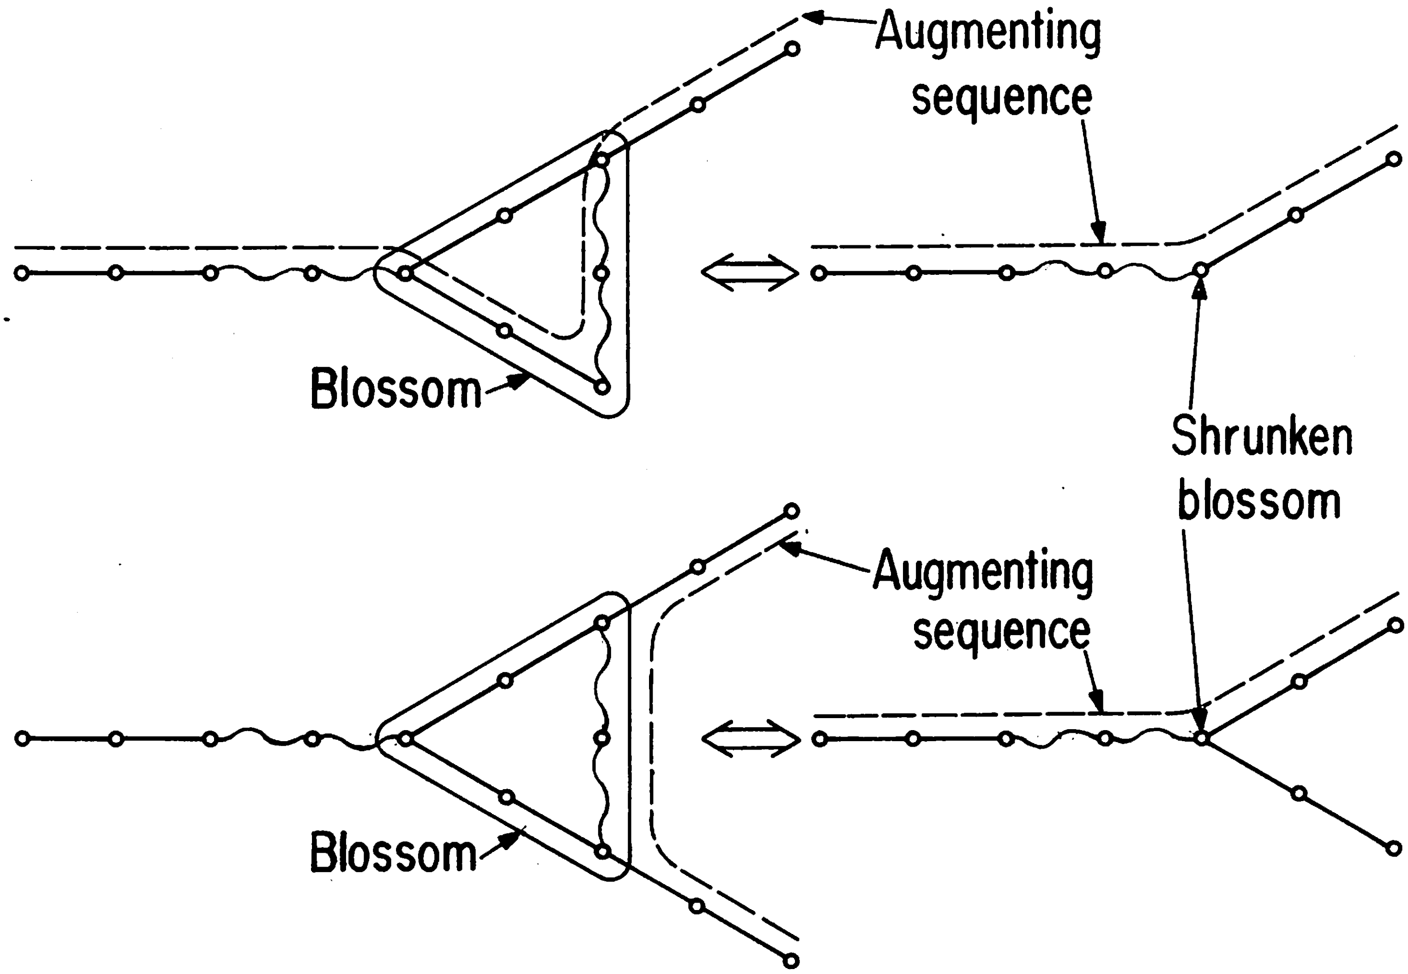

Fig. 6. Effect of Shrinking on Augmenting Sequences