Chapter 13

Max-Flow Min-Cut Matroids

13.1 Overview

Concepts, theorems, or algorithms in one area of mathematics often inspire new approaches in another area. In turn, the ensuing new developments in the latter area may lead to new ideas in the former one. In this chapter, we describe an interesting instance of this cyclic process.

We start with the *max flow problem* for undirected graphs, which may be defined as follows. Given is a connected and undirected graph G. One of the edges of G, say l, is declared to be *special*. To each edge e of G other than l, a nonnegative integer h_e is assigned and called the *capacity* of e. Define G to have *flow value* F if there are F cycles, not necessarily distinct, that satisfy the following two conditions. Each cycle of the collection must contain the special edge l, and any other edge e of G is allowed to occur altogether in at most h_e of the cycles. The max flow problem demands that one solve the problem $\max F$. The solution value must be accompanied by a collection of cycles producing that value.

A companion of the max flow problem is the following *min cut problem*. For any cocycle D of G containing the special edge l, define the *capacity* of D to be the sum of the capacities of the edges in D other than l. Denote the capacity of D by $h(D)$. The min cut problem asks one to solve $\min h(D)$. The solution value must be accompanied by a cocycle D containing the edge l and having the solution value as capacity.

The famous max-flow min-cut theorem for graphs says that $\max F = \min h(D)$ no matter how the nonnegative integral edge capacities are selected.

The max flow and min cut problems for graphs have obvious matroid translations. In the above description, one replaces the graph G by a matroid M and specifies elements instead of edges, and circuits and cocircuits instead of cycles and cocycles. One might conjecture that the max-flow min-cut theorem still holds in the expanded setting. But this is not so in general. Counterexamples can be produced with small matroids, in particular with U_4^2, the rank 2 uniform matroid on four elements, and with F_7^*, the Fano dual matroid. Indeed, because of the symmetry of U_4^2 as well as of F_7^*, the conjecture is false for these matroids no matter which element is declared to be special. In general, there are matroids where the equality $\max F = \min h(D)$ does or does not hold for all capacity vectors h, depending on the selection of the special element. Thus, we are motivated to define a matroid M with a special element l to have the *max-flow min-cut property*, or to be a *max-flow min-cut matroid*, if $\max F = \min h(D)$ no matter which nonnegative integral values are assigned as capacities. Absence or presence of the max-flow min-cut property for M is evidently governed by the connected component of M containing the element l. Thus, for the purposes of characterizing the max-flow min-cut matroids, we might as well restrict ourselves to connected matroids.

Using an ingenious but complicated induction hypothesis, Seymour proved in a long paper that the connected max-flow min-cut matroids are precisely the connected binary matroids where the special element l is not contained in any F_7^* minor. In this chapter, we prove Seymour's result using a quite different approach. Let us define $\mathcal{M}$ to be the class of connected binary matroids where each matroid has a special element l such that l is not contained in any F_7^* minor. In Section 13.2, we establish 2- and Δ-sum decomposition theorems for the matroids of $\mathcal{M}$. In Section 13.3, we use these theorems to prove Seymour's result that $\mathcal{M}$ is precisely the class of connected max-flow min-cut matroids. In Section 13.4, we use a part of the proof to validate a construction for $\mathcal{M}$ that involves certain 2- and Δ-sums. We show that the construction can be determined for any matroid of $\mathcal{M}$ in polynomial time, and thus conclude that one can test for the max-flow min-cut property of binary matroids in polynomial time. We also describe polynomial algorithms for the following problems involving the matroids of $\mathcal{M}$: the max flow problem, the min cut problem, and a certain shortest circuit problem.

In Section 13.5, we examine an interesting graph application of the above results for $\mathcal{M}$. Let H be an undirected graph each of whose edges is declared to be odd or even. Recall that K_4 is the complete graph on four vertices. Declare a K_4 minor of H to be an odd-K_4 minor if it has a certain property that is defined via the relative position of its even and odd edges. Then let $\mathcal{G}$ be the class of 2-connected graphs without odd-K_4 minors. The graphs of $\mathcal{G}$ have pleasant properties, so an understanding of their structure is desirable. It turns out that the graphs of $\mathcal{G}$ can be linked to the class $\mathcal{M}$. As a result, we can apply the cited construction for $\mathcal{M}$ to understand the structure of the graphs of $\mathcal{G}$. In that way, we obtain a construction for the graphs of $\mathcal{G}$. At the same time, we produce a polynomial test for membership in $\mathcal{G}$. Evidently, we have moved from the max-flow min-cut property of graphs to the max-flow min-cut matroids, and then to the graphs without odd-K_4 minors. Thus, we have an instance of the cyclic process mentioned in the introductory paragraph.

The final section, 13.6, contains additional applications, extensions, and references.

The chapter makes use of Chapters 2, 3, and 5–11. It is also assumed that the reader has a basic knowledge of linear programming. Relevant references are included in Section 13.6.

13.2 2-Sum and Delta-Sum Decompositions

Recall that $\mathcal{M}$ is the class of connected binary matroids where each matroid has a special element l such that l is not contained in any F_7^* minor. In this section, we first show that any 2-separable matroid of $\mathcal{M}$ has a certain 2-sum decomposition where both component matroids are also in $\mathcal{M}$. Then we establish that any 3-connected nonregular matroid of $\mathcal{M}$ has a particular Δ-sum decomposition where again the components are in $\mathcal{M}$. These two results will be used in the next section to show that $\mathcal{M}$ is precisely the class of max-flow min-cut matroids. Before proceeding, the reader may want to review briefly the results for 2- and Δ-sums in Sections 8.2 and 8.5, respectively.

We begin with some definitions. By symmetry, there is essentially just one way to declare an element of the Fano matroid F_7 or of its dual F_7^* to be the special element l. When this is done, we get the matroid F_7 *with* l or F_7^* *with* l. Suppose two binary matroids M and M' contain the element l. If an isomorphism exists between M and M' that takes the element l of one of the matroids to l of the other one, then the two matroids are *l-isomorphic*. Finally, we emphasize that l is always the special element of any matroid in $\mathcal{M}$. Note that F_7 with l is in $\mathcal{M}$, while F_7^* with l is not.

We are ready for the detailed discussion of the 2- and Δ-sum results for $\mathcal{M}$.

2-Sum Decomposition

The structure theorem for the 2-sum case is as follows.

(13.2.1) Theorem. *Any 2-separable matroid $M \in \mathcal{M}$ on a set E has a 2-sum decomposition where both components M_1 and M_2 are connected* *minors of M, contain l, and thus are in $\mathcal{M}$. In addition, M_2 has an element $y \notin l$ so that any set $C \subseteq E$ is a circuit of M with l if and only if* (i) *or* (ii) *below holds.*

(i) *C is a circuit of M_2 with l but not y.*
(ii) *$C = (C_1 - \{l\}) \cup (C_2 - \{y\})$ where C_1 is a circuit of M_1 with l, and where C_2 is a circuit of M_2 with both l and y.*

Proof. Using the results of Section 8.2 for 2-separations and 2-sums, as well as the path shortening technique of Chapter 5, one readily shows that any 2-separable $M \in \mathcal{M}$ has a representation matrix B of the form

(13.2.2)

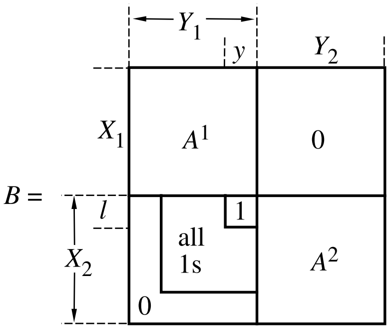

Matrix B of $M \in \mathcal{M}$ with exact 2-separation

Note the position of the element l in X_2, and the indicated element $y \in Y_1$. By Section 8.2, M is a 2-sum where both components M_1 and M_2 are connected minors of M, have the element l, and are represented by the matrices B^1 and B^2 of (13.2.3) below. Routine arguments using B, B^1 and B^2 confirm the claims of the theorem concerning the circuits of M, M_1, and M_2. □

(13.2.3)

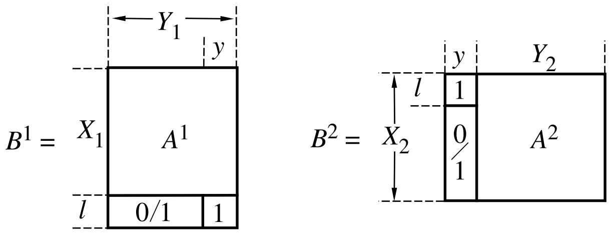

Matrices B^1 and B^2 of 2-sum decomposition of M

Delta-Sum Decomposition

We turn to the much more challenging case where M is 3-connected. In Section 13.3, it will be shown that any regular matroid of $\mathcal{M}$ has the max-flow min-cut property. Thus, we concentrate here on the situation where M is not regular. For that case, one can prove M either to be isomorphic to F_7, or to have a particular Δ-sum decomposition. The precise statement is as follows.

(13.2.4) Theorem. *Any 3-connected nonregular matroid $M \in \mathcal{M}$ on a set E is isomorphic to the Fano matroid F_7, or has a Δ-sum decomposition where the components M_1 and M_2 are connected minors of M, contain l, and thus are in $\mathcal{M}$. In the Δ-sum decomposition, both connecting triangles of M_1 and M_2 contain l. Let the remaining elements of the connecting triangle of M_1 (resp. M_2) be a and b (resp. v and w). Then any set $C \subseteq E$ is a circuit of M with l if and only if* (i), (ii), (iii), *or* (iv) *below holds.*

(i) *C is a circuit C_1 of M_1 with l but without a and b.*
(ii) *C is a circuit C_2 of M_2 with l but without v and w.*
(iii) *$C = (C_a - \{a\}) \cup (C_v - \{v\})$ where C_a is a circuit of M_1 with l and a but without b, and where C_v is a circuit of M_2 with l and v but without w.*
(iv) *$C = (C_b - \{b\}) \cup (C_w - \{w\})$ where C_b is a circuit of M_1 with l and b but without a, and where C_w is a circuit of M_2 with l and w but without v.*

The proof of Theorem (13.2.4) takes up the remainder of this section. We proceed as follows. First, we show that any 3-connected nonregular $M \in \mathcal{M}$ is isomorphic to F_7 or has a minor that is l-isomorphic to the matroid N_8 of $\mathcal{M}$ defined by the matrix B^8 below.

(13.2.5)

$$B^8 = \begin{array}{cc|ccc|c} & & & Y_1 & & Y_2 \\ & & u & v & w & l \\ \hline X_1 & d & 1 & 1 & 1 & 0 \\ \hline & a & 1 & 0 & 1 & 1 \\ X_2 & b & 1 & 1 & 0 & 1 \\ & c & 0 & 1 & 1 & 0 \end{array}$$

Matrix B^8 for the matroid N_8

Evidently, $(X_1 \cup Y_1, X_2 \cup Y_2)$ is a 3-separation of N_8. Note that Y_2 contains just the element l. It is easy to confirm that N_8 is indeed in $\mathcal{M}$.

Next, we establish that M has a minor N with the following properties. N contains l and is l-isomorphic to N_8, and a 3-separation of N corresponding to $(X_1 \cup Y_1, X_2 \cup Y_2)$ of N_8 under one of the l-isomorphisms induces a 3-separation of M. From that induced 3-separation of M, we finally derive the Δ-sum decomposition claimed in Theorem (13.2.4).

Induced 3-Separation

We begin the detailed discussion. The first lemma deals with the presence of N_8 minors in M.

(13.2.6) Lemma. *Let M be a 3-connected nonregular matroid of $\mathcal{M}$. Then M is isomorphic to the Fano matroid F_7, or has a minor that is l-isomorphic to N_8.*

Proof. If M has seven elements, then it must be isomorphic to F_7 since membership in $\mathcal{M}$ rules out F_7^*. So assume that M has at least eight elements. If M does not have any F_7^* minors, then by the splitter result of Lemma (11.3.19), M has an F_7 minor and is 2-separable, or is regular. The assumptions made here rule out these cases. Thus, M has an F_7^* minor, say N. Obviously, the minor N does not contain the element l.

According to Lemma (5.2.4), M has a connected N' minor that is a 1-element extension of N by l. If in N' the element l is parallel to or in series with some other element, then clearly l is part of an F_7^* minor of N' and thus of M, a contradiction. Thus, N' is 3-connected. Now F_7^* has no 3-connected 1-element expansion, so N' is obtained from N by addition of l. Routine calculations confirm that up to l-isomorphism, there is just one addition case that does not have an F_7^* minor with l. That case is represented by the matrix of (13.2.5). $\square$

From now on, we assume that M is a 3-connected nonregular matroid of $\mathcal{M}$ with at least eight elements. By Lemma (13.2.6), M has a minor that is l-isomorphic to N_8. For one such minor, we will exhibit a 3-separation that corresponds to the 3-separation $(X_1 \cup Y_1, X_2 \cup Y_2)$ of N_8 under one of the l-isomorphisms, and that induces a 3-separation of M. For this task, we invoke Corollary (6.3.25) and Theorem (6.3.28). Recall that Corollary (6.3.25) contains sufficient conditions for induced k-separations, while Theorem (6.3.28) extends those conditions to the case where M has a special subset L of elements. For the case at hand, we define L to be the set containing just l. We combine Corollary (6.3.25) and Theorem (6.3.28) for this special set L to the following theorem.

(13.2.7) Theorem. *Suppose a 3-connected N given by the matrix B^N*

(13.2.8)
$$B^N = \begin{array}{c|c|c|} & Y_1 & Y_2 \\ \hline X_1 & A^1 & 0 \\ \hline X_2 & D & A^2 \\ \hline \end{array}$$

Partitioned matrix B^N for N

is in $\mathcal{M}$ and has $l \in (X_2 \cup Y_2)$. Assume that $(X_1 \cup Y_1, X_2 \cup Y_2)$ is a k-separation of N. Furthermore, assume that $N/(X_2 \cup Y_2)$ has no loops and that $N \backslash (X_2 \cup Y_2)$ has no coloops. Finally, assume for every 3-connected 1-element extension of N in $\mathcal{M}$, say by an element z, that the pair $(X_1 \cup Y_1, X_2 \cup Y_2 \cup \{z\})$ is a k-separation of that extension. Then for any 3-connected matroid $M \in \mathcal{M}$ with a minor l-isomorphic to N, the following holds. Any k-separation of any such minor that corresponds to $(X_1 \cup Y_1, X_2 \cup Y_2)$ of N under one of the l-isomorphisms induces a k-separation of M.

We use Theorem (13.2.7) plus the recursive decomposition scheme explained in Section 10.5 to establish 3-separations for the nonregular 3-connected matroids of $\mathcal{M}$. We do not repeat here details of the decomposition scheme, so the reader may want to review that material before proceeding. The first iteration of that scheme is effectively accomplished by the following lemma.

(13.2.9) Lemma. *Let M be a 3-connected matroid of $\mathcal{M}$ with N_8 as minor. Then the 3-separation $(X_1 \cup Y_1, X_2 \cup Y_2)$ of N_8 as given by (13.2.5) induces a 3-separation of M, or M has a minor that is l-isomorphic to the matroid N_9 represented by B^9 below.*

(13.2.10)

$$B^9 = \begin{array}{cc|ccc|c|c|} & & \multicolumn{3}{c|}{Y_1} & Y_2 & \\ & & u & v & w & l & z \\ \hline X_1 & d & 1 & 1 & 1 & 0 & 1 \\ \hline & a & 1 & 0 & 1 & 1 & 1 \\ X_2 & b & 1 & 1 & 0 & 1 & 0 \\ & c & 0 & 1 & 1 & 0 & 0 \\ \hline \end{array}$$

Matrix B^9 for the matroid N_9

Proof. We assume absence of the specified N_9 minors in M, and apply Theorem (13.2.7). The matroid N_8 plays the role of N of that theorem. Since in N_8 we have $Y_2 = \{l\}$, we have $l \in (X_2 \cup Y_2)$. It also is easily checked that N_8 satisfies the conditions of Theorem (13.2.7) involving $N/(X_2 \cup Y_2)$ and $N \backslash (X_1 \cup Y_1)$. For verification of the remaining conditions of the theorem, we must compute the 3-connected 1-element extensions N' of N_8, say by an element z, such that in N' the pair $(X_1 \cup Y_1, X_2 \cup Y_2 \cup \{z\})$ is not a 3-separation. We are done by contradiction once we show that any such case of N' is l-isomorphic to N_9 or is not in $\mathcal{M}$.

We first consider the addition of z. For this, we adjoin to B^8 of (13.2.5) a column $[g/h]$ representing the element z. The partition of $[g/h]$ corresponds to that of B^8. By assumption, $(X_1 \cup Y_1, X_2 \cup Y_2 \cup \{z\})$ is not a 3-separation of N'. Thus, we must have $g = 1$. By the 3-connectedness of N', the vector h must have one or three 1s. Routine checking confirms that all such instances are l-isomorphic to N_9. We turn to the expansion by z. We adjoin to B^8 of (13.2.5) a row $[e \mid f]$ representing the element z. The partition of $[e \mid f]$ corresponds to that of B^8. The conditions imposed on z demand that e has one or three 1s and that $f = 1$. In each case, the element l of N' can be placed into an F_7^* minor, and thus is not in $\mathcal{M}$. $\square$

In the terminology of the decomposition scheme of Section 10.5, we have just completed the first iteration. We begin the second iteration by analyzing N_9 for 3-separations. That matroid has several such separations, some of them useful for our purposes, and others not. Indeed, in the initial investigation into the class $\mathcal{M}$, a 3-separation of N_9 was selected that led to a large number of iterations of the decomposition scheme. But eventually a better choice was found. It is the 3-separation $(\{a,b,d,u\},\{c,l,v,w,z\})$. It can be displayed in the accustomed format once we compute the representation matrix for N_9 corresponding to the base $\{b,c,v,w\}$. That matrix for N_9 is as follows.

(13.2.11)
$$\begin{array}{c|ccc|c|c|} & a & d & u & l & z \\ \hline b & 1 & 1 & 1 & 0 & 0 \\ w & 1 & 0 & 1 & 1 & 1 \\ v & 1 & 1 & 0 & 1 & 0 \\ c & 0 & 1 & 1 & 0 & 1 \\ \hline \end{array}$$

Matrix for N_9 corresponding to base $\{b,c,v,w\}$

The second iteration of the decomposition scheme is accomplished by the following lemma.

(13.2.12) Lemma. *Let M be a 3-connected matroid of $\mathcal{M}$ with N_9 as minor. Then the 3-separation $(\{a,b,d,u\},\{c,l,v,w,z\})$ of N_9 induces a 3-separation of M.*

Proof. The arguments are virtually identical to those for Lemma (13.2.9), except that this time each 3-connected 1-element extension satisfies the conditions of Theorem (13.2.7) or is not in $\mathcal{M}$. We leave it to the reader to fill in the details. $\square$

Lemmas (13.2.6), (13.2.9), and (13.2.12) imply the following theorem, which thus is the result of two iterations of the decomposition scheme.

(13.2.13) Theorem. *Let M be a 3-connected nonregular matroid of $\mathcal{M}$ with at least eight elements. Then M has a minor N with the following properties. N contains the element l and is l-isomorphic to N_8, and a 3-separation of N corresponding to the 3-separation $(X_1 \cup Y_1, X_2 \cup Y_2)$ of N_8 under one of the l-isomorphisms induces a 3-separation of M.*

Proof. By Lemma (13.2.6), M has a minor that contains l and that is l-isomorphic to the matroid N_8. We may suppose that N_8 itself is that minor. Then by Lemma (13.2.9), M has a 3-separation induced by the 3-separation $(X_1 \cup Y_1, X_2 \cup Y_2)$ of N_8, or M has a minor with l that is l-isomorphic to the matroid N_9. In the former case, we are done. In the latter case, we apply Lemma (13.2.12). Accordingly, M has a 3-separation induced by the 3-separation $(\{a,b,d,u\},\{c,l,v,w,z\})$ of N_9. From the matrix of (13.2.11) for N_9, we now delete the last column. Evidently, up to indices other than l, a matrix for N_8 results. The matroid N represented by that matrix is thus l-isomorphic to N_8. Furthermore, by the derivation of N from N_9, the 3-separation $(\{a,b,d,u\},\{c,l,v,w\})$ of N also induces the 3-separation of M derived earlier from N_9. But that 3-separation of N corresponds to the 3-separation $(X_1 \cup Y_1, X_2 \cup Y_2)$ of N_8. Thus, the case involving N_9 also leads to the desired conclusion. □

From 3-Separation to Delta-Sum

The reader probably anticipates that the conversion of the just-proved 3-separation of M to the Δ-sum decomposition claimed in Theorem (13.2.4) is straightforward. Unfortunately, the situation is not quite as simple. In particular, we must obtain some insight into the position of l relative to the 3-separation of M before we can proceed to the Δ-sum decomposition. We obtain this insight next.

By (13.2.5), the minor $N = N_8/c$ of N_8 is l-isomorphic to the Fano matroid F_7 with l. Indeed, the groundset of N is $\{a, b, d, l, u, v, w\}$, and the 3-separation $(\{d, u, v, w\}, \{a, b, l\})$ of N induces the same 3-separation in M that N_8 induces.

From now on, we do not need the matroid N_8 any more. Thus, we can switch notation, and can utilize the sets X_1, X_2, Y_1, and Y_2 so far employed for N_8 to denote induced 3-separations of M. We do this next in a representation matrix B of M. That matrix displays the just-defined F_7 minor N, with indices a, b, d, u, v, w, l and 3-separation $(\{d, u, v, w\}, \{a, b, l\})$.

(13.2.14)

$$B = \begin{array}{cc|ccc|c|c|}
 & & \multicolumn{3}{c|}{Y_1} & \multicolumn{2}{c}{Y_2} \\
 & & u & v & w & l & \\ \hline
X_1 & & \multicolumn{3}{c|}{A^1} & \multicolumn{2}{c|}{0} \\
 & d & 1 & 1 & 1 & 0 & \\ \hline
X_2 & a & 1 & 0 & 1 & 1 & \\
 & b & 1 & 1 & 0 & 1 & \\
 & & \multicolumn{3}{c|}{D} & \multicolumn{2}{c|}{A^2} \\ \hline
\end{array}$$

Matrix B of M displaying N
and induced 3-separation $(X_1 \cup Y_1, X_2 \cup Y_2)$

To capture the role of l in the 3-separation $(X_1 \cup Y_1, X_2 \cup Y_2)$ of M, we define l to *straddle* the 3-separation if a shift of l from Y_2 to Y_1 results in another 3-separation of M. The next theorem says that l must straddle the 3-separation. This fact will be essential for the Δ-sum decomposition to come.

(13.2.15) Theorem. *Let a 3-connected matroid $M \in \mathcal{M}$ be represented by the matrix B of* (13.2.14)*, where $(X_1 \cup Y_1, X_2 \cup Y_2)$ is a 3-separation of M. Then the element l straddles that 3-separation.*

Proof. Evidently, the element l straddles the 3-separation of M if and only if column l of the submatrix A^2 of B and column u of the submatrix

D of B are parallel. The former column vector we call g^l, and the latter one g^u. Note that g^u is the sum (in GF(2)) of the columns v and w of D.

Assume that l does not straddle the 3-separation. Thus, g^l and g^u are not parallel. Equivalently, one of these vectors contains a 1 in a row $x \in X_2$ where the other one has a 0. Thus, two cases are possible.

In the first case, $g^l_x = 0$ and $g^u_x = 1$. Then the submatrix of B indexed by a, b, d, x and l, u, v, w is either

$$
\text{(13.2.16)} \qquad
\begin{array}{c|ccc|c|}
 & u & v & w & l \\
\hline
d & 1 & 1 & 1 & 0 \\
\hline
a & 1 & 0 & 1 & 1 \\
b & 1 & 1 & 0 & 1 \\
x & 1 & 0 & 1 & 0 \\
\hline
\end{array}
$$

Submatrix of B

or is obtained from the matrix of (13.2.16) by exchanging the row indices a and b, and the column indices v and w. Thus, we may assume that the situation depicted by (13.2.16) is at hand. Let $\overline{M}$ be the minor represented by the matrix of (13.2.16). The minor $\overline{M} \backslash b$ of $\overline{M}$ turns out to be an F_7^* minor with l, and thus $M \notin \mathcal{M}$, a contradiction.

In the second case, $g^l_x = 1$ and $g^u_x = 0$. If row x of D is nonzero, or equivalently, if row x has 1s in both columns v and w, then rows a, b, d, x and columns l, v, w prove l to be in an F_7^* minor, a contradiction. Thus, row x of D is zero. By this fact and by the completion of the first case, we can narrow down the second case as follows. There is a nonempty subset $\overline{\overline{X}}_2 \subseteq X_2$ such that for all $x \in \overline{\overline{X}}_2$, we have $g^l_x = 1$ while row x of D is equal to 0. Furthermore, for all $x \in (X_2 - \overline{\overline{X}}_2)$, we have $g^l_x = g^u_x$.

Let $\overline{X}_2$ be the index set of the nonzero rows of D. By the above definition of $\overline{\overline{X}}_2$, we know $\overline{X}_2 \subseteq (X_2 - \overline{\overline{X}}_2)$. We now redraw B of (13.2.14) by enlarging and repartitioning the submatrices A^2 and D. The revised B is as follows.

(13.2.17)

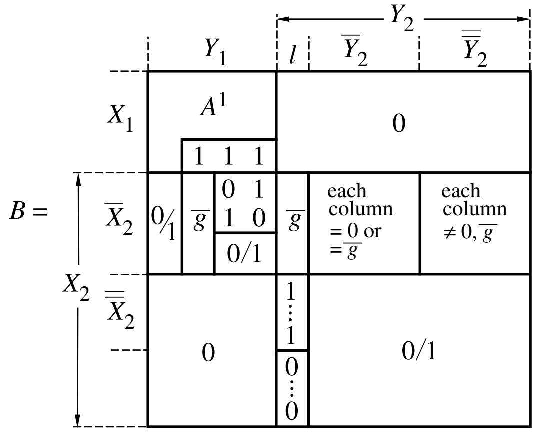

Repartitioned matrix B

The partition of the rows of B of (13.2.17) is induced by the sets $\overline{X}_2$ and $\overline{\overline{X}}_2$ introduced above. The partition of Y_2 into the sets l, $\overline{Y}_2$, and $\overline{\overline{Y}}_2$ is based on the information given in the submatrices indexed by $\overline{X}_2$, $\overline{Y}_2$ and $\overline{X}_2$, $\overline{\overline{Y}}_2$. Thus, each column of the former submatrix is zero or parallel to the subvector $\overline{g}$ of column l. By the definition of $\overline{X}_2$, the subvector $\overline{g}$ also occurs in D, as shown.

Suppose $\overline{\overline{Y}}_2 = \emptyset$. Then it is readily checked via B of (13.2.17) that $X_1 \cup Y_1 \cup \overline{X}_2$ is one side of a 2-separation of the 3-connected M, a contradiction. Thus, $\overline{\overline{Y}}_2 \neq \emptyset$.

Derive a matrix $\overline{B}$ from B by deleting all 1s in positions (x, y) where $x \in \overline{X}_2$ and $y \in (\{l\} \cup \overline{Y}_2)$. Consider the bipartite graph $\mathrm{BG}(\overline{B})$. If that graph does not have a path from l to $\overline{\overline{Y}}_2$, then arguments analogous to those of the proof of Lemma (5.2.11), prove M to be 2-separable. Thus, a path exists from l to some $y \in \overline{\overline{Y}}_2$. Because of path-shortening pivots, we may assume that the path has exactly two arcs. Thus, we can extract from $\overline{B}$ the following submatrix $\overline{\overline{B}}$,

(13.2.18)

$$
\begin{array}{cc|c|cc|c|c|}
 & & u & v & w & l & y \\
\hline
 & & 1 & 1 & 1 & \multicolumn{2}{c|}{0} \\
\hline
\overline{X}_2 & a & \overline{g} & 0 & 1 & \overline{g} & h \\
 & b & & 1 & 0 & & \\
 & & & \multicolumn{2}{c|}{0/1} & & \\
\hline
 & & \multicolumn{3}{c|}{0} & 1 & 1 \\
\hline
\end{array}
$$

Submatrix $\overline{\overline{B}}$ of $\overline{B}$

where $h \neq 0$ and $h \neq \overline{g}$. Reduce $\overline{X}_2$ to a minimal set containing a and b such that the just-mentioned condition is still satisfied by the correspondingly reduced $\overline{g}$ and h. Thus, $|\overline{X}_2| = 2$ or 3. Since $\overline{g}$ has two 1s in rows a and b, we have $|\overline{X}_2| = 3$ if and only if h has either two or no 1's in rows a and b. If h is a unit vector, then a pivot on the 1 of h plus deletion of the pivot column produces a previously resolved instance where $g_x^l \neq g_x^u$ and where a row x of D is nonzero. Otherwise, $|\overline{X}_2| = 3$, and h has two 1s in rows a and b. In each one of the three possible cases, a pivot in h plus deletion of the pivot column and of one row proves that l is in an F_7^* minor, a contradiction. □

As described in Section 8.5, a Δ-sum decomposition is derived from a 3-sum decomposition via a certain ΔY exchange. We first carry out the 3-sum decomposition of M using the matrix of (13.2.14), except that this time we ignore the column index u and the details of that column. We are guided in that decomposition by the matrices of (8.3.10) and (8.3.11). Below, we list the matrix B for M, as well as the representation matrices B^1 and $\tilde{B}^2$ for the components M_1 and $\tilde{M}_2$ of the 3-sum decomposition of M.

(13.2.19)

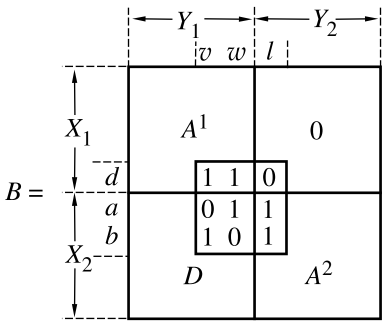

Matrix B for M

(13.2.20)

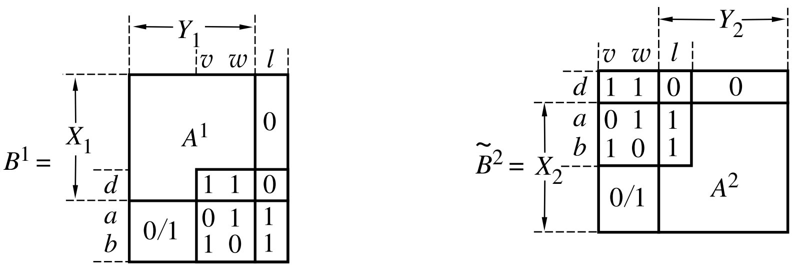

Matrices B^1 and $\tilde{B}^2$ of 3-sum decomposition of M

The Δ-sum decomposition has M_1 and a second matroid, say M_2, as components, where M_2 is derived from $\tilde{M}_2$ by the exchange of the triad $\{d, v, w\}$ with a triangle, say Z, according to the rule of (4.4.5). For the particular case at hand, we modify that exchange using the following arguments.

By Theorem (13.2.15), the element l straddles the 3-separation $(X_1 \cup Y_1, X_2 \cup Y_2)$ of M. Thus, in B, the columns v and w of the submatrix D span the column l of the submatrix A^2. This implies that one element of M_2 in the triangle Z is parallel to l. Thus, we can delete that element from M_2, and still have enough information to compose M_1 and the reduced M_2 to M again. Accordingly, we redefine M_2 to be that reduced matroid. From $\tilde{B}^2$ of (13.2.20), we see that the just-defined matroid M_2 may be taken to be $\tilde{M}_2/d$. Note that M_2 may have just six elements, in which case it is isomorphic to the wheel matroid $M(W_3)$.

For later reference, we include the matrices B^1 and B^2 of the Δ-sum decomposition of M into M_1 and M_2 in (13.2.21) below. We have simplified the matrices by omitting indices that are of no consequence for the subsequent discussion. With B^1 and B^2 of (13.2.21), one readily confirms the statements (i)–(iv) of Theorem (13.2.4) about certain circuits of M, M_1 and M_2. We leave the simple calculations to the reader. Thus, we have completed the proof of that theorem.

(13.2.21)

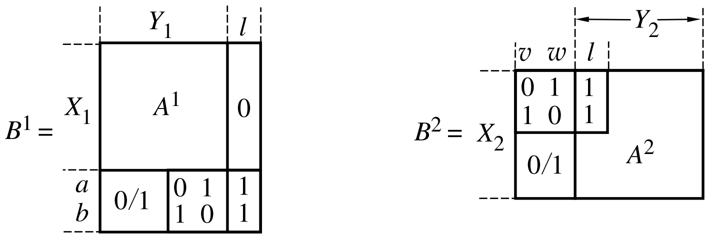

Matrices B^1 and B^2 of Δ-sum decomposition of M

In the next section, we use Theorems (13.2.1) and (13.2.4) to characterize the max-flow min-cut matroids with special element l by the exclusion of U_4^2 minors and of F_7^* minors with l.

13.3 Characterization of Max-Flow Min-Cut Matroids

Recall from the introduction that a max-flow min-cut matroid is a connected matroid with a special element l such that for any nonnegative integral edge capacity vector h, we have $\max F = \min h(D)$. Here $\max F$ is the optimal value of the max flow problem, where one must find a maximum number of cycles with l such that each element $e \neq l$ is contained in at most h_e of these cycles. Then $\min h(D)$ is the optimal value of the min cut problem, where one must determine a cocircuit D with l such that $h(D)$, defined to be the sum of the values h_e of the elements $e \in (D - l)$, is minimum.

In this section, we show that a connected matroid with a special element l has the max-flow min-cut property if and only if M has no U_4^2 minors and has no F_7^* minors with l. By Theorem (3.5.2), a matroid is binary if and only if it has no U_4^2 minors. Thus, the claimed characterization is equivalent to the statement that $\mathcal{M}$, the class of connected binary matroids having no F_7^* minors with l, is the class of connected max-flow min-cut matroids.

The characterization is summarized by the following theorem and corollary.

(13.3.1) Theorem. *A connected matroid M with a special element l is a max-flow min-cut matroid if and only if M has no U_4^2 minors and has no F_7^* minors with l.*

(13.3.2) Corollary. *$\mathcal{M}$ is the class of connected max-flow min-cut matroids.*

We begin the proof of Theorem (13.3.1) by noting that the result is trivially true for matroids whose groundset contains only the special element l. Thus, we assume from now on that all matroids M examined below have at least one additional element besides l. We first formulate the max flow problem and the min cut problem as linear programs with an integrality condition.

Max Flow Problem

Let M be a connected but not necessarily binary matroid whose groundset E contains a special element l. We construct the following matrix H from M. The rows of H correspond to the elements of E other than l, and the columns to the circuits C of M containing l. The entry of H in row e and column C is then 1 if e occurs in circuit C, and 0 otherwise. Consider the following linear program involving H and an arbitrary nonnegative integral vector h. Recall that $\underline{1}$ is a vector containing only 1s. All vectors are assumed to be column vectors of appropriate dimension. Below, the abbreviation "s. t." stands for "subject to."

$$\text{(13.3.3)} \qquad \begin{array}{ll} \max & \underline{1}^t \cdot r \\ \text{s. t.} & H \cdot r \leq h \\ & \quad\;\; r \geq 0 \end{array}$$

We call this problem the *fractional max flow problem*, since it becomes the max flow problem with capacity vector h when we require the solution vector r to be integral. Indeed, for any integral solution vector r, the entry in position C specifies the number of times the circuit C is to be selected.

Min Cut Problem

The linear programming dual of (13.3.3) is

$$\text{(13.3.4)} \qquad \begin{array}{ll} \min & h^t \cdot s \\ \text{s. t.} & H^t \cdot s \geq \underline{1} \\ & \qquad\; s \geq 0 \end{array}$$

We call (13.3.4) the *fractional min cut problem*. We justify the term next. Suppose that we require the solution vector s of (13.3.4) to be integral. According to the constraints of (13.3.4), any optimal solution s can then be assumed to be a $\{0,1\}$ vector. We do this from now on when we impose the integrality condition on s of (13.3.4). By the constraint $H^t \cdot s \geq \underline{1}$ of (13.3.4), the vector s is thus the incidence vector of a subset Z of E that intersects every circuit encoded by a column of H. Let $C^* = Z \cup \{l\}$. We claim that C^* contains a cocircuit of M with l. For a proof, select a base X of M that contains l and that avoids the set Z as much as possible. Collect in a set $\overline{Y}$ each element $y \in (E - X)$ whose fundamental circuit C_y with X contains l. By this definition, $\overline{Y} \cup \{l\}$ is the fundamental cocircuit of M that l forms with $(E - X)$. Furthermore, by the derivation of Z from the vector s, Z intersects, for each $y \in \overline{Y}$, the fundamental circuit C_y in some element z different from l. Suppose no such z is equal to y. Then for some $z \in Z$, $(X - \{z\}) \cup \{y\}$ is a base, which proves that X does not avoid Z as much as possible, a contradiction. Thus, $\overline{Y} \subseteq Z$, and C^* contains a cocircuit of M with l as claimed. Since the vector h is nonnegative, we may assume C^* to be that cocircuit.

By Lemma (3.4.25), any circuit of M with l and any cocircuit of M with l cannot intersect just on the element l. Thus, any cocircuit of M with l is a candidate for producing an integral solution for (13.3.4), and the best such candidate corresponds to a $\{0, 1\}$ vector s solving (13.3.4) with integrality condition. Thus, that problem represents the min cut problem, and we are justified in calling (13.3.4), without the integrality requirement, the fractional min cut problem for M.

Necessity of Excluded Minors Condition

We are ready to show the "only if" part of Theorem (13.3.1), which says that a connected max-flow min-cut matroid cannot have U_4^2 minors, or F_7^* minors with l. The proof is based on two reductions. The first one is accomplished by the following lemma and corollary. We omit the proof of the lemma, since it is just a particular version of the so-called duality theorem of linear programming.

(13.3.5) Lemma. *For any feasible vectors r and s of* (13.3.3) *and* (13.3.4), *respectively, we have $\underline{1}^t \cdot r \leq h \cdot s$, with equality holding if and only if both vectors r and s are optimal.*

(13.3.6) Corollary. *If every optimal solution vector r of* (13.3.3) *is nonintegral, then the matroid defining that problem is not a max-flow min-cut matroid.*

Proof. Let r and s be optimal solutions for (13.3.3) and (13.3.4), respectively. Because of the assumptions and Lemma (13.3.5), we must have $\max F < \underline{1}^t \cdot r = h^t \cdot s \leq h(D)$. Thus, $\max F \neq \min h(D)$, and the matroid defining (13.3.3) cannot be a max-flow min-cut matroid. □

For the second reduction, we rewrite Lemma (5.4.3) to get the following result.

(13.3.7) Lemma. *Let M be a connected matroid with an element l. If M has U_4^2 minors, then M has a U_4^2 minor with l.*

By the preceding reductions, we may prove the necessity of the excluded minors condition of Theorem (13.3.1) by producing, for any connected matroid with a U_4^2 minor with l or with an F_7^* minor with l, a nonnegative integral capacity vector h such that any optimal solution vector for (13.3.3) is nonintegral. The next lemma says that such a vector h can always be found.

(13.3.8) Lemma. *Let a connected matroid M with an element l have a U_4^2 minor with l or an F_7^* minor with l. Then there is a nonnegative integral vector h such that all optimal solution vectors r for* (13.3.3) *with that h are nonintegral.*

Proof. Let N be a minor of M with l and isomorphic to U_4^2 or F_7^*. By the discussion of Section 3.4, we may assume that $N = M/U \backslash W$, where U does not contain any cycle of M, and W does not contain any cocycle. Let M have n elements. Define the vector h by $h_e = 1$ for each element of N except l, $h_e = n$ for each $e \in U$, and $h_e = 0$ for each $e \in W$. Routine calculations show that any optimal solution vector r for (13.3.3) with this h is nonintegral. We leave the verification to the reader. $\square$

Sufficiency of Excluded Minors Condition

We prove the "if" part of Theorem (13.3.1), which says that any connected matroid M with a special element l, without U_4^2 minors, and without F_7^* minors with l is a max-flow min-cut matroid.

We invoke a result of polyhedral combinatorics to simplify the proof. That result concerns a certain integrality property of linear programs called *total dual integrality* and is due to Edmonds and Giles.

(13.3.9) Theorem. *Suppose the matrix A and the vector c of the linear program*

$$\text{(13.3.10)} \qquad \begin{array}{ll} \max & c^t \cdot f \\ s.\ t. & A \cdot f \leq b \\ & \quad\ \ f \geq 0 \end{array}$$

are integral and permit a feasible solution for the dual linear program of (13.3.10), *which is*

$$\text{(13.3.11)} \qquad \begin{array}{ll} \min & b^t \cdot g \\ s.\ t. & A^t \cdot g \geq c \\ & \quad\ \ g \geq 0 \end{array}$$

Furthermore, assume that the linear program (13.3.10) *has an integral optimal solution for every integral vector b for which it has a feasible solution. Then all extreme point solutions for the dual linear program* (13.3.11) *are integral.*

We apply Theorem (13.3.9) as follows. We view the fractional max flow problem (13.3.3) as an instance of (13.3.10). Then the linear programming dual of (13.3.3), which is the fractional min cut problem (13.3.4), is the problem (13.3.11). Note that (13.3.3) has a feasible solution if and only if $h \geq 0$. Suppose for a given M with l, and for any integral $h \geq 0$, we can show that the fractional max flow problem has an optimal solution vector that is integral. By Theorem (13.3.9), the fractional min cut problem then has, for any integral $h \geq 0$, an optimal solution that is integral. By Lemma (13.3.5), the optimal objective function values of the two problems agree, and thus we have $\max F = \min h(D)$. We conclude that M is a max-flow min-cut matroid.

By the above arguments, we may establish the max-flow min-cut property for M by showing that the fractional max flow problem (13.3.3) has, for any integral $h \geq 0$, an optimal solution vector that is integral. We divide the latter task into two parts. In the first one, we assume M to be regular. In the second one, M is assumed to be nonregular.

Regular Matroid Case

We begin with the first part. So assume that a connected matroid M on a set E and with special element l is regular. Let h be any nonnegative integral capacity vector. By Theorem (9.2.9), M has a real totally unimodular representation matrix B. Adjoin an identity matrix I to B, getting a real matrix $A = [I \mid B]$ whose columns are indexed by E. Consider the following linear program, where the solution vector f is a real column vector indexed by E.

$$\begin{array}{ll} \max & f_l \\ \text{s. t.} & A \cdot f = 0 \\ & \quad\; f \geq -h \\ & \quad\; f \leq h \end{array} \tag{13.3.12}$$

The linear program is clearly bounded and has the zero vector as solution. Thus, there is an optimal solution vector, say $\tilde{f}$. By linear programming results, $\tilde{f}$ may be assumed to be the solution to a system of equations of the form $\tilde{A} \cdot f = \tilde{h}$, where $\tilde{A}$ is a square nonsingular matrix consisting of some rows of A plus some unit vector rows, and where $\tilde{h}$ contains zeros and some entries of h and $-h$. Evidently, $\tilde{A}$ is totally unimodular. By Lemma (9.2.2), $\tilde{A}^{-1}$ is totally unimodular as well, and thus $\tilde{f} = \tilde{A}^{-1} \cdot \tilde{h}$ is integral. We may suppose $\tilde{f} \geq 0$ since any negative entry $\tilde{f}_e$ of $\tilde{f}$ can be transformed to a positive one by a clearly permissible scaling of column e of A by -1. Thus, $\tilde{f}$ is nonnegative and integral.

Until stated otherwise, we assume that $\tilde{f}_l > 0$. From $\tilde{f}$, we derive $\tilde{f}_l$ circuits that we later show to correspond to an optimal solution of the fractional max flow problem (13.3.3). We obtain these circuits by repeatedly solving a certain linear inequality system. In the first iteration, that system is

$$
\begin{aligned}
A \cdot g &= 0; \\
g_e &\geq 0; \text{ if } e \neq l \text{ and } \tilde{f}_e > 0 \\
g_e &= 0; \text{ if } e \neq l \text{ and } \tilde{f}_e = 0 \\
g_l &= 1
\end{aligned}
\tag{13.3.13}
$$

The vector $g = \tilde{f}/\tilde{f}_l$ is feasible for (13.3.13). Arguing analogously to the case involving (13.3.12), we are thus assured of an integral nonnegative solution $\tilde{g}$ for (13.3.13). Indeed, this time the vector playing the role of the earlier $\tilde{h}$ is a unit vector, and $\tilde{g}$ may be taken to be the characteristic vector of a circuit of M with l.

We now replace $\tilde{f}$ in (13.3.13) by $\tilde{f}' = \tilde{f} - \tilde{g}$ and deduce from the so-modified system (13.3.13) another circuit of M with l. We repeat this iterative derivation of circuits until $\tilde{f}_l$ has been reduced to 0. At that time, we have $\tilde{f}_l$ circuits, each containing the element l. By the derivation, any element e of M occurs in at most h_e of these circuits.

Now assume that $\tilde{f}_l = 0$. Then we do not select any circuit at all. The discussion to follow applies to this case as well as to the one where $\tilde{f}_l > 0$ and where we do select circuits.

Define a vector $\tilde{r}$ from the circuits of M with l that we just have selected by setting $\tilde{r}_C$ equal to the number of times the circuit C occurs in the collection. Thus, $\underline{1}^t \cdot \tilde{r} = \tilde{f}_l$. We claim that $\tilde{r}$ solves the fractional max flow problem (13.3.3). By the construction, $\tilde{r}$ is feasible for (13.3.3). We prove optimality as follows. Take any feasible solution r for (13.3.3). From each nonzero entry r_C of r, we derive a nonzero solution f^C for the equation $A \cdot f = 0$ as follows. First, we obtain a nontrivial $\{0, \pm 1\}$ solution for $A \cdot f = 0$ such that the support vector of that solution is the characteristic vector of C. This can clearly be done, since A is totally unimodular. Next, we scale that $\{0, \pm 1\}$ solution by r_C to get a vector f^C. We define f to be the sum of the vectors f^C if $f_l \neq 0$, and to be the zero vector otherwise. By the derivation, f is a solution for (13.3.12), with objective function value $f_l = \underline{1}^t \cdot r$. Recall that the vector $\tilde{f}$ is optimal for (13.3.12). Thus, $\underline{1}^t \cdot \tilde{r} = \tilde{f}_l \geq f_l = \underline{1}^t \cdot\ r$, which proves $\tilde{r}$ to be optimal for (13.3.3). We conclude that M is a max-flow min-cut matroid.

Nonregular Matroid Case

We turn to the second part, where M is not regular. We divide the proof into three subcases. First, we assume M to be F_7 with l; second, to have at least eight elements and to be 2-separable; and third, to have at least eight elements and to be 3-connected.

Fano Matroid Subcase

Routine calculations prove that F_7 with l is a max-flow min-cut matroid. To assist the reader with the checking, we list the matrix H of (13.3.3) for F_7 with l below, but otherwise omit all details. We have partitioned H to exhibit its structure.

$$
\begin{array}{|ccc|ccc|c|}
\hline
1 & 0 & 0 & 1 & 0 & 0 & 1 \\
0 & 1 & 0 & 0 & 1 & 0 & 1 \\
0 & 0 & 1 & 0 & 0 & 1 & 1 \\
\hline
1 & 0 & 0 & 0 & 1 & 1 & 0 \\
0 & 1 & 0 & 1 & 0 & 1 & 0 \\
0 & 0 & 1 & 1 & 1 & 0 & 0 \\
\hline
\end{array}
\tag{13.3.14}
$$

Matrix H of (13.3.3) for F_7 with l

For both the second and third subcase of the proof, we introduce the following terminology in connection with the fractional max flow problem (13.3.3). For any matroid under consideration, we define a *collection of weighted circuits* to be a collection of the circuits with l where a real nonnegative weight has been assigned to each circuit. We say that a collection of weighted circuits is *feasible* or *optimal* if the vector r with the weights of the circuits as entries is feasible or optimal for (13.3.3). The collection has *flow value* α if the corresponding objective function value of (13.3.3) is α. Finally, we say that a collection of weighted circuits *uses* an edge e α times if the weights of the circuits containing the element e sum to α.

The arguments of the second and third subcase use induction. The smallest instance involves the already treated F_7. Thus, for some $n \geq 7$, we assume the desired conclusion for matroids with at most n elements and take M to have $n+1$ elements. By the excluded minors condition, we know that M is in $\mathcal{M}$. Thus, the decomposition results of Section 13.2 apply.

2-Sum Decomposition Subcase

We use the decomposition Theorem (13.2.1), which we repeat below.

(13.3.15) Theorem. *Any 2-separable matroid $M \in \mathcal{M}$ on a set E has a 2-sum decomposition where both components M_1 and M_2 are connected minors of M, contain l, and thus are in $\mathcal{M}$. In addition, M_2 has an element $y \notin l$ so that any set $C \subseteq E$ is a circuit of M with l if and only if* (i) *or* (ii) *below holds.*

(i) *C is a circuit of M_2 with l but not y.*
(ii) *$C = (C_1 - \{l\}) \cup (C_2 - \{y\})$ where C_1 is a circuit of M_1 with l, and where C_2 is a circuit of M_2 with both l and y.*

The notation below is that of Theorem (13.3.15). In particular, M_1 and M_2 are the two components of a 2-sum decomposition of M.

First we find a collection of weighted circuits $\mathcal{C}$ that solves the fractional max flow problem for M for an arbitrary nonnegative integral vector h. Using parts (i) and (ii) of Theorem (13.3.15), we derive from $\mathcal{C}$ two collections $\mathcal{C}_1$ and $\mathcal{C}_2$ of weighted circuits for M_1 and M_2, as follows. Let C be a circuit of $\mathcal{C}$ of M with positive weight. By Theorem (13.3.15), C is a circuit of M_2 with l but not y, or $C = (C_1 - \{l\}) \cup (C_2 - \{y\})$ where C_1 is a circuit of M_1 with l, and where C_2 is a circuit of M_2 with both l and y. In the first case, we assign the weight of C to that circuit of M_2. In the second case, we assign the weight of C to both the circuit C_1 of M_1 and the circuit C_2 of M_2. By this construction, $\mathcal{C}_2$ has the same flow value as $\mathcal{C}$. Furthermore, $\mathcal{C}_1$ uses the element l of M_1 just as often as $\mathcal{C}_2$ of M_2 uses the element y, say α times.

Round up α to the next integer, getting, say, α'. Declare α' to be the capacity of the element y of M_2. Except for the element l of M_1 and for the elements l and y of M_2, assign the entries of the capacity vector h as capacities to the elements of M_1 and M_2.

Solve the fractional max flow problem for M_1. By induction and the earlier proof for regular matroids, we may assume that a collection $\mathcal{C}_1'$ of circuits with integral weights is found. Since $\mathcal{C}_1$ produces a feasible solution with flow value α for that problem, the new collection $\mathcal{C}_1'$ must have a flow value of at least α'.

Solve the fractional max flow problem for M_2. Once more, we may suppose that a collection $\mathcal{C}_2'$ of circuits with integral weights is found. The flow value of $\mathcal{C}_2'$ must be at least as large as that of $\mathcal{C}_2$, which in turn we know to be equal to that of $\mathcal{C}$.

Suppose $\mathcal{C}_2'$ uses the element y α'' times, for some $\alpha'' \leq \alpha'$. Derive a collection $\mathcal{C}_1''$ from $\mathcal{C}_1'$ by arbitrarily reducing weights of some circuits until the element l is used exactly α'' times. We combine $\mathcal{C}_1''$ with $\mathcal{C}_2'$ via Theorem (13.3.15) to a collection of circuits for M with integral weights and with flow value at least as large as that of $\mathcal{C}$, as desired.

Delta-Sum Decomposition Subcase

We turn to the final subcase, where the nonregular M has at least eight elements and is 3-connected. We use Theorem (13.2.4), which we list again below.

(13.3.16) Theorem. *Any 3-connected nonregular matroid $M \in \mathcal{M}$ on a set E is isomorphic to the Fano matroid F_7, or has a Δ-sum decomposition where the components M_1 and M_2 are connected minors of M, contain l, and thus are in $\mathcal{M}$. In the Δ-sum decomposition, both connecting triangles of M_1 and M_2 contain l. Let the remaining elements of the connecting* *triangle of M_1 (resp. M_2) be a and b (resp. v and w). Then any set $C \subseteq E$ is a circuit of M with l if and only if* (i), (ii), (iii), *or* (iv) *below holds.*

(i) *C is a circuit C_1 of M_1 with l but without a and b.*
(ii) *C is a circuit C_2 of M_2 with l but without v and w.*
(iii) *$C = (C_a - \{a\}) \cup (C_v - \{v\})$ where C_a is a circuit of M_1 with l and a but without b, and where C_v is a circuit of M_2 with l and v but without w.*
(iv) *$C = (C_b - \{b\}) \cup (C_w - \{w\})$ where C_b is a circuit of M_1 with l and b but without a, and where C_w is a circuit of M_2 with l and w but without v.*

The proof to come is similar to that for the 2-sum case. But there are subtle differences, as we shall see. We start again with a collection $\mathcal{C}$ of weighted circuits for M that solves the fractional max flow problem for an arbitrary nonnegative integral vector h. By Theorem (13.3.16), M has a Δ-sum decomposition into two matroids M_1 and M_2. We derive from $\mathcal{C}$ two collections $\mathcal{C}_1$ and $\mathcal{C}_2$ of weighted circuits for M_1 and M_2 by processing each circuit C of $\mathcal{C}$ as follows.

Suppose C is a circuit of case (i) or (ii) of Theorem (13.3.3). Then C occurs in M_1 or M_2, and we assign the same weight to that circuit in the corresponding collection $\mathcal{C}_1$ or $\mathcal{C}_2$. Suppose C is a circuit of case (iii) of Theorem (13.3.3). Thus, $C = (C_a - \{a\}) \cup (C_v - \{v\})$, where C_a is a circuit of M_1 with l and a but without b, and C_v is a circuit of M_2 with l and v but without w. Then we assign to both C_a and C_v the weight of C. Finally, case (iv) of Theorem (13.3.3) is handled analogously to case (iii).

Suppose the collection $\mathcal{C}_1$ uses the element a (resp. b) α times (resp. β times). By the above derivation, $\mathcal{C}_2$ uses the element v (resp. element w) also α times (resp. β times). If there are several choices for $\mathcal{C}$, we prefer one that minimizes $\alpha + \beta$.

We claim that $\alpha = 0$ or $\beta = 0$. Suppose both α and β are positive. Thus, there are, in the notation of parts (iii) and (iv) of Theorem (13.3.16), four circuits C_a, C_b, C_v, and C_w with positive weights, where C_a and C_b occur in M_1, and C_v and C_w in M_2. Let γ be the minimum of these four weights. Recall that $\{a, b, l\}$ is a triangle of M_1. Lemma (3.3.8) says that the symmetric difference of two disjoint unions of circuits of a binary matroid is a disjoint union of circuits. By two applications of this result, we see that the elements of $C_a \cup C_b \cup \{a, b, l\}$ not contained in exactly two of the circuits C_a, C_b, and $\{a, b, l\}$ of M_1 form a disjoint union of circuits that contains l but not a or b. Let C' be the circuit of that disjoint union containing l.

Derive two collections $\mathcal{C}_1'$ and $\mathcal{C}_2'$ from $\mathcal{C}_1$ and $\mathcal{C}_2$ as follows. First, add γ to the weight of the circuit C' in $\mathcal{C}_1$. Then reduce the weights of the four circuits C_a, C_b, C_v, and C_w by γ. Derive a collection $\mathcal{C}'$ for M from $\mathcal{C}_1'$ and $\mathcal{C}_2'$ of M_1 and M_2 in the by now obvious way. By the construction, $\mathcal{C}$ and

$\mathcal{C}'$ have the same flow value. Define α' and β' for $\mathcal{C}'$ analogously to α and β of $\mathcal{C}$. By the construction, $\alpha' + \beta' < \alpha + \beta$, a contradiction.

We thus may assume without loss of generality that $\beta = 0$. In the notation of Theorem (13.3.16), $\mathcal{C}$ produces circuits of type C_1 and C_a for $\mathcal{C}_1$, and of type C_2 and C_v for $\mathcal{C}_2$. Round up α to the next integer, getting α'. Assign α' as capacity to the element a of M_1 (resp. element v of M_2), and declare the element b of M_1 (resp. element w of M_2) to have capacity zero. To all other elements of M_1 and M_2, assign the capacity according to the vector h for M.

Solve the fractional max flow problem for M_1. By induction and the earlier proof for regular matroids, we may assume that a collection $\mathcal{C}_1'$ of circuits with integral weights is found. Suppose that collection does not use the element a exactly α' times. Then, analogously to the earlier situation where both $\alpha, \beta > 0$, one easily shows that $\mathcal{C}$ does not minimize $\alpha + \beta$.

Now solve the fractional max flow problem for M_2. We may assume that a collection $\mathcal{C}_2'$ of circuits with integral weights is found that uses the element v exactly α' times. It is a simple matter to combine $\mathcal{C}_1'$ and $\mathcal{C}_2'$ of M_1 and M_2 to a collection $\mathcal{C}'$ of circuits with integral weights for M. The latter collection is readily seen to have a flow value that is at least as large as that of $\mathcal{C}$. Thus, we have completed the last step in the proof of the excluded minors characterization of the connected max-flow min-cut matroids of Theorem (13.3.1).

In the next section, we utilize a part of the preceding proof of Theorem (13.3.1) to validate a construction of the connected max-flow min-cut matroids. We also describe polynomial algorithms for several problems involving these matroids.

13.4 Construction of Max-Flow Min-Cut Matroids and Polynomial Algorithms

In this section, we establish a construction of the connected max-flow min-cut matroids. We utilize the 2-sum and Δ-sum of Section 13.2, and rely on the construction of the regular matroids given by Theorem (11.3.16). From the proof of the construction, we derive a polynomial algorithm for deciding whether a binary matroid has the max-flow min-cut property. Finally, we describe polynomial algorithms that, for any connected max-flow min-cut matroid, solve the max flow problem, the min cut problem, and a certain shortest circuit problem. We begin with the construction of the connected max-flow min-cut matroids.

Construction of Max-Flow Min-Cut Matroids

We use the same terminology as for the construction of the regular matroids described in Theorem (11.3.16). In particular, the two initial matroids of any construction sequence, as well as all matroids that recursively are composed with the matroid already on hand, are the *building blocks*. Among the building blocks is the by now familiar matroid R_{10}. The two representation matrices for that matroid are given by (10.2.8) and (10.4.5).

We split the description of the construction into two cases, depending on whether the resulting max-flow min-cut matroid is to be regular. We begin with the case where this is so.

(13.4.1) Theorem. *Any connected regular max-flow min-cut matroid with special element l is graphic, cographic, or isomorphic to R_{10}, or may be constructed recursively by 2-sums and Δ-sums using as building blocks graphic matroids, cographic matroids, or matroids isomorphic to R_{10}. In the case of the recursive construction, one of the two initial building blocks contains the special element l, and no other building block contains that element.*

Proof. The theorem is nothing but Theorem (11.3.16) except for the occurrence of the special element l. It is a trivial matter to adapt the proof of Theorem (11.3.16) to account for that element. We leave it to the reader to fill in the details. □

Next we deal with the nonregular case.

(13.4.2) Theorem. *Any connected nonregular max-flow min-cut matroid with special element l is isomorphic to F_7, or may be recursively constructed by 2-sums given by* (13.2.2), (13.2.3), *and by Δ-sums given by* (13.2.19), (13.2.21). *In the case of the recursive construction, one of the two initial building blocks is isomorphic to F_7 and contains l. Furthermore, each additional building block also contains l, and is isomorphic to F_7 or is a connected regular max-flow min-cut matroid.*

A key result for the proof of Theorem (13.4.2) is the following composition theorem.

(13.4.3) Theorem. *Let M_1 and M_2 be two connected max-flow min-cut matroids that are represented by the matrices B^1 and B^2 of* (13.2.3) *or* (13.2.21). *Then the 2-sum or Δ-sum M of M_1 and M_2 as represented by the matrix B of* (13.2.2) *or* (13.2.19) *is a max-flow min-cut matroid.*

Proof. The part of the proof of Theorem (13.3.1) concerned with 2-sum and Δ-sum decompositions proves the result. □

Proof of Theorem (13.4.2). By Theorem (13.4.3), the compositions specified in Theorem (13.4.2) maintain the max-flow min-cut property if it is present in the components. Except for this observation, the remaining arguments are analogous to those proving Theorem (11.3.16). Indeed, the Δ-sum situation is easier to handle than the Δ-sum case of Theorem (11.3.16) since the element l straddles any 3-separation induced by an N_8 minor. Thus, the reader should have no difficulties in filling in the details. □

Polynomial Test for Max-Flow Min-Cut Property

The proofs of Theorems (13.4.1) and (13.4.2) just discussed are easily translated to polynomial algorithms that find the constructions claimed by these theorems. Any such method may be used to test for the max-flow min-cut property in binary matroids. The next theorem and corollary record these facts.

(13.4.4) Theorem. *There is a polynomial algorithm that for any connected max-flow min-cut matroid finds an applicable construction as described by Theorems* (13.4.1) *and* (13.4.2).

(13.4.5) Corollary. *There is a polynomial algorithm for determining whether an arbitrary binary connected matroid with a special element l has the max-flow min-cut property.*

Next we devise polynomial algorithms for the max flow problem, the min cut problem, and a certain shortest circuit problem. In each case, we assume the matroid to have the max-flow min-cut property. In general, these problems seem to be difficult for arbitrary binary matroids. For example, the min cut problem and the as-yet-unspecified shortest circuit problem can be shown to be $\mathcal{NP}$-hard.

Polynomial Algorithm for Shortest Circuit Problem

The *shortest circuit* problem is as follows. For each element e of a connected binary matroid M with special element l, a nonnegative rational distance d_e is given. Then one must find a circuit C of M containing l such that the sum of the distances d_e, $e \in C$, called the *length* of C, is minimal. This problem is well solved for the max-flow min-cut matroids, as follows.

(13.4.6) Theorem. *There is a polynomial algorithm for the shortest circuit problem of the max-flow min-cut matroids.*

Proof. Let M be the given connected max-flow min-cut matroid, with given nonnegative distance vector d. We first identify a construction as specified by Theorems (13.4.1) and (13.4.2). If M is isomorphic to F_7, then enumeration solves the problem. If M is regular, then we represent M by a real totally unimodular matrix B, define $A = [I \mid B]$, and solve the following linear program using any one of several known polynomial algorithms for linear programs.

$$
\begin{array}{lll}
 & \min & d^t \cdot f + d^t \cdot g \\
 & \text{s. t.} & A \cdot f - A \cdot g = 0 \\
 & & f \geq 0 \\
(13.4.7) & & g \geq 0 \\
 & & f_l = 1 \\
 & & g_l = 0
\end{array}
$$

Arguing as in Section 13.3, the linear program (13.4.7) must have $\{0, \pm 1\}$ solution vectors f and g. Indeed, the set $\{e \mid f_e = \pm 1 \text{ or } g_e = \pm 1\}$ may be assumed to be the desired shortest circuit of M with l.

In the remaining case, M is not regular. Then by Theorem (13.4.2), M is a 2-sum or Δ-sum where one component is regular or isomorphic to F_7. We discuss the Δ-sum case, and leave the easier 2-sum situation for the reader. Let M_1 and M_2 be the components of the Δ-sum decomposition of M. In agreement with the discussion of Section 13.2, let $\{a, b, l\}$ and $\{v, w, l\}$ be the triangles of M_1 and M_2 created by the decomposition.

We may assume that M_1 is isomorphic to F_7 or is regular. Except for a, b, and l, we assign the appropriate distance values of the vector d to the elements of M_1. To the former elements, we assign 0 as distance. Temporarily, we delete the element b from M_1, and solve the shortest circuit problem for that minor of M_1. Let C_a be the circuit so found, with length m_a. Similarly, by deletion of a from M_1, we obtain C_b and m_b. Finally, we delete both a and b, and get C_l and m_l. Note that $m_l \leq m_a + m_b$, as is readily proved by comparing m_l with the sum of the distances of the symmetric difference of C_a, C_b, and $\{a, b, l\}$.

Until stated otherwise, we assume that both $a \in C_a$ and $b \in C_b$. By Theorem (13.2.4), we conclude the following. M has C_l as shortest circuit, or M has a shortest circuit that intersects M_1 in C_a or C_b, or M has a shortest circuit that is contained in M_2. To decide which case applies, we assign m_a and m_b as distances to the elements v and w, respectively, of M_2 and recursively solve the shortest circuit problem in that matroid using the construction we already have on hand for it.

Let C be the circuit so found for M_2, with length m_2. If $m_2 \geq m_l$, then C_l is the desired shortest circuit of M. Otherwise, C contains just v, or just w, or none of v and w, or both of v and w. In the first three cases, the respective shortest circuit for M is $(C_a - \{a\}) \cup (C - \{v\})$, or $(C_b - \{b\}) \cup (C - \{w\})$, or C. In the fourth case, the previously derived inequality $m_l \leq m_a + m_b$ implies that $m_2 \geq m_l$, a situation already treated.

We now discuss the case for M_1 where $a \notin C_a$ or $b \notin C_b$. Thus, the shortest circuit of M_1 using a or b is at least as long as C_l, and hence we are justified to assign a distance larger than m_a or m_b to v or w of M_2 to get the correct conclusion by the preceding algorithm. It is easy to see that this method is polynomial.

Polynomial Algorithm for Max Flow and Min Cut Problems

We turn to the max flow problem and the min cut problem. The reader has surely recognized that the min cut problem is the shortest circuit problem for the dual matroid of the given max-flow min-cut matroid. Thus, we could adapt the preceding proof procedure to directly solve the min cut problem. We will not do so. Instead, we invoke Theorem (13.3.9) and a result of linear programming to achieve a brief presentation. The latter result is concerned with the efficient solution of linear programs and is due to Grötschel, Lovász, and Schrijver.

(13.4.8) Theorem. *Suppose there is a polynomial algorithm with the following features. For any linear program of a given class and any rational vector, the algorithm decides whether the vector is feasible for the linear program. If the answer is negative, the algorithm also obtains a violated constraint. Then there is a polynomial algorithm for the linear programs of the class and their duals. If any such linear program or its dual has extreme points, then an extreme point solution will be produced for that linear program.*

The next theorem contains the result that the max flow problem and the min cut problem can be solved in polynomial time.

(13.4.9) Theorem. *There is a polynomial algorithm that for any connected max-flow min-cut matroid and for any nonnegative integral capacity vector solves the max flow problem and the min cut problem.*

Proof. Let M be a connected max-flow min-cut matroid, and let h be a given vector of nonnegative integral capacities for the max flow problem and the min cut problem.

For the solution of the max flow problem, we first find a construction as described in Theorem (13.4.4). Note that the decompositions of the construction always involve at least one component that is isomorphic to F_7 or regular. Thus, we can carry out the proof procedure of Theorem (13.3.1), knowing that we only need to solve max flow subproblems for matroids that are isomorphic to F_7 or regular. The F_7 case is straightforward. For regular components, linear programs of the form (13.3.12) and inequality systems given by (13.3.13) must be solved. These tasks are handled efficiently by any polynomial algorithm for linear programs. Thus, the entire proof procedure for Theorem (13.3.1) can be converted to a polynomial algorithm for the max flow problem.

We turn to the min cut case. By the definition of the max-flow min-cut property, any max flow solution also solves the fractional max flow problem (13.3.3). Thus, by Theorem (13.3.9), the fractional min cut problem (13.3.4), which is the dual of the fractional max flow problem (13.3.3), has only integral extreme solutions. Thus, by Theorem (13.4.8), a polynomial solution algorithm for the min cut problem exists if we have a polynomial algorithm that decides whether a given rational vector s is feasible for (13.3.4). The latter algorithm must also identify a violated constraint of (13.3.4) if s is not feasible.

The feasibility test for s is trivial if that vector has any negative entries. So assume s to be nonnegative. View the entries of s as distances for the elements of M different from l, and assign 0 as distance to the latter element. Then deciding whether s is feasible for (13.3.4) is clearly equivalent to deciding whether the length of a shortest circuit of M with l is greater than or equal to 1. To answer the latter question, we compute a shortest circuit C of M using the polynomial algorithm of Theorem (13.4.6). If the length of C is at most 1, then the vector s is feasible for (13.3.4). Otherwise, the characteristic vector of C defines an inequality of (13.3.4) that is violated by s. □

In the next section, we meet an important graph application of the max-flow min-cut matroids.

13.5 Graphs without Odd-K_4 Minors

In this section, we transform a graph problem involving certain signed graphs into a matroid problem involving max-flow min-cut matroids. As a result, we can apply the construction and polynomial testing algorithm for max-flow min-cut matroids to obtain a construction and polynomial testing algorithm for the signed graphs. We begin with an application that gives rise to the signed graphs. Let P be a bounded polyhedron of the form $\{x \in \mathbb{R}^n \mid A \cdot x \leq a;\ x \geq 0\}$. We are to determine an inequality system $D \cdot x \leq d$ such that the polyhedron $Q = \{x \mid A \cdot x \leq a;\ D \cdot x \leq d;\ x \geq 0\}$ has only integral vertices and contains all integral points of P. The polyhedron Q is usually called the *integer hull* of P. The arrays D and d can have complex structure even for small arrays A and a. Thus, determination of the structure of D and d for special cases of A and a, either by some combinatorial description or in terms of a constructive scheme, has become one of the basic problems of polyhedral combinatorics.

An equally important converse problem is as follows. One postulates some construction for D and d and characterizes the cases of A and a for which the construction does work. For the description of one such construction, we need a rounding operation for rational vectors. The operation rounds down each entry of a given vector b to the next integer. We denote the resulting vector by $\lfloor b \rfloor$. The construction is as follows. Each row D_i of D is obtained from some linear combination of the rows of A by rounding down the entries. Thus, for some rational row vector $\lambda^i \geq 0$, we have $D_i = \lfloor \lambda^i \cdot A \rfloor$. The corresponding value d_i of the vector d is $\lfloor \lambda^i \cdot b \rfloor$. We call this construction *simple rounding*.

Just a few classes of combinatorially interesting polyhedra are known for which simple rounding does produce the desired D and d. The most important case is due to Edmonds and has as A any $\{0, \pm 1\}$ matrix with two ± 1s in each column. The vector a may contain any integers. The polyhedron $P = \{x \in \mathbb{R}^n \mid A \cdot x \leq a;\ x \geq 0\}$ for this case is known as the *fractional matching polyhedron*. On the other hand, it is well known that simple rounding generally does not work for the polyhedra produced by $\{0, \pm 1\}$ matrices A and integral vectors a where A has two ± 1s in each row instead of each column. Polyhedra of the latter variety are very important, since they arise from a number of combinatorial problems. An example is the vertex cover problem, where one must cover the edges of an undirected graph by vertices. The matrix A is then the real and negated version of the transpose of the node/edge incidence matrix of the graph, and the vector a contains only -1s.

There are special cases, though, where simple rounding does work for the latter polyhedra. To describe one such class, we derive a special graph $G(A)$ from any $\{0, \pm 1\}$ matrix A with two ± 1s in each row. Each column of A corresponds to a node of $G(A)$, and each row to an edge. Specifically, if a row of A has two entries of opposite (resp. same) sign, say in column j and k, then an edge connects the nodes j and k of $G(A)$ and is declared to be *even* (resp. *odd*). We call $G(A)$ a *signed* graph.

We define a cycle C of a signed graph to be *odd* if it contains an odd number of odd edges. If we scale a column j of A by -1, getting, say, A', then the graph $G(A')$ may be obtained from $G(A)$ by declaring each even (resp. odd) edge incident at node j to be odd (resp. even). We also say that $G(A')$ is obtained from $G(A)$ by *scaling*. Note that any cycle of $G(A)$ is even or odd if and only if this is so for the corresponding cycle of $G(A')$. Any graph obtained from $G(A)$ by repeated column scaling we call *equivalent* to $G(A)$. As expected, "is equivalent to" is an equivalence relation.

We derive minors from a signed graph by a sequence of deletions and contractions interspersed with scaling steps, with the restriction that any contraction may involve even edges only. It is easily checked that two minors produced by the same deletions and contractions, in any order, are equivalent, provided the reductions obey the cycle/cocycle condition of Section 2.2.

Consider the signed graph G obtained from K_4, the complete graph on four vertices, by declaring the edges of one triangle to be odd and the remaining three edges to even. Evidently, the triangles of G are precisely its odd circuits. Declare any signed graph that is equivalent to the just-specified one to be an *odd* K_4.

The signed graphs $G(A)$ without odd-K_4 minors have very pleasant properties. Indeed, the preceding simple rounding construction does produce the desired integer hull for a variant of the previously specified polyhedron involving the matrix A. The result, due to Gerards and Schrijver, is as follows.

(13.5.1) Theorem. *Let A be a $\{0, \pm 1\}$ matrix with two ± 1s in each row, and let P be the polyhedron $\{x \in \mathbb{R}^n \mid a^1 \leq A \cdot x \leq a^2;\ c^1 \leq x \leq c^2\}$. Then simple rounding can derive, for all integer vectors a^1, a^2, c^1, and c^2, a description of the integer hull of P if and only if $G(A)$ has no odd-K_4 minors.*

At this point, we have a nice theorem of polyhedral combinatorics involving the graphs without odd-K_4 minors, but actually have no clue what these graphs look like. Nor do we know how to decide efficiently whether a signed graph has odd-K_4 minors.

To gain the desired insight, we move from signed graphs G to certain binary matroids M. Let G be given. Define the edge set of G plus an additional element l to be the groundset of M. Next we produce a representation matrix for M. We begin with the usual binary node/edge incidence matrix F for G, and record whether edges are even or odd by appending to F an additional row having 0s (resp. 1s) in the columns corresponding to the even (resp. odd) edges of G. We adjoin a column unit vector with index l, having its 1 in the new row. Let the resulting matrix be F'. Finally, we declare the circuits of M to be the index sets of the GF(2)-mindependent column subsets of F'. Since scaling in G does not affect evenness or oddness of cycles, all graphs equivalent to G produce the same matroid M.

From F', we obtain a matrix $[I \mid B]$ by elementary row operations (in GF(2)). Thus B is a representation matrix for M. The subsequent analysis is simplified if we let l index a row of B, or equivalently, if we accept the unit vector of F' indexed by l as part of the identity I of $[I \mid B]$. Indeed, in that case we only need to choose a tree T of G and scale G so that all tree edges become even, and finally use the fundamental cycles of the new G to write down the columns of B. By the derivation, any such fundamental cycle is even (resp. odd) if the out-of-tree edge producing that cycle is even (resp. odd). It is easy to see that M is connected if G is 2-connected and has at least one odd cycle. Also, every 2-connected minor of G with at least one odd cycle corresponds to a minor of M with l, as expected. But what property does M have when G has no odd-K_4 minors? The next lemma, also due to Gerards and Schrijver, gives the surprising answer.

(13.5.2) Lemma. *Let G be a 2-connected signed graph, and let M*

be the corresponding binary connected matroid. Then G has no odd-K_4 minors if and only if M has no F_7^* minors with l, and thus, if and only if M is a max-flow min-cut matroid.

Proof. Assume that G has an odd-K_4 minor, say $\overline{G}$. Because of scaling, we may presume that a triangle of that minor has all edges odd, and that the remaining edges, which form a 3-star, are even. Derive a representation matrix $\overline{B}$ for the corresponding minor $\overline{M}$ of M as described above, using the 3-star as tree. Then one readily sees that $\overline{M}$ is an F_7^* minor of M with l. By reversing the arguments, we see that any F_7^* minor of M with l corresponds to an odd-K_4 minor of G. □

By Lemma (13.5.2) and Corollary (13.4.5), we already have a polynomial algorithm for deciding whether a signed graph has no odd-K_4 minors. We also have, indirectly, a complete analysis of the graphs without odd-K_4 minors, in the form of the decomposition Theorems (13.2.1) and (13.2.4) and the construction Theorems (13.4.1) and (13.4.2) for the max-flow min-cut matroids. To understand the structure of the graphs without odd-K_4 minors, we only need to translate these results into graph language. We carry out this task next.

Construction of Graphs without Odd-K_4 Minors

Let M be the connected max-flow min-cut matroid corresponding to a 2-connected signed graph G without odd-K_4 minors. We assume that G has at least one odd circuit. As before, l is the special element of M. Define G' to be the unsigned version of G. Since M/l is the graphic matroid of the 2-connected G', that matroid is connected. We will work out a description of G that involves scaling. That way, a particularly simple description results. We divide the discussion into five cases, depending on whether G or M has a 2-separation, and whether M is graphic, cographic, regular nongraphic and noncographic, or nonregular.

2-Separation Cases

Suppose $(E_1 \cup \{l\}, E_2)$ is a 2-separation of M. Suppose that $|E_1| = 1$, say $E_1 = \{e\}$. Recall that M/l is connected and graphic. Thus, e and l must be two series elements of M. Accordingly, G can be scaled so that e is the only odd edge. Let G be that scaled graph. We call a signed graph with all odd edges incident at one vertex, and thus G, a graph with one *partially odd* vertex. Later, we will see another instance of such a graph.

Suppose $|E_1| \geq 2$. Then (E_1, E_2) is a 2-separation of G. Select a base $X_1 \cup X_2$ of M where $X_1 \subseteq E_1$ and $l \in X_1$, and where X_2 is a maximal independent subset of E_2. In the corresponding representation matrix for

M, the row l has only 0s in the columns of $E_2 - X_2$. Correspondingly, G has, up to scaling, odd edges only in E_1. Assume such scaling. We derive from G the 2-sum components G_1 and G_2 as depicted in (8.2.8), where the explicitly shown connecting edges are declared to be even. Note that G_2 has only even edges. Later, we will encounter another 2-sum decomposition, so we define the present one to be a 2-*sum of type* 1.

As a result of the preceding discussion, we may assume from now on that M is 3-connected. This implies that any parallel class of edges of G contains at most two edges, one even and one odd.

Suppose the deletion of parallel edges from G creates a 2-separable graph. Then G has a 2-separation (E_1, E_2) where neither E_1 nor E_2 is just a set of parallel edges. Assume $|E_1| = 2$. Since E_1 is not just a set of parallel edges, the two edges of E_1 must be incident at a degree 2 node. But by scaling, both edges can be made even, so $(E_1, E_2 \cup \{l\})$ is a 2-separation of M contrary to the assumption that M is 3-connected. Hence, $|E_1|, |E_2| \geq 3$. Clearly, $(E_1 \cup \{l\}, E_2)$ and $(E_1, E_2 \cup \{l\})$ are 3-separations of M. In the terminology of Section 13.2, the element l straddles the 3-separation $(E_1, E_2 \cup \{l\})$ of M. Furthermore, that 3-separation induces a Δ-sum decomposition of M into M_1 and M_2, as depicted by the matrices B, B^1, and B^2 of (13.2.19) and (13.2.21). Since l straddles the 3-separation of M, the column l of the submatrix A^2 of B of (13.2.19) is spanned by the columns of the submatrix D of B. So if we GF(2)-pivot in column l of B and subsequently delete the pivot row, then the matrix D of B is reduced to a matrix with GF(2)-rank 1. We rely on this fact next when we interpret the Δ-sum decomposition of M in terms of a decomposition of G.

Suppose in each one of the matrices B, B^1, and B^2 of (13.2.19) and (13.2.21) we perform a GF(2)-pivot on the 1 in row a and column l. Let $\tilde{B}$, $\tilde{B}^1$, and $\tilde{B}^2$ result. Delete row l from $\tilde{B}$, $\tilde{B}^1$, and $\tilde{B}^2$, getting B', B'^1, and B'^2. Clearly, the matrix B' corresponds to G', the unsigned version of G. Since l straddles the 3-separation of M and B, the matrices B'^1 and B'^2 are readily seen to correspond to a somewhat unusual 2-sum decomposition of G. Indeed, the component graphs are as depicted in (8.2.8) once we replace the connecting edge x of G_1 in (8.2.8) by two parallel edges a and b, and the connecting edge y of G_2 in (8.2.8) by two parallel edges v and w. Assume such a replacement has been done, and denote by G_1' and G_2' the graphs corresponding to B'^1 and B'^2. Because of scaling, we may suppose that the signature of G agrees with that given by row l of $\tilde{B}$. Then according to row l of $\tilde{B}^1$ and $\tilde{B}^2$, we may sign G_1' and G_2' to obtain graphs for the matroids M_1 and M_2 of the Δ-sum decomposition, as follows. In G_1' (resp. G_2'), we declare the edge a (resp. w) to be odd, and the edge b (resp. v) to be even. The remaining edges of G_1' and G_2' are signed in agreement with the signature of the edges of G. Let G_1 and G_2 be the graphs so obtained for M_1 and M_2. We call this special 2-sum decomposition a 2-*sum of type* 2.

Theorem (11.2.10), Corollary (11.2.12), and Theorem (13.4.3), which cover the composition of regular matroids and of max-flow min-cut matroids, imply that the above 2-sum decompositions of type 1 and 2 are reversible. That is, if both components G_1 and G_2 correspond to regular matroids (resp. max-flow min-cut matroids), then the 2-sum G of type 1 or 2 also corresponds to a regular matroid (resp. max-flow min-cut matroid).

We have completed the 2-separation cases. From now on, we assume that M is 3-connected, and that G is 3-connected up to parallel edges. In fact, any parallel class of edges of G consists of at most two edges, one even and one odd.

Graphic Matroid Case

We assume M to be graphic. Thus, there is a graph H so that M is the graphic matroid $M(H)$ of H. Recall that G' is the unsigned version of G. Since M/l is the graphic matroid of H/l as well as of G', the graphs H/l and G' must be 2-isomorphic. We know that G is 3-connected up to parallel edges, so by Theorem (3.2.36), $H/l = G'$. The graph H may thus be derived from G' by splitting some vertex v into two nodes, which are then connected by the edge l. Pick a spanning tree X' of G' that has the node v as tip node. Then $X = X' \cup \{l\}$ is a tree of H. The representation matrix of M corresponding to the base X has l as row index. The 1s in that row correspond to a subset $\overline{E}$ of the edges of G incident at v. Thus, up to scaling, the odd edges of G are precisely the edges of $\overline{E}$. Assuming such scaling, G is therefore a graph with one partially odd vertex. By reversing the construction, we see that any signed graph G with one partially odd vertex produces a graphic M.

Cographic Matroid Case

We assume M to be cographic but not graphic. Hence, M^* is graphic. As before, G' represents M/l, so both M/l and $(M/l)^* = M^*\backslash l$ are planar. We conclude that G' and G may be taken to be 3-connected plane graphs plus parallel edges, and that there is a graph H such that $M(H) = M^*$ and $G' = (H\backslash l)^*$. In H, let v_1 and v_2 be the endpoints of the edge l. Pick a tree of H that does not contain the edge l. In the corresponding representation matrix for M^*, the 1s of column l represent a path P from v_1 to v_2. Thus, in $H\backslash l$, each vertex other than v_1 and v_2 has an even number of edges of P incident. By scaling in the plane graph G, the edges of P become precisely the odd edges. Thus, up to scaling, exactly the two faces of G corresponding to v_1 and v_2 of $H\backslash l$ have an odd number of odd edges incident. The latter property is invariant under scaling, so the same property holds for G. We call a planar graph with this property, and thus G, a graph with *two odd faces*.

Regular Nongraphic and Noncographic Matroid Case

We turn to the case where the 3-connected matroid M is regular but not graphic and not cographic. By Theorem (13.4.1), M is isomorphic to R_{10}, or is a Δ-sum with components M_1 and M_2 where one component does not contain l. The R_{10} case is not possible, since M/l is graphic, and since the contraction of any element of R_{10} produces an $M(K_{3,3})^*$ minor, as shown in the proof of Lemma (10.4.4). Thus, we only need to examine the case of a Δ-sum. Let $(E_1 \cup \{l\}, E_2)$ be the corresponding 3-separation of M. By the derivation of Theorem (13.4.1) via Theorem (11.3.16), we may assume that $|E_1 \cup \{l\}|, |E_2| \geq 6$. The pair (E_1, E_2) cannot be a 2-separation of G since otherwise $|E_1|, |E_2| \geq 5$ implies that the earlier discussed 2-separation case applies.

Suppose $(E_1, E_2 \cup \{l\})$ is also a 3-separation of M. In the terminology of Section 13.2, the element l straddles the 3-separation $(E_1 \cup \{l\}, E_2)$. Let B of (13.2.19) be a representation matrix for M exhibiting the 3-separation $(E_1, E_2 \cup \{l\})$, except that we ignore in B the explicitly shown 0s and 1s in the submatrix indexed by $\{a, b, d\}$ and $\{l, v, w\}$. Since l straddles the 3-separation, the submatrix D of B spans the column l of A^2. Then a GF(2)-pivot on any 1 in column l of A^2, followed by the deletion of the pivot row and pivot column, reduces D, which has GF(2)-rank $D = 2$, to a matrix with GF(2)-rank 1. We conclude that (E_1, E_2) is a 2-separation of M/l, and hence of G, contrary to assumption. Hence, l does not straddle the 3-separation. Thus, analogously to the case of a 2-sum of type 1, there is a base $X_1 \cup X_2$ for M where $X_1 \subseteq E_1$ and $l \in X_1$, and where X_2 is a maximal independent subset of E_2. Furthermore, up to scaling, all edges of G in E_2 are even. For the remainder of this case, we suppose G to be of this form.

Assume that (E_1, E_2) is a graph 3-separation of G. Arguing as in the 2-sum case, G has a Δ-sum decomposition where we declare all edges of the connecting triangles to be even, and where the component defined from E_2 has even edges only. Without chance of confusion, we call this process a Δ-*sum* decomposition of G. By the composition Corollary (11.2.12) for regular matroids, the Δ-sum decomposition of G is reversible.

So suppose (E_1, E_2) is not a graph 3-separation of G. This is possible, by Theorem (3.2.25). Recall that each parallel class of G has at most two edges, one even and one odd. Since the edges of E_2 are even, all parallel edges must occur in E_1. Define G'' to be G minus all odd parallel edges. Correspondingly, define E_1'' and E_2'' from E_1 and E_2, respectively. Since $|E_1|$, $|E_2| \geq 5$, we have $|E_1''|$, $|E_2''| \geq 3$. Let M'' be the minor of M corresponding to G''. Then (E_1'', E_2'') is a 3-separation of M'' that induces the 3-separation $(E_1 \cup \{l\}, E_2)$ of M. If (E_1'', E_2'') is a graph 3-separation of G'', then (E_1, E_2) is a graph 3-separation of G contrary to assumption. Thus, by Theorem (3.2.25), E_1'' is a cutset of cardinality 3, and removal of that cutset from G produces two connected graphs H_{21} and H_{22} whose edge sets form a partition of E_2. Furthermore, each one of the latter graphs contains a cycle.

If H_{21} or H_{22} has at least four edges, then a graph 3-separation of G exists where one side contains only even edges, and the earlier case applies. Thus, each one of H_{21} and H_{22} is just a triangle, and G consists of a pair of triangles whose nodes are pairwise joined by three edges, say e, f, and g, plus edges parallel to the latter edges. Then G corresponds to a cographic matroid case unless G is of the form

(13.5.3)

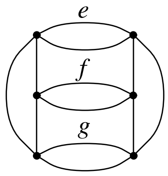

Exceptional signed graph case

where e, f, and g are odd, and where all other edges are even. The matroid M for this graph is regular, as is easily checked.

At this point, we want to summarize the above analysis for the situation where M is regular. For that summary, in Theorem (13.5.4) below, we need a characterization of the regularity of M in terms of signed graphs. To this end, we define an *odd double triangle* to be the following signed graph. We start with a triangle, all of whose edges are even. Then we add to each edge one parallel odd edge. Using a representation matrix for F_7, for example the matrix of (10.2.4), one readily confirms that M has an F_7 minor with l if and only if G has an odd double triangle minor. It is easy to check that M is regular if and only if M has no F_7 minors with l and no F_7^* minors with l. Indeed, this claim follows almost immediately from the discussion of Section 13.2. We thus have the following theorem for the case where M is regular.

(13.5.4) Theorem. *Let G be a 2-connected signed graph without odd-K_4 minors and without odd double triangle minors, but with an odd circuit. Then up to scaling, G has exactly one partially odd vertex, or is planar and has exactly two odd faces, or is the graph of (13.5.3), or may be constructed recursively by 2-sums of type 1 or 2, and by Δ-sums. Up to scaling, the building blocks are the cited graphs and graphs having only even edges.*

Proof. The above analysis and a simple inductive argument prove the result.

□

We turn to the final case, where M is nonregular.

Nonregular Matroid Case

Suppose G produces a nonregular max-flow min-cut matroid M. The translation of the construction Theorem (13.4.2) for M gives the following result.

(13.5.5) Theorem. *Let G be a 2-connected signed graph without odd-K_4 minors, but with an odd double triangle minor. Then up to scaling, G is an odd double triangle, or may be recursively constructed by 2-sums of type 1 or 2. In the case of the recursive construction, one of the initial building blocks is an odd double triangle. Furthermore, each additional building block is also an odd double triangle, or corresponds to a signed graph having no odd-K_4 minors and no double triangle minors.*

Proof. As shown above, matroid 2-sums and matroid Δ-sums with l as straddling element correspond to graph 2-sums of type 1 or 2. Thus, the theorem follows from Theorem (13.4.2). □

In the next section, we include applications, extensions, and references.

13.6 Applications, Extensions, and References

The first comprehensive treatment of the max flow problem for graphs is contained in Ford and Fulkerson (1962). The definition and characterization of max-flow min-cut matroids by excluded minors is due to Seymour (1977a). Related is a characterization of matroid ports in Seymour (1976), (1977b). The decomposition of the max-flow min-cut matroids is described in Tseng and Truemper (1986). That reference proves a decomposition result slightly stronger than Theorem (13.2.13): Any N_8 minor of a max-flow min-cut matroid M has a 3-separation that up to indices other than l is given by (13.2.5) and that induces a 3-separation of M. A short proof of Theorem (13.2.13) is given in Bixby and Rajan (1989).

Basic linear programming results may be found in any standard text, for example in Dantzig (1963), Chvátal (1983), or Schrijver (1986). The concept of total dual integrality and the related Theorem (13.3.9) are due to Edmonds and Giles (1977).

The existence of polynomial algorithms for the various problems of Section 13.4 is established in Truemper (1987a). That reference proves the min cut problem and the shortest circuit problem to be $\mathcal{NP}$-hard, and also shows that the max flow problem is solved by an integer extreme point solution of the linear program (13.3.3), not just an integer solution as proved here. Theorem (13.4.8) is due to Grötschel, Lovász, and Schrijver (1988). Actually, simpler machinery suffices to prove the results of

Section 13.4, as shown in Truemper (1987a). The recognition problem of max-flow min-cut path matrices is treated in Hartvigsen and Wagner (1988).

The simple rounding scheme is analyzed in Chvátal (1973). The most important case treatable by simple rounding is the matching problem. Its solution is due to Edmonds (1965c), (1965d); see Lovász and Plummer (1986) for an excellent exposition of this result and of many related ones. Theorem (13.5.1) and Lemma (13.5.2) are from Gerards and Schrijver (1986). General results for packing and covering problems involving matroid circuits are described in Seymour (1980c). Additional decomposition theorems and other results for signed graphs are given in Gerards (1988), (1989b). Gerards, Lovász, Schrijver, Seymour, and Truemper (1991) summarize various decomposition results for signed graphs.

In Section 13.4, it is shown that the fractional min cut problem (13.3.4) has only integer extreme point solutions, provided max-flow min-cut matroids are involved. The latter condition is sufficient, but not necessary. Thus, one may want to characterize when precisely (13.3.4) has only integer extreme point solutions. Much progress has been made on this difficult problem (see Lehman (1981), Seymour (1990), and Cornuéjols and Novick (1994)), but finding a complete characterization seems to be very difficult.