A COMBINATORIAL BASE FOR SOME OPTIMAL MATROID INTERSECTION ALGORITHMS

by

Stein Krogdahl

STAN-CS-74-468
NOVEMBER 1974

COMPUTER SCIENCE DEPARTMENT
School of Humanities and Sciences
STANFORD UNIVERSITY

A Combinatorial Base for Some Optimal Matroid Intersection Algorithms

Stein Krogdahl

Abstract

E. Lawler has given an algorithm for finding maximum weight intersections for a pair of matroids, using linear programming concepts and constructions to prove its correctness. In this paper another theoretical base for this algorithm is given which depends only on the basic properties of matroids, and which involves no linear programming concepts.

Keywords: intersection algorithms, optimal intersections, weighted matroids.

This work was supported in part by The Norwegian Research Council for Science and the Humanities and by the Office of Naval Research contract NR 044-402. Reproduction in whole or in part is permitted for any purpose of the United States Government.

A Combinatorial Base for Some Optimal Matroid Intersection Algorithms

Stein Krogdahl

1. Introduction

The algorithms for which a theoretical base is given in this paper, have been known for some time, and were first developed by E. Lawler. However, the proofs given for the correctness of these algorithms have used linear programming concepts such as primal and dual solutions, and have been rather difficult to understand. Hopefully, the proofs given here will be easier to understand, and thereby will give deeper insight into the nature of the problems involved.

2. Some Properties of Matroids

In this section we shall develop the properties of matroids on which the following theory is built. We denote a matroid $M(E)$, and thereby mean a matroidian structure defined on the finite set E . It is assumed that the reader knows the basic properties of matroids, and we will use the following notation: If $A \subseteq E$ and $e \in E$, then $A-e$ and $A+e$ shall mean $A-\{e\}$ and $A \cup \{e\}$ respectively. The closure of a set $A \subseteq E$, or the span of A , will be denoted $sp(A)$. If $I \subseteq E$ is independent and $e \in sp(I)-I$, then the unique circuit in $I+e$ will be denoted $C(e,I)$.

Our first theorem is the following:

Theorem 1. Assume that $M(E)$ is a matroid and that I and J are subsets of E such that I is independent and $J \subseteq sp(I) -1$. Further assume that a one-to-one mapping $d: J \to I$ is defined such that for all $e \in J$, $d(e) \in C(e,I)$, and for all nonempty sets $A \subseteq J$ there is an $e \in A$ such that for all $e' \in A - e$ we have $d(e) \notin C(e',I)$. Then we can conclude:

A: the set $I' = I \cup J - d(J)$ is independent;

B: $sp(I') = sp(I)$, thus $d(J) \subseteq sp(I') -1'$;

c: for all $e \in J$ we have $e \in C(d(e),I')$.

Proof. For later convenience let us first choose one element e in each A c_ J such that for all $e' \in A-e$ we have $d(e) \notin C(e',I)$, and. call it $S(A)$.

Part B of the conclusion is a simple consequence of part A, since III $= |I'|$ and everything goes on within $sp(I)$. To prove part A we assume that I' contains a circuit C_0 . Because I is independent, $A_0 = C_0 \cap J$ is not empty, and we set $e = S(A_0)$. Because $d(e) \in C(e,I)-C_0$ and $e \in C(e,I) \cap C_0$, we can find a circuit C_1 within $C(e,I) \cup C_0-e$ such that $d(e) \in C_1$, and we know $A_1 = C_1 \cap J \subseteq A_0 - e$. If $A_1 \neq \emptyset$ we'choose $e_1 \in A_1$. Now we have $d(e) \in C_1-C(e_1,I)$ and $e_1 \in C_1 \cap C(e_1,I)$ and we can find a circuit $C_2 \subseteq C_1 \cup C(e_1,I)-e_1$ such that $d(e) \in C_2$, and we have $A_2 = C_2 \cap J \subseteq A_1-e_1$. If A_2 is not empty, we pick an $e_2 \in A_2$ and repeat this process again, and for some k , A_k must become empty, since $|A_i| < |A_{i\ 1}|$. But then C_k (which contains at least $d(e)$) must be a circuit entirely within I , and this is a contradiction outruling the existence of C_0 .

The proof of part C is done by a similar construction as the one used to prove part A, but a little more care is needed. First we order the elements of J in a sequence $J = \{j_1, j_2, \ldots, j_n\}$, by the following definition: $j_1 = S(J)$, $j_2 = S(J-j_1)$, . . . , $j_{n-1} = S(J-j_1-\ldots-j_{n-2})$, $j_n = S(\{j_n\})$. We then observe that if $p < q$, then $d(j_p) \notin C(j_q, I)$.

Now assume $e \in J$, and we will show $e \in C(d(e), I')$. Suppose $e = j_r$, then we know $B_0 = d(J) \cap C(e,I) \subseteq \{d(e), d(j_{r+1}), \ldots, d(j_n)\}$. If now $B_0 = \{d(e)\}$ we must have $C(d(e),I') = C(e,I)$, and we are done. If not, let p_1 be the smallest q such that $d(j_q) \in B_0 - d(e)$. Now $d(j_{p_1}) \in C(e,I) \cap C(j_{p_1}, I)$ and $e \in C(e,I) - C(j_{p_1}, I)$. Thus we can find a circuit $C_1 \subseteq C(e,I) \cup C(j_{p_1}, I) - d(j_{p_1})$ such that $e \in C_1$, We now know that $B_1 = d(J) \cap C_1 \subseteq \{d(e), d(j_{p_1+1}), \ldots, d(j_n)\}$. Now if $B_1 - d(e) \neq \emptyset$, we choose p_2 as the smallest q such that $d(j_q) \in B_1 - d(e)$, and let C_2 be a circuit in $C_1 \cup C(j_{p2}, I) - d(j_{p2})$ containing e. Define $B_2 = d(J) \cap C_2$, and if again $B_2 - d(e) \neq \emptyset$ we choose p_3 as minimum such that $d(j_{p_3}) \in B_2 - d(e)$ and go on in this way. We must then get $p_i < p_{i+1}$, and for some k, $B_k - d(e)$ must become empty. Then C_k is a circuit, containing e, whose only element outside I' can be d(e). Since I' is independent we must have $d(e) \in C_k$, and also $C(d(e), I') = C_k$. Thus $e \in C_k = C(d(e), I')$ as claimed.

Note that we also can add, to I' in the theorem, at least one element not in sp(I), and remove as many elements as we wish, and still rely on the independence of I'.

Theorem 2. Assume that I and J are independent sets in a matroid such that $J \subseteq sp(I)$. Then for each $e \in J-I$, $C(e,I) \cap (I-J) \neq \emptyset$, and for each $J' \subseteq J-I$ the set $I' = (I-J) \cap \left(\bigcup_{e \in J'} C(e,I) \right)$ is such that $|J'| \leq |I'|$.

Proof. The first statement must be true, or else J would contain a circuit. To prove the second we observe that $J' \cup (I \cap J) \subseteq sp(I' \cup (I \cap J))$. Since both these sets are independent we must have $|J' \cup (I \cap J)| \leq |I' \cup (I \cap J)|$, and thus $|J'| \leq |I'|$.

3. The Simple Border Graph (SBG) of an Independent Set

Assume that I is an independent set in a matroid M(E) . Then we can construct a bipartite graph, called the "simple bordergraph" (SBG) of I , in the following way:

The nodes of the graph are (in one-to-one correspondence with) the elements of E , and there is an arc between the nodes $e_1 \in E-I$ and $e_2 \in I$ if and only if $e_1 \in sp(I)$ and $e_2 \in C(e,I)$. This means that if e is not a self-circuit-element and $e \in E-I$, then e has no arcs onto it if and only if $e \in E-sp(I)$. Also if $e_1 \in E-I$ and $e_2 \in I$ and there is an arc between e_1 and e_2, then $I+ e_1 - e_2$ is independent.

We now note that the function d used in Theorem 1 corresponds to a matching in the SBG of I , and that part A of this theorem says something about when the interchanges indicated by the arcs in a certain matching can be performed simultaneously without destroying the independence of I .

To be able to formulate part A of Theorem 1 in these graphic terms we define the graph "induced" by a matching in the SBG of I , as the graph with node set equal to the set of end-nodes of the matching, and with arc set equal to the set of all arcs between these nodes in the original SBG. Indeed this graph contains the arcs of the matching itself, and we call these the "main arcs" of the induced graph. Further a matching is said to be "usable" if the interchange in I of its end-nodes inside I with those outside I makes a new independent set.

Part A of Theorem 1 then says that a matching D in the SBG of I is usable if the induced graph of any submatching D' of D has at least one main arc which is the only arc to its end node in I .

To get this condition on a, for us, more convenient form, we define a "main cycle" in the induced graph of a matching as a simple cycle that uses a main arc exactly each second time. We will say that a matching "induces a main cycle" if its induced graph contains a main cycle, and if a matching does not induce any main cycle it is said to be "clean".

Theorem 3. A clean matching D in the SBG of an independent set I is usable.

Proof. We will show that if D does not have the property that every submatching D' of D has at least one arc whose node inside I has degree 1 in the induced graph of D' , then the induced graph of D must contain a main cycle. Therefore assume that D' is a submatching of D inducing a graph where all the nodes inside I has at least two arcs onto it. Then start at any node of D' outside I , pass along the main arc to its endpoint inside I , and take any of the other arcs from here. Then we are back at the outside end of another main arc, and we repeat the process. This process must eventually lead back to a main arc which is used before, and then a main cycle in the induced graph of D' is found, and this is also a main cycle in the induced graph of D .

We conclude this section by giving a theorem that assures the existence of matchings under certain conditions.

Theorem 4. Let I and J be independent sets of a matroid such that $J \subseteq sp(I)$. Then there is a matching in the SBG of I such that J-I is exactly the set of end-nodes of the matching outside I , and all the inside end-nodes are within I-J .

Proof. By Theorem 2 each node in J-I must have at least one arc leading into I-J , and these arcs must be so arranged that for every $J' \subseteq J-I$ the set in I-J directly reachable through an arc from nodes in J' has greater cardinality than J' . Thus, by a well known theorem about matchings in bipartite graphs the matching required by our theorem exists.

4. Intersections of Matroids

In the following we shall deal with two matroids here called M_R and M_B (the red and the blue matroid), both defined on the same set E . A subset of E which is independent in both M_R and MB is traditionally called an "intersection" of M_R and M_B , and our task shall be to develop algorithms for finding optimal (in a certain sense) intersections. In a later chapter weights are given to the elements of E and the task is to find an intersection with the maximum sum of weights. However, we shall first treat the simpler case where all the weights are one, that is, to find a maximum cardinality intersection.

For simplicity we assume that neither M_R nor MB has self-circuit-elements. If any of them has, these elements (which cannot figure in any intersections) can be deleted first, or they can simply be ignored by any algorithm.

5. The Bordergraph (BG) of Intersections, and Alternating Paths

If I is an intersection of M_R and M_B we define the "bordergraph" (BG) of I to be, in a certain sense, the union of the SBG of I in M_R and in M_B. That is, the set of nodes of the BG of I is again in one-to-one correspondence with E , and the arcs are exactly those from the SBG of I in M_R colored red, and exactly those from the SBG of I in M_B colored blue, and only these.

If a node outside I has no blue arcs onto it, it is said to be "unicolored", with color red, and if it has no red arcs onto it, it is unicolored with color blue. If it has neither red nor blue arcs onto it, it is also said to be unicolored, now with color white. Obviously (if the matroids contain no self-circuit elements) an element is outside $sp_B(I)$ if it is unicolored with color red or white, and outside $sp_R(I)$ if it is unicolored with color blue or white.

We now define an "alternating path" in the BG of an intersection I as follows: Either it is a single unicolored white node outside I , or it is a simple path or cycle of length at least one which uses red and blue arcs alternately and which is such that any end-node of the path outside I is unicolored. ("Simple" is here used in the sense that no node is "used twice" along the path.)

If P is an alternating path in the BG of I we denote its set of nodes outside I as " out(P) ", the set of those inside I as " in(P) ", and the set I −in(P) ∪ out(P) as " P(1) ". Further we say that P is "usable" if P(I) is an intersection of the two matroids M_R and M_B.

For later use we will classify the alternating paths in four groups: W-paths, N-paths, M-paths and O-paths. An O-path is a cyclic path, a W-path is one with both (unicolored) endpoints outside I (with the single white-node-path as a special case), an N-path is a path with one (unicolored) endpoint outside I and one inside I, and an M-path is one with both end-points inside I.

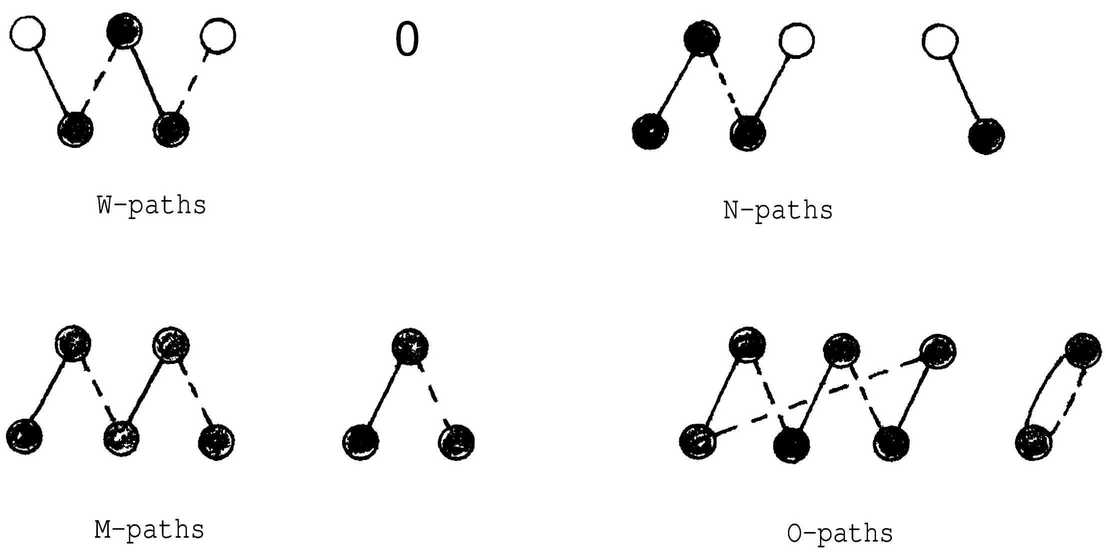

Examples of alternating paths. The unfilled nodes must be unicolored, and the lower nodes are assumed to be inside I .

By considering only the red or only the blue arcs of an alternating path, a red and a blue matching (of which one or both may be empty) is defined. If both these matchings are clean in their own SBG, then the alternating path is also said to be clean.

By Theorem 3, and the comments after Theorem 1, we get the following theorem:

<u>Theorem 5</u>. A clean alternating path is usable.

Our main interest is in W-paths, because if P is a usable W-path, then P(I) is a new intersection with one element more than I . The following theorem assures-the existence of W-paths in the bordergraph of I , if greater intersections exist at all.

Theorem 6. Let I and J both be intersections such that $|I| < |J|$. Then there is an alternating path P of W-type in the BG of I such that $out(P) \subseteq J-I$ and $in(P) \subseteq I-J$.

Proof. If J contains an element outside $sp_R(I) \cup sp_B(I)$, then this is a unicolored white element and it is usable as a W-path alone, and we are done. Therefore suppose $J \subseteq sp_R(I) \cup sp_B(I)$ and define $J' = J-I$. Now partition J' into the sets J_R , J_B and J_0 as being the unicolored red elements, the unicolored blue elements and the rest of the elements of J' . By Theorem 4 we can now find a red matching using exactly the nodes of $J_R \cup J_0$ outside I and only nodes in I-J inside I , and a blue matching using exactly the nodes of $J_B \cup J_0$ outside I and only nodes in I-J inside I . Now define I_R , I_B and I_0 as the nodes in I-J which have only a red arc onto it, only a blue arc onto it and one red and one blue arc onto it respectively in this matching. We know that $|I_R \cup I_0 \cup I_B| \leq |I-J| < |J-I| = |J_R \cup J_0 \cup J_B|$, $|I_R \cup I_0| = |J_R \cup J_0|$ and $|I_B \cup I_0| = |J_B \cup J_0|$. Therefore we must have $|I_R| < |J_R|$ and $|I_B| < |J_B|$.

The arcs we have got now must obviously form a set of alternating paths of various types. However since every O-path will '*consume" nodes only from J_0 and I_0 and each N-path will consume exactly one node in J_R and one in I_R or one in J_B and one in I_B (plus possibly some in J_0 and I_0) at least one path must extend from a node in J_R to a node in J_B . This path is a W-path in the BG of I , and it obviously meets the requirements of the theorem.

6. Shortcutting of Alternating Paths

Suppose we go along an alternating path P from one end to the other, or around an O-path, and are just about to leave a node by an arc of color X . If we then, from where we are now, find another arc, also of color X , leading to a node further ahead on our path, we can delete all nodes and arcs lying between these two nodes on the path and insert this new arc instead, thus obtaining a new alternating path of the same type. This operation is called "shortcutting", and the resulting alternating path is called a shortcut of P . We obtain the following theorem.

Theorem 7. If P is any alternating path where no further shortcutting is possible, then P is clean, and thus usable.

Proof. It is easy to see that if either the red or the blue matching in the graph induce a red or blue main cycle, then at least one shortcutting edge must exist.

7. Maximum Cardinality Intersection Algorithm

We can now construct a rather straight-forward algorithm for finding a maximum cardinality intersection of two matroids.

If we have an intersection I_k with k elements, then Theorem 6 tells us that if its BG contains no W-paths, then I_k is a maximum cardinality intersection. If not, we can take any W-path in the BG, shortcut it until no shortcut is possible, and by Theorem 7 we know that it is now usable and will bring us to an intersection I_{k+1} with k+1 elements. Starting with $I_0 = \emptyset$ this will give us an algorithm for finding a maximum cardinality intersection.

We shall not elaborate on how such an algorithm can be implemented or optimized here, but only notice that if we have a way of determining in polynomial time if a set is independent in M_R or in M_B, then we can also find $C_R(e,I)$ and $C_B(e,I)$, and thus build the BG in polynomial time. The search for W-paths and a possible shortcutting process can obviously also be carried out in polynomial time, which altogether gives a maximum cardinality algorithm working in polynomial time.

8. Weighted Matroids

We will now consider the case where the elements of the set E over which M_R and M_B are defined has weights. That is, a mapping w from E to the real numbers is given.

We also define the weight of a set $A \subseteq E$ as the sum

$$w(A) = \sum_{e \in A} w(e) \quad .$$

Also we define the weight of an alternating path P in the BG of I as $w(P) = w(P(I)) - w(I)$. This is obviously equivalent to

$$w(P) = w(out(P)) - w(in(P)) \quad .$$

An intersection I is said to be k-maximal if $|I| = k$ and for all intersections I' such that $|I'| = k$ we have $w(I') \leq w(I)$. Our aim in the following is to show that if I is k-maximal and P is a weightiest W-path in the BG of I which cannot be shortcutted without lowering its weight, then P(I) is a k+1 -maximal intersection.

If I is k-maximal-then the weight of any usable alternating path of type 0 or N must be less than or equal to zero, or else a weightier intersection with k elements could have been found.
A consequence of this is the following:

Theorem 8. If I is a k-maximal intersection and P is any alternating path in the BG of I such that any shortcut of P will give a path with less weight, then P is clean, and thereby usable.

Proof. We will show that the assumption that one (or both) of the matchings given by the path P induces a main cycle, (that is, P is not clean) leads to a contradiction. Therefore assume that e.g. the red matching induces a main cycle. Now, if the alternating path is of O-type remove any blue arc to obtain a linear structure.

The idea of the proof now is really quite simple, namely that each time the main cycle contains a "shortcut-arc" in P like this

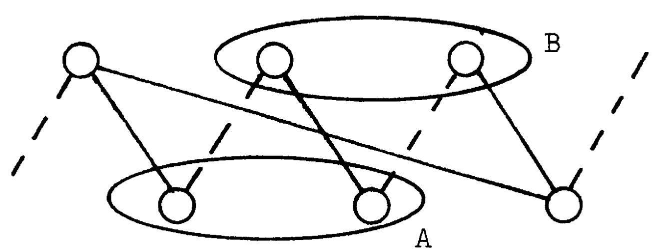

we know that w(A) <w(B) because any shortcut is supposed to give less weight. (We assume that the red arcs are fully drawn, and that the lower nodes are inside I.) On the other hand, if we have a "cross-over" structure like this:

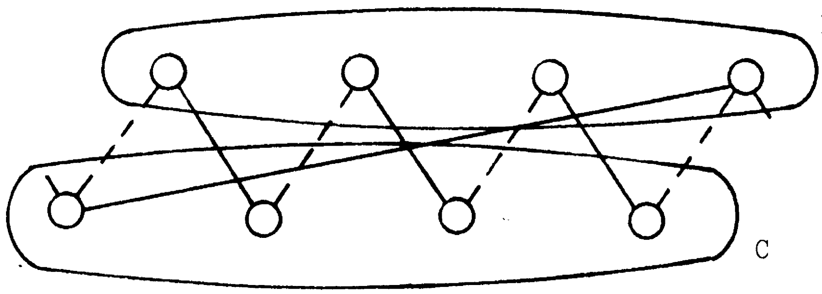

and if we know that neither the blue matching nor the red matching (including the crossover-arc) involved contains any main cycle, then the local O-path formed is usable and $w(C) \geq w(D)$ or else I would not be k-maximal.

We will afterwards show if there exists any main cycle at all then we can find one (in the same or in the other color) in which each crossover arc obeys the conditions above. First, however, we will show that this would lead to a contradiction.

Suppose that we have obtained such a main cycle in red, and that we, as above, draw the alternating path so that its red arcs go down to the right. Then every red crossover arc in the main cycle will go up to the right, and every red shortcut arc will go up to the left. The main cycle must obviously contain a leftmost and a rightmost main arc in this drawing, and we will use this to partition the main cycle as follows: Part 1 is what you pass if you start by going down the leftmost main arc and follow the main cycle until the top of the rightmost main arc is met. Part 2 is the rest. Part 1 must in a way be dominated by crossover arcs, although it may have many back-steps by shortcut arcs. Part 2 must likewise be dominated by shortcut arcs. For example, the two parts can look like this:

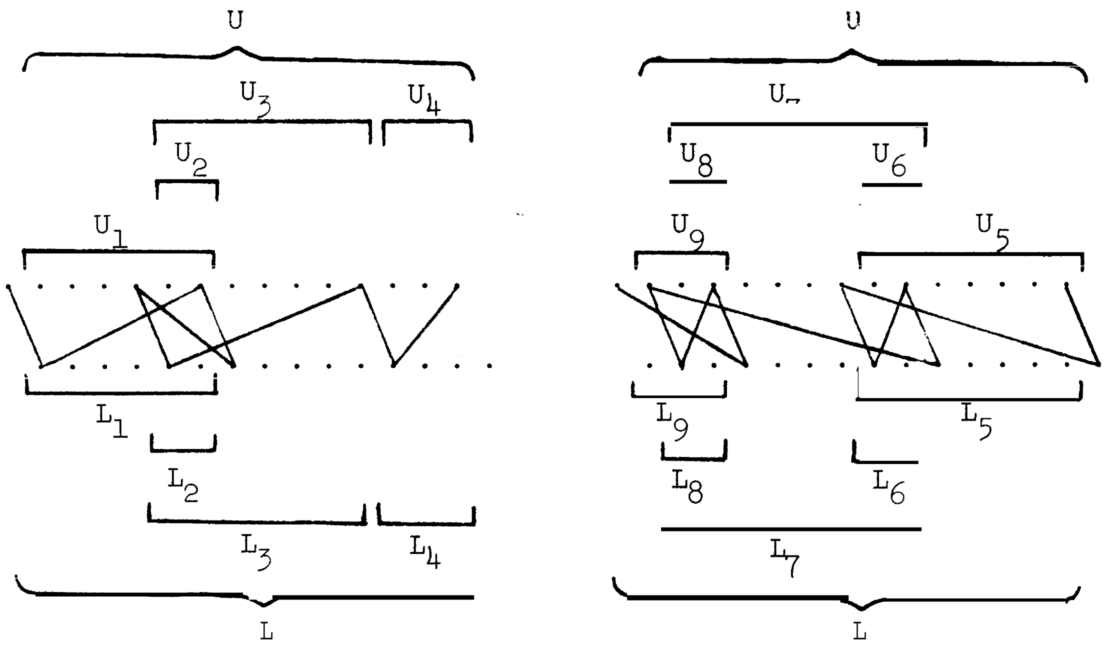

Part 1 Part 2

Note that the two parts cannot use the same main arcs, and there will generally be many red main arcs not in the main cycle in between those used by the cycle. Now let U be the set of all upper nodes on the part of the path covered by our main cycle, except the leftmost one, and let L be the lower ones except the rightmost node. We have $|L| = |U| > 0$. By summing along Part 1 we get $w(L) \geq w(U)$ but by summing over Part 2 we get $w(L) < w(u)$. For example in the above illustration, $w(L_1) \geq w(U_1)$, $w(L_2) < w(U_2)$, $w(L_3) \geq w(U_3)$, $w(L_4) \geq w(U_4)$, $w(L_5) < w(U_5)$, $w(L_6) \geq w(U_6)$, $w(L_7) < w(U_7)$, $w(L_8) \geq w(U_8)$, and $w(L_9) < w(U_9)$. This contradiction would now complete the proof, if we knew that the existence of an induced main cycle implied the existence of one in which each crossover arc formed a local clean (and thus usable) 0-path.

To see that this is correct, assume that there is an induced main cycle in one of the colors. If, inside the subpath that this main cycle covers, there are other main cycles in this or the other color, then choose one for which there is no other main cycle that covers a strict subpath of the subpath that this one covers.

Now we must look at each crossover arc in this main cycle.

Consider the following possible picture within a "great main cycle" which includes the crossover arc AB :

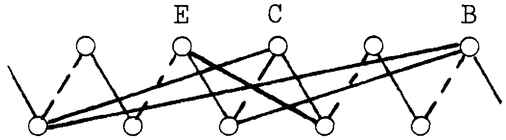

Here there will now be no main cycle induced by these blue (dotted) arcs, nor by the red main arcs crossed by the AB-arc. However, if we also consider AB as a main arc, as we do when we look at this as a local O-path, there can still be a main cycle as indicated above: BAC DEFB . Then, however, we can delete AB from our great main cycle and insert ACDEFB instead to form a new great main cycle. The new crossover arcs formed by this process (AC and FB) must be shorter than the original one, and therefore a repetition of this process must terminate. When this happens all crossover arcs must form clean O-paths, and the proof is complete.

9. O-paths, N-paths, and Cut N-paths

An immediate consequence of Theorem 8 is that the BG of a k-maximal intersection cannot contain any O-path or N-path with positive weight at all. For if one such positive path existed, we could go on performing such shortcuts on it that would not lower its weight until such shortcuts no longer were possible. (Note that such a process cannot make the path shorter than two nodes.) The resulting path could then only be even more positive, and according to Theorem 8 it would be usable. This, however, is impossible if the intersection is k-maximal.

For our next main theorem we also need a slightly different fact, namely that if we, in the BG of a k-maximal intersection, have a W-path and an M-path which are node-disjoint, then the sum of the weight of these two paths must be zero or less. To see this we can put them together to form a special N-path with one arc missing in the middle. We will call such a path a "cut" N-path. As a cut N-path we will also accept a W-path together with a single node inside I .

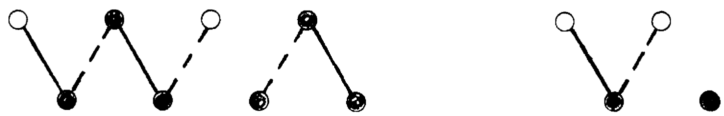

Cut N-paths. Unfilled nodes must be unicolored.

Now assume that the weight of such a cut N-path is positive and shortcut it exactly as we did above. If a shortcut crosses the cut, this leads to an immediate contradiction since we obtain a positive N-path. If not, we can show that both matchings involved are clean, by arguments similar to those used to prove Theorem 8. However, now a crossover edge may also form a local N-path (which is equally good), and we do not have to worry about the usability of the local N- or O-paths formed. Since both matchings are clean the cut N-path is obviously usable and cannot have positive weight.

We state these results as a theorem.

<u>Theorem 9</u>. If I is a k-maximal intersection then the BG of I contains no O-path, N-path or cut N-path with positive weight.

We now prove the -following theorem, which is the weighted counterpart of Theorem 6

Theorem 10. Suppose that I is a k-maximal intersection, and that J is any intersection with k+1 elements. Then there is a W-path P in the BG of I such that $w(P) \geq w(J) - w(I)$, and this P can be chosen so that $out(P) \subseteq J-I$ and $in(P) \subseteq I-J$.

Proof. The proof given for this theorem is very similar to that given for Theorem 6.

First look at the case that J has elements outside $sp_R(I) \cup sp_B(I)$, and let e be any of these elements. In the BG of I , e will now be a white unicolored node and can serve as a W-path P alone. Since J-e is another intersection with k elements, we must have $w(J-e) \leq w(I)$ which implies $w(P) \geq w(J) - w(I)$ as required.

Then assume that $J \subseteq sp_R(I) \cup sp_B(I)$ and partition J-I into parts J_R , J_B and J_0 , find a red and a blue matching, and define disjoint subsets I_R , I_B and I_0 in I-J exactly as we did in the proof of Theorem 6. The figures formed must now, as then, be a set of disjoint alternating paths of the four different types, and there must be at least one W-path.

Now there may be elements in I-J which are in neither I_R , I_B nor in I_0 , that is, they are not reached by any of the arcs in the matchings. Call the set of these I_1 . As an example look at the following picture:

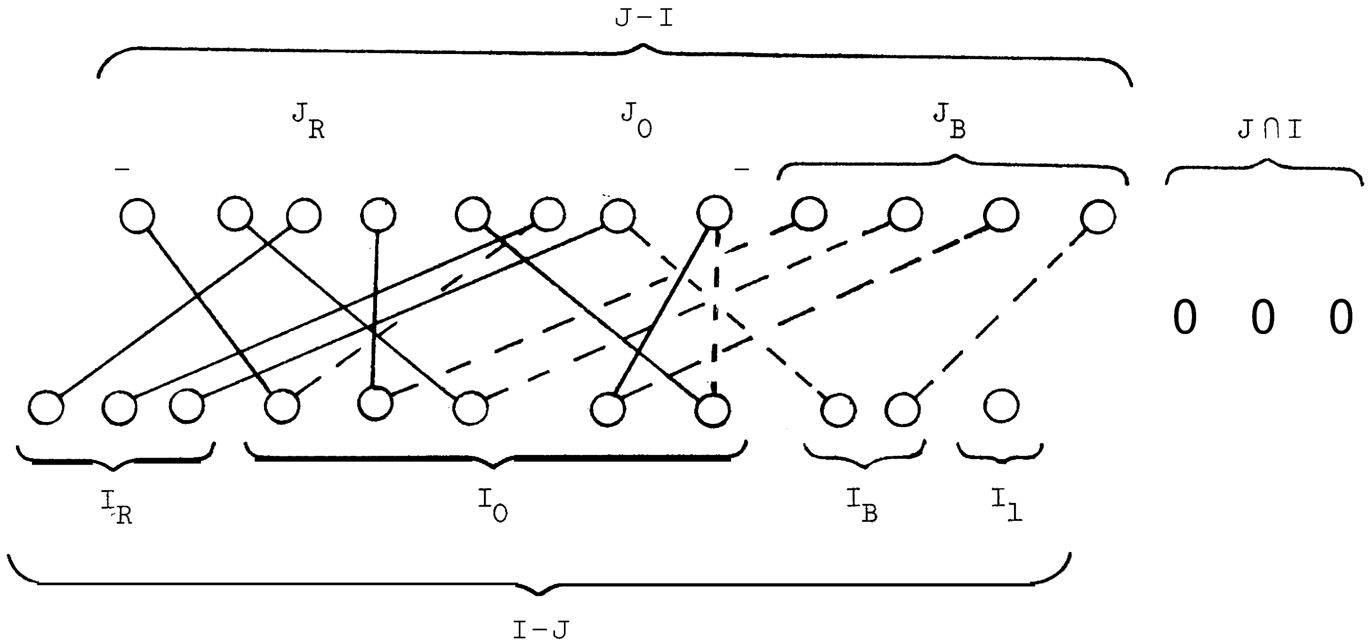

It follows now that the number of M-paths plus $|I_1|$ must always be exactly one less than the number of W-paths. An easy way to see this is first to observe that the following equation must hold:

$$|J_R| - |I_R| = |J_B| \quad |I_B| = 1 + |I_1| .$$

Then "remove" all N-paths. Since each N-path has either exactly one node in J_R and one in I_R, or one in J_B and one in I_B, the equations above must be kept true.

But now we must also have

$$|J_R| = |J_B| = \text{number of W-paths}$$

$$|I_R| = |I_B| = \text{number of M-paths.}$$

Also we have (from the first equation)

$$|J_R| - 1 = |I_R| + |I_1|$$

which gives exactly what we want.

Now it is easy to see that any of the W-paths present is good for our purpose. Choose one of them as P , and pair the rest of them to either an M-path or an element in $|I_1|$. This must fit exactly, and form a set of cut N-paths. The rest of the present paths must be either N-paths or O-paths. By Theorem 9 this implies that $w(J - I - out(P)) \leq w(I - J - in(P))$. This is equivalent to (since $out(P) \subseteq J-I$ and $in(P) \subseteq I-J$): $w(out(P)) - w(in(P)) \geq w(J-I) - w(I-J)$. This is again equivalent to what we want, namely:

$$w(P) \geq w(J) - w(I) .$$

10. Concavity

In this section we shall prove that the weight-increase we can obtain from a k-maximal to a k+l -maximal intersection cannot be greater than the increase obtained from a k-l - to a k-maximal one in the same pair of matroids. This property could suitably be called "concavity", and it will help us to determine when a maximal weight intersection is found. From what is proved until now we can easily construct an algorithm giving us a k+l -maximal set if we have a k-maximal one, and if intersections with k+l elements at all exist. However, not even if all the weights are positive, will the weight of a k-maximal set always increase with k , and the concavity will guarantee that we have obtained a maximal weight intersection the first time we" cannot get a weightier intersection by taking in one more element.

We will prove two theorems, whose combination immediately will give us the concavity. The first is a stronger version of Theorem 9, namely that even if a W-path and an M-path are not disjoint the sum of their weights cannot be positive in a k-maximal intersection. The other is that if P is a clean W-path in the BG of I , then P will appear as an M-path in the BG of P(1) , with weight -w(P) .

For the proof of the stronger version of Theorem 9, we introduce the concept of an "alternating walk" as being exactly the same as an alternating path, except that it may use the same nodes and arcs more than once on its way.

The weight of an alternating walk is defined so that the weight of a node is counted as many times as the node is used by the walk. The walks are classified as W-walks, N-walks, M-walks and O-walks exactly as for paths. We can then prove the following theorem.

Theorem 11. Assume that I is a k-maximal intersection. Then there is no O-walk or N-walk in the BG of I with positive weight.

Proof. We will show this by induction on the length of the walk, expressed as the number of nodes used, in the sense that each node is counted once each time it is used.

Any N-walk or O-walk has at least length 2 , and if the length is 2 then it is obviously also an N- or O-path. Therefore the theorem holds in this case.

Now assume that the theorem holds for all lengths less than L , and that we have a 0- or N-walk of length L . If none of the nodes of the walk is used twice (or more) we have a 0- or N-path, and the theorem must hold. If not, start in one end of the N-walk, or anywhere on the O-walk, and pass on to the first meeting with a multiple-used node.

Here skip over to a later passage of this node, and follow the walk in the old direction until we reach the end of the N-path, or are back at the starting point of the O-path. We have then passed through an N- or O-walk with length less than L , so its weight is not positive. The part of the walk we skipped must have been an O-walk with length less than L , forcing its weight not to be positive. However, the weight of the original path is the sum of the weight of the two parts, which gives our theorem.

Now the theorem we need follows quite easily.

<u>Theorem 12</u>. Assume that I is a k-maximal intersection, that P is a W-path and that Q is an M-path (or possibly only a single node in I), in the BG of I. Then $w(P) + w(Q) \le 0$.

<u>Proof</u>. If the two paths are disjoint, the theorem follows immediately from Theorem 9. Now assume that they are not disjoint, and assume first that Q is a single node. Then the theorem follows from summing the weight of the two N-paths starting in each end of the W-path, and ending at the M-path-node.

If the M-path is a real one, we do a generalization of this. We choose any node which is used both by the W- and the M-path, and obtain two N-walks, R_1 and R_2 , by starting in one and the other end of the W-path, and shifting over to the M-path at the chosen node. Note that the direction in which you shall proceed in the M-path after the shift is determined by the direction in which you come to the shift-node in the W-path. Thus these two N-walks cover exactly what the M- and the W-path covered and $w(R_1) + w(R_2) = w(P) + w(Q)$. But by Theorem 11 $w(R_1) \le 0$ and $w(R_2) \le 0$ so the theorem follows.

Theorem 13. If P is a clean W-path in the BG of an intersection I then this W-path will appear as an M-path in the BG of P(I) .

Proof. This theorem is a direct consequence of part C of Theorem 1, which simply says that the arcs used in P will turn up also in the BG of P(1) , since both the matchings in P are clean. This is exactly what we need.

Theorem 14. If I_{k-1} , I_k and I_{k+1} are k-1-, k- and k+1 -maximal intersections respectively, then $w(I_k) - w(I_{k-1}) \geq w(I_{k+1}) - w(I_k)$.

Proof. By the earlier results we know that we can find a clean W-path P_1 in the BG of I_{k-1} such that if $I'_k = P_1(I_{k-1})$, then $w(I'_k) = w(I_k)$. Further we can find a clean W-path P_2 in the BG of I'_k such that if $I'_{k+1} = P_2(I'_k)$ then $w(I'_{k+1}) = w(I_{k+1})$. We now want to prove that $w(P_1) \geq w(P_2)$. This follows from the fact that in the BG of I'_k , P_1 appears as an M-path with weight $-w(P_1)$, "by Theorem 13. Therefore since I'_k is k-maximal we know by Theorem 12 that $-w(P_1) + w(P_2) \leq 0$, which is exactly what we want.

11. Maximum Weight Intersection Algorithm

We can now construct an algorithm for finding a maximum weight intersection of two matroids, which is very similar to the maximum cardinality algorithm given earlier. The only change we have to make is that we now each time must find a weightiest W-path in the BG of the intersection we have. If the weight of this path is negative, or if no W-path exists at all, we now have a maximum weight intersection by Theorems 10 and 14. If this is not the case, we perform weight-preserving shortcuts on this W-path as far as possible, and know then by Theorem 8 that the resulting W-path is usable. Thus, by Theorem 10, the performance of this path will bring us to a k+1 -maximal intersection.

We will not here go into further details on this algorithm, only notice that under the same conditions as for the maximum cardinality algorithm, we can make this algorithm work in polynomial time.

It may be interesting to notice that if all subsets of E are independent in one of the matroids, then the algorithm above will degenerate to the well known greedy algorithm for the other matroid.

12. A Characterization of Optimal Intersections

We conclude by giving necessary and sufficient conditions for an intersection to be k-maximal and to be of maximum weight.

Theorem 15. An intersection is k-maximal if and only if $|I| = k$ and its BG contains no 0-, N- or cut N-paths with positive weight. An intersection has maximum weight if and only if its BG has no M-, N-, 0-, or W-paths with positive weight.

Proof. By earlier results the above conditions are obviously all necessary. To get the sufficiency in the first part, assume that I is an intersection with k elements which is not k-maximal. We will show that there must exist a positive 0-, N- or cut N-path in its BG. Since I is not k-maximal there is an intersection J so that $w(J) > w(I)$ and $|J| = |I|$. Then we make a construction similar to the one used in the proof of

Theorems 6 and 10, and can easily verify that every W-path formed can be exactly paired to one M-path (or one element in J-I not met by any arc in the matchings). We have then obtained a set of disjoint O-, N- and cut N-paths whose out-parts and in-parts exactly form J-I and I-J respectively. Since $w(J-I) > w(I-J)$ at least one of these paths must have positive weight.

For the sufficiency of the second part we first observe that if the BG of I has no positive M- or W-path, then it cannot have any positive cut N-path. Thus I is k-maximal, with $|I| = k$.

By Theorem 10 and the concavity we know that there cannot be any intersection I' such that $|I'| > |I|$ and $w(I') > w(I)$. However, by again using the same technique as in the proof of Theorems 6 and 10 we obtain that if $k \geq 1$ and I is a k-maximal intersection then we can find an M-path P in the BG of I such that P(I) is a k-1 -maximal intersection. Thus, by the concavity again, we know that there cannot be any intersection I" such that $|I''| < |I|$, and $w(I'') > w(I)$. Hence the theorem is proved.

Reference

E. L. Lawler, "Optimal matroid intersections," in Combinatorial Structures and Their Applications, Eds. R. Guy, H. Hanani, N. Sauer, J. Schonheim, (Gordon and Breach, 1970), p. 233.