A Counting Problem in Linear Programming

Jim Lawrence

Department of Mathematical Sciences
George Mason University
Fairfax, VA 22030–4444
lawrence@gmu.edu

Summary. Using a popular setup, in solving a linear programming problem one looks for a tableau of the problem which has no negative elements in the last column and no positive elements in the last row. We study a matrix whose (i, j)–th entry counts the tableaux for such a problem (here taken to be totally nondegenerate) which have i negative elements in the last column and j positive elements in the last row. It is shown that this matrix possesses a certain symmetry, which is described.

Key words: Linear programming; oriented matroid; tableau enumeration; polytope; h–vector; f–vector.

1 Introduction

It is assumed that the reader has some familiarity with linear programming. If not, there are a great many books which were written to serve as textbooks on linear programming; they are all, to some extent, descendants of the first such book, [2], written by Saul Gass.

Suppose we have a linear programming problem (assumed to be "totally nondegenerate"), having s nonnegative variables and r additional inequality constraints. In order to solve the problem using the simplex method, we may construct a suitable tableau and, by pivoting, attempt to move to a tableau in which the last column has no negative entries and the last row has no positive entries (ignoring the entry they share in common). A crude measure of progress toward this solution tableau is indicated by the pair of numbers (a, b), where a is the number of negative entries in the last column and b is the number of positive entries in the last row of the current tableau. Of course, these numbers may go up and down during the trek; it's not at all clear what this information tells us about getting to the solution. Even so, when these numbers aren't too big, intuition seems to dictate that we are "getting warmer," and that a tableau with $a = b = 0$ may not be many steps away. Thus we are led to the question of what can be said about the

$(r+1) \times (s+1)$ matrix N whose (a,b)–th entry (where $0 \leq a \leq r$, $0 \leq b \leq s$) is the number of tableaux for the problem having a negative entries in the last column and b positive entries in the last row. The main theorem below gives a simple property of this matrix.

This question is related to certain enumerative results concerning convex polytopes, arrangements of hyperplanes, and oriented matroids. Since our linear programming problem is totally nondegenerate, and under the assumption that the feasible region for the problem is nonempty and bounded, this feasible region is a simple convex polytope P of dimension s. ("Simple" means that, at each vertex, exactly s of the $r+s$ inequality constraints are satisfied with equality — each vertex lies on exactly s facets. Equivalently, exactly s edges emanate from each vertex.) Letting f_k denote the number of k–dimensional faces of P, so that f_0 is the number of vertices, f_1 is the number of edges, f_{s-1} is the number of facets, and $f_s = 1$, the vector $(f_0, \ldots, f_s)$ is the *f–vector of P*.

Each vertex of P is on exactly s edges, since P is simple. The objective function is not constant on any edge of this polytope, again by total nondegeneracy. We say that an edge which emanates from a vertex is *leaving* if the value of the objective function at that vertex is less than its value at the other vertex of the edge (so that the objective function increases when one moves along the edge away from the given vertex). Let h_j $(0 \leq j \leq s)$ denote the number of vertices of P for which there are exactly j leaving edges. The vector $(h_0, \ldots, h_s)$ is called the *h–vector*.

The f–vector can be determined from the h–vector, by the following system of linear equations.

$$f_k = \sum_{j=0}^{s} \binom{j}{k} h_j \quad \text{for } 0 \leq k \leq s.$$

These equations result from the fact that each k–dimensional face of P possesses a unique vertex at which the objective function is minimized, and, if a given vertex has exactly j "leaving" edges, so that it contributes to h_j, then exactly $\binom{j}{k}$ faces of dimension k achieve their minima at the vertex.

Since the objective function factors into the determination of the h–vector in the above, one might guess that there would be lots of different h–vectors, depending upon the objective function chosen. In fact, there is only one, as can be seen by noting that the above system of equations is triangular with 1's on the diagonal, and therefore invertible. Upon inversion, it is clear that the h–vector can be determined from the f–vector. Since the f–vector does not depend upon the choice of objective function, any objective function (for which the nondegeneracy assumption is satisfied) yields the same h–vector. In particular, if the objective function is replaced by its negative, one sees that the h–vector is symmetric: $h_j = h_{s-j}$ $(0 \leq j \leq s)$. When these equations are written in terms of the f–vector, they are known as the *Dehn–Somerville equations*.

The tableaux for our problem which have no negative entries in the last column correspond to the vertices of P. The number of out–going edges at the vertex is the number of positive entries in the last row. It follows that the h–vector is the first row of N. (It can be shown that all entries of the last row of N are zero, in this case.) Similarly, if the dual of our problem has a feasible region which is nonempty and bounded, then the h-vector for the polytope which is the dual feasible region appears as the first column of N (and the last column of N consists of 0's).

It is the symmetry of the h–vector noted above that is generalized in this paper for the entire matrix N (with a modification that removes the need for the condition that the feasible region be nonempty and bounded).

The h–vector figured prominently in McMullen's proof in [6] of the "Upper Bound Conjecture" concerning the maximum number of vertices that a polytope of given dimension and with a given number of facets might have, and in the proof, jointly by Billera and Lee [1] and Stanley [9], of the characterization of the f–vectors of the simple polytopes, conjectured by McMullen in [7]. Related counting problems for arrangements of hyperplanes and, more generally, for oriented matroids have been considered, for example, in [4], [5]. The "mutation count matrices" of [5] are related to, but not the same as, the matrices N considered here.

In what follows, A denotes an $r \times s$ matrix of real numbers; B is a column vector of length r; C is a row vector of length s; D is a real number; and M is the $(r+1) \times (s+1)$ composite block matrix

$$M = \begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

The matrix M, augmented by row and column labels indicating the variables associated with them (denoted by $x_1, \ldots, x_n$, where $n = r + s$), is the Tucker tableau for the linear programming problem (described by Goldman and Tucker, [3]):

$$\begin{aligned} \text{Maximize} \quad & CX_{(0)} - D \\ \text{subject to} \quad & AX_{(0)} \leq B, \\ & X_{(0)} \geq 0. \end{aligned}$$

Here $X_{(0)}$ denotes the column vector $X_{(0)} = [x_1, \ldots, x_s]^T$ of real variables.

The simplex method solves such a problem by performing "pivot" steps beginning with the matrix M and ending at a matrix from which the solution can be easily obtained. A *(Tucker) pivot* on the (nonzero) entry M_{i_0,j_0} of M consists of changing M by performing the following operations:

- Each entry $M_{i,j}$ not in the same row or column as M_{i_0,j_0} (so that $i \neq i_0, j \neq j_0$), is replaced by $M_{i,j} - M_{i,j_0} M_{i_0,j} / M_{i_0,j_0}$;
- Each entry M_{i,j_0} in the same column as the pivot entry M_{i_0,j_0} but different from it is replaced by $-M_{i,j_0}/M_{i_0,j_0}$;

- Each entry $M_{i_0,j}$ in the same row as the pivot entry but different from it is replaced by $M_{i_0,j}/M_{i_0,j_0}$;
- The entry itself is replaced by its inverse, $1/M_{i_0,j_0}$; and finally
- The labels of the row and column of the pivot element are exchanged.

For example, if

$$A = \begin{pmatrix} 1 & 2 & 3 \\ 6 & 3 & 1 \end{pmatrix}, \qquad B = \begin{pmatrix} 6 \\ 24 \end{pmatrix},$$

$$C = \begin{pmatrix} 1 & 1 & 1 \end{pmatrix}, \quad \text{and} \quad D = 0,$$

then we have the following Tucker tableau. (Here and later, T_{ijk}, where $1 \leq i < j < k \leq 5$, represents a tableau having x_i, x_j, x_k, and u as its column labels, not necessarily in that order.)

T_{123} :

	x_1	x_2	x_3	u
x_4	1	2	3	6
x_5	6	3	1	24
v	1	1	1	0

Upon performing a pivot on the entry whose row is labeled by x_5 and whose column is labeled by x_1, we get the following new tableau.

T_{235} :

	x_5	x_2	x_3	u
x_4	-1/6	3/2	17/6	2
x_1	1/6	1/2	1/6	4
v	-1/6	1/2	5/6	-4

The order of the rows is irrelevant, as is the order of the columns; the labels keep track of required information. This tableau is repeated again in the list below. Each tableau appears with rows and columns sorted according to their labels. Some labeling schemes include labels for dual variables. Here only primal variables are used as labels.

The simplex method performs Tucker pivots only for pivot elements which are in the portions of the matrices occupied by A (henceforth termed the "A part"). For the example, all the possible matrices obtainable from the initial tableau by sequences of such pivots are listed here:

T_{123} :

	x_1	x_2	x_3	u
x_4	1	2	3	6
x_5	6	3	1	24
v	1	1	1	0

T_{124} :

	x_1	x_2	x_4	u
x_3	1/3	2/3	1/3	2
x_5	17/3	7/3	-1/3	22
v	2/3	1/3	-1/3	-2

T_{125} :

	x_1	x_2	x_5	u
x_3	6	3	1	24
x_4	-17	-7	-3	-66
v	-5	-2	-1	-24

T_{134} :

	x_1	x_3	x_4	u
x_2	1/2	3/2	1/2	3
x_5	9/2	-7/2	-1/2	15
v	1/2	-1/2	-1/2	-3

T_{135} :

	x_1	x_3	x_5	u
x_2	1/3	2	1/3	8
x_4	-2/3	-3	7/3	-10
v	-1/3	-1	2/3	-8

T_{145} :

	x_1	x_4	x_5	u
x_2	17/7	-1/7	3/7	66/7
x_3	-9/7	3/7	-2/7	-30/7
v	-1/7	-2/7	-1/7	-36/7

T_{234} :

	x_2	x_3	x_4	u
x_1	2	3	1	6
x_5	-9	-17	-6	-12
v	-1	-2	-1	-6

T_{235} :

	x_2	x_3	x_5	u
x_1	1/2	1/6	1/6	4
x_4	3/2	17/6	-1/6	2
v	1/2	5/6	-1/6	-4

T_{245} :

	x_2	x_4	x_5	u
x_1	7/17	-1/17	3/17	66/17
x_3	9/17	6/17	-1/17	12/17
v	1/17	-5/17	-2/17	-78/17

T_{345} :

	x_3	x_4	x_5	u
x_1	-7/9	-1/3	2/9	10/3
x_2	17/9	2/3	-1/9	4/3
v	-1/9	-1/3	-1/9	-14/3

.

These are all the tableaux that can be obtained by pivoting in the A part, beginning from the original one of our example.

Let $X_{(1)} = [x_{1+s}, \ldots, x_{r+s}]^T$, and consider the vector space $\mathcal{W}$ consisting of all vectors

$$X = \begin{pmatrix} X_{(0)} \\ X_{(1)} \\ u \\ v \end{pmatrix} \in R^{r+s+2}$$

such that

$$X_{(1)} = uB - AX_{(0)}, \qquad v = -uD + CX_{(0)}.$$

The problem can now be written as one of maximizing v subject to $u = 1$, the nonnegativity of the x_i's, and $X \in \mathcal{W}$. In these equations, $X_{(1)}$ and v are determined linearly from $X_{(0)}$ and u; so $\mathcal{W}$ is the graph of a linear function from R^{s+1} to R^{r+1}. The tableau determines the linear function. The graph $\mathcal{W}$ is a linear subspace of $R^{r+s+2} = R^{n+2}$ having dimension $s+1$. Given another set of s of the x_i's, unless there is a linear relation relating them, it is possible to transform the equation system in such a way that v and the other r x_i's are determined linearly from those original s, and u. Indeed, Tucker pivoting gives the tableau of such a function. The subspace $\mathcal{W}$ is unchanged; we look at it differently, as the graph of a function of a different set of s variables.

There are only $\binom{n}{s}$ sets of s of the n variables, so there cannot be more than this number of Tucker tableaux. In the nondegenerate case of the example, we do indeed have all $\binom{5}{2} = 10$ distinct tableaux. There is sufficient nondegeneracy that no tableau for the problem has a zero entry in the A, B, or C part. We say that the problem exhibits *total nondegeneracy.* We refer to the set of tableaux in our list as the *complete set of tableaux* for the problem. In general, given an initial tableau for a problem, the complete set of tableaux for the problem consists of all the tableaux which can be derived from the initial one by sequences of pivots on elements in the A part of the tableaux.

By a basic theorem of linear programming applied to the totally nondegenerate case, exactly one of the following three possibilities holds for the complete set of tableaux for the problem:

- There is a unique tableau having only positive entries in the B part and only negative entries in the C part, from which the optimal solution can be obtained (primal and dual feasibility);
- There is a tableau having only positive entries in the B part and having another column for which the entry in the C part is positive and all other entries negative (primal feasibility, dual infeasibility);
- There is a tableau having only negative entries in the C part and having another row for which the entry in the B part is negative and all other entries are positive (primal infeasibility, dual feasibility).

It is not possible for the totally nondegenerate problem to be both primal and dual infeasible.

Let N be the $(r+1)\times(s+1)$ matrix of nonnegative integers whose (i,j)-th entry is the number of tableaux in the complete set which have exactly i negative elements in the B part of the tableaux and exactly j positive elements in the C part. Then $N(0,0)$ is either 0 or 1; it is 1 (in the presence, as we assume, of total nondegeneracy) if and only if our problem has a solution – if and only if the feasible set is nonempty and the objective function is bounded above on it. The *solution tableau* is the one, if it exists, having 0 negative elements in the B part and 0 positive elements in the C part. Let N' be the $(r+1)\times(s+1)$ matrix whose (i,j)-th entry is the number of tableaux in the complete set which have exactly i positive elements in the B part of the tableaux and exactly j negative elements in the C part. For our example, we have

$$N = \begin{pmatrix} 1 & 2 & 2 & 1 \\ 3 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \text{ and } N' = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 3 \\ 1 & 2 & 2 & 1 \end{pmatrix}.$$

Thanks to total nondegeneracy, there is no 0 entry in the B or C part of any tableau, so, for $0 \le i \le r$ and $0 \le j \le s$, $N'_{i,j} = N_{r-i,s-j}$. Therefore, the sum $Q = N + N'$ has some symmetry: $Q_{i,j} = Q_{r-i,s-j}$, for all i, j.

The sum is:

$$Q = \begin{pmatrix} 1 & 2 & 2 & 1 \\ 3 & 1 & 1 & 3 \\ 1 & 2 & 2 & 1 \end{pmatrix},$$

which actually has even more symmetry: $Q_{i,j} = Q_{r-i,j} = Q_{i,s-j} = Q_{r-i,s-j}$, for all i, j. Our objective is to show that this is true, for all totally nondegenerate linear programming problems.

Mulmuley, in [8], has considered a matrix which is a submatrix of N, in a more special case. Let k be a nonnegative integer, and let X_k denote the set of points satisfying at least $n-k$ of the inequality constraints (so that X_0 is the feasible region). Under the assumptions that (i) X_k is bounded and (ii) X_0 is nonempty, Mulmuley's "h–matrix" is the submatrix of N consisting of its first $k+1$ rows, and he has shown that N satisfies the equations $N(i,j) = N(i,s-j)$ when $0 \le i \le k$ and $0 \le j \le s$. It is not difficult to show that, when (i) and (ii) are satisfied, the last $k+1$ rows of N consist of 0's, so that the first $k+1$ rows of N and Q coincide.

2 Derived Problems

There are several operations which can be performed on a tableau which lead to different problems, having different solutions, but for which the complete set of tableaux for the original problem provide sufficient information. We consider some of these, one-by-one.

Suppose that the row or column of the tableau which is indexed by x_i is multiplied by -1. If it is a row, the new tableau corresponds to a problem identical to the original, except that the sense of the inequality constraint corresponding to that row is reversed; if it is a column, then the variable labeling that column is required to be nonpositive instead of nonnegative. This operation commutes with the operation of performing a pivot on the tableau: If we multiply the row or column indexed by x_i by -1 and then perform a pivot on the entry in row and column labeled by x_j and x_k, the resulting tableau is the same as the tableau that results when we first perform the pivot and then perform the multiplication by -1. This being the case, it becomes clear that the complete set of tableaux for the new problem can be obtained by replacing each tableau for the original problem by the result upon multiplying its x_i–labeled row or column by -1.

Suppose next that x_i is a row label in our tableau and that row is deleted. This corresponds to removing the corresponding inequality from the constraint set. Pivoting again commutes with this operation, as long as we do not pivot in the row labeled x_i. The complete set of tableaux for the new problem is obtained by replacing each tableaux for the original, having the property that x_i labels a row, by the result upon deleting that row. The tableaux in which x_i labels a column are discarded.

Suppose that x_i is a column label in the original tableau, and that that column is deleted. This amounts to replacing the original problem by one in which x_i is set to zero: $x_i = 0$. The complete set of tableaux for the new problem is obtained by removing the column labeled x_i from each original tableau which has a column labeled x_i. The tableaux in which x_i labels a row are discarded.

Of course, linear programming duality amounts to replacing the original tableau by the negative of its transpose; and this also commutes with the pivot operations.

We term the problems obtained by the operations of multiplication of rows, columns by -1, deletion of rows, columns, the *derived problems* of the original. The duals of the derived problems are the derived problems of the dual. For each derived problem, the complete set of tableaux can be obtained from those of the original problem, as above. If P, Q, and R are pairwise disjoint sets of variables x_i, with $|P| \leq r$ and $|Q| \leq s$, then $\mathcal{D}(P,Q,R)$ is the derived problem for which the complete set of tableaux is obtained by starting with those tableaux in the original set for which the variables in P appear as row labels and those in Q appear among the column labels, and deleting from them the rows labeled by variables in P, deleting the columns labeled by variables in Q, and multiplying rows and columns labeled by elements of R by -1. (We leave it to the reader to sort out the meaning of those derived problems $\mathcal{D}(P,Q,R)$ for which $|P| = r$ or $|Q| = s$. These cases are useful in some of the proofs below.)

We consider the number $\mathcal{F}(p,q,t)$ which is the count of the derived problems $\mathcal{D}(P,Q,R)$ which are feasible and such that the objective function is bounded above on the feasible region, and having $|P| = p$, $|Q| = q$, and $|R| = t$. We call p, q, and t the *parameters* of the derived problem $\mathcal{D}(P,Q,R)$.

Theorem 1. *The number $\mathcal{F}(p,q,t)$ is given by a linear combination of the entries of the matrix N:*

$$\mathcal{F}(p,q,t) = \sum_{\substack{0\leq i\leq r,\\ 0\leq j\leq s}} \gamma(p,q,t;i,j)N(i,j).$$

The coefficients $\gamma(p,q,t;i,j)$ are nonnegative integers.

Proof. Given a tableau T, it is easy to see that the number of derived problems having parameters p, q, and t which have T as the solution tableau is a function only of the numbers i and j which count the negative entries in the B part and positive entries in the C part of T; $\gamma(p,q,t;i,j)$ is this number. Since each derived problem which has a nonempty feasible set on which the objective function is bounded above has a unique solution tableau, the sum yields the total number. □

Suppose T is a tableau, and suppose its B part has i negative entries and its C part has j positive entries, so that it contributes to the number $N(i,j)$. How many derived problems $\mathcal{D}(P,Q,R)$ having parameters p, q, and t have T as the solution tableau? We may construct them all, as follows. Let k and ℓ be nonnegative integers whose sum is t. Choose a set R_1 of k variables from among the i variables labeling the rows of T which have a negative entry in the B part; there are $\binom{i}{k}$ ways to do this. Next choose a set R_2 of ℓ variables from among the j variables labeling the columns of T which have a positive entry in the C part. There are $\binom{j}{\ell}$ ways to do this. Let $R = R_1 \cup R_2$. From among the row labels, choose a set P having p variables. Include in P all the variables not in R_1 which label a row having a negative entry in the B part; choose the remaining $p-i+k$ elements of P from among the $r-i$ row labels having a positive entry in the B part. There are $\binom{r-i}{p-i+k}$ ways to do this. Finally, choose a set Q having q elements, including the labels of the $j-\ell$ columns having positive elements in the C part, with the other $q-j+\ell$ elements chosen from the $s-j$ labels of columns having negative elements in the C part. There are $\binom{s-j}{q-j+\ell}$ ways to choose Q in this fashion. It follows that

$$\gamma(p,q,t;i,j) = \sum_{\substack{k+\ell=t,\\ k,\ell\geq 0}} \binom{i}{k}\binom{j}{\ell}\binom{r-i}{p-i+k}\binom{s-j}{q-j+\ell}.$$

There are usually many more values of $\mathcal{F}(p,q,t)$ than entries in the matrix N, so there are linear relations among the values $\mathcal{F}(p,q,t)$.

The next theorem shows that N can be obtained from those values of $\mathcal{F}$ for which: $q=0$, $0\leq p\leq r$, and $0\leq t\leq s$. For such p, q, and r, we have:

$$\gamma(p,q,t;i,j) = \binom{i}{t-j}\binom{r-i}{p-i+t-j}$$

and

$$\mathcal{F}(p,0,t) = \sum_{\substack{0\leq i\leq r,\\ 1\leq j\leq s}} \binom{i}{t-j}\binom{r-i}{p-i+t-j} N(i,j).$$

Theorem 2. *In Theorem 1, upon restricting $p,q,$ and t to satisfy $q=0$, $0\leq p\leq r$, and $0\leq t\leq s$, and providing a suitable ordering, we obtain a triangular system of linear equations having 1's on the diagonal. Thus it is invertible, so that, given arbitrarily the values of the function $\mathcal{F}$ for p, q, and t so restricted, there is a unique matrix N for which the equations of Theorem 1 are satisfied.*

Proof. We linearly order the set of pairs (a, b) of integers: $(a, b) \prec (a', b')$ if $b < b'$, or $b = b'$ and $a < a'$. Consider $(p, t) = (i, j)$; then, in the expression yielding $\mathcal{F}(p, 0, t)$ above, the coefficient of $N(i, j)$ is 1. Suppose $(p, t) \prec (i, j)$; then this coefficient is 0. $\square$

3 The Proof

Two lemmas will be useful in the proof of the main result.

Let Π denote the problem corresponding to the original tableau, T. Let Π' denote the problem corresponding to the result upon multiplying the column of T labeled u by -1; let Π'' denote the problem corresponding to the result upon multiplying the row of T labeled v by -1; and let Π''' denote the problem corresponding to the result upon performing both of those operations. Since the pivot operations commute with these operations, it is clear that, from the complete set of tableaux for Π, we obtain a complete set for each of the other problems by performing the same operation on all of the tableaux in the complete set Π.

For the corresponding derived problem $\mathcal{D}(P, Q, R)$, function $\mathcal{F}$, and matrix N we shall now write $\mathcal{D}_\Pi(P, Q, R)$, $\mathcal{F}_\Pi$, and N_Π; and similarly for the other three problems.

Given the sets P, Q, R of variables, pairwise disjoint, let

$$\iota_\Pi(P, Q, R) = \begin{cases} 1 & \text{if } \mathcal{D}_\Pi(P, Q, R) \text{ is feasible and dual-feasible} \\ 0 & \text{otherwise.} \end{cases}$$

The functions $\iota_{\Pi'}$, $\iota_{\Pi''}$, and $\iota_{\Pi'''}$ are similarly defined.

Lemma 1. *For each choice of P, Q, and R, we have*

$$\iota_\Pi(P, Q, R) + \iota_{\Pi'''}(P, Q, R) = \iota_{\Pi'}(P, Q, R) + \iota_{\Pi''}(P, Q, R).$$

Proof. We make the argument in the case $P = Q = R = \emptyset$; however it will be clear that the same argument applies in general. Note that $\mathcal{D}(\emptyset, \emptyset, \emptyset)$ is simply the original problem, Π.

We designate the following statement by $B+$: There is a tableau in the complete set of tableau for Π which has only positive elements in the B part. Similarly define statements, $B-$, $C+$, and $C-$.

Either $B+$ and $B-$ both hold, or $C+$ and $C-$ both hold; for otherwise one of the four problems Π, Π', Π'', or Π''' is neither primal feasible nor dual feasible. Then it is easy to see that, depending upon whether neither, exactly one, or both of the other two statements hold, both sides in the above equation must be 0, 1, or 2. (Actually, 2 is not possible; but this fact is not needed here.) $\square$

Lemma 2. *We have*

$$\mathcal{F}_{\Pi}(p,q,t) + \mathcal{F}_{\Pi'''}(p,q,t) = \mathcal{F}_{\Pi'}(p,q,t) + \mathcal{F}_{\Pi''}(p,q,t).$$

Proof. Observe that

$$\mathcal{F}_{\Pi}(p,q,t) = \sum_{(P,Q,R)} \iota_{\Pi}(P,Q,R),$$

where the summation extends over triples (P,Q,R) of pairwise disjoint sets of variables for which $|P| = p$, $|Q| = q$, and $|R| = t$. The same holds for the other three problems, so this lemma is a consequence of the preceding lemma. □

Theorem 3. *Letting $Q(i,j) = N(i,j) + N(r-i, s-j)$, we have $Q(i,j) = Q(r-i,j) = Q(i,s-j) = Q(r-i,s-j)$.*

Proof. Equivalently, we must show that $N_{\Pi}(i,j) + N_{\Pi}(r-i,s-j) = N_{\Pi}(r-i,j) + N_{\Pi}(i,s-j)$. It is clear that:

$$N_{\Pi'}(i,j) = N_{\Pi}(r-i,j),$$

$$N_{\Pi''}(i,j) = N_{\Pi}(i,s-j),$$

and

$$N_{\Pi'''}(i,j) = N_{\Pi}(r-i,s-j).$$

The equations of Theorem 1 also give

$$\mathcal{F}_{\Pi}(p,q,t) + \mathcal{F}_{\Pi'''}(p,q,t) - \mathcal{F}_{\Pi'}(p,q,t) - \mathcal{F}_{\Pi''}(p,q,t)$$

in terms of

$$N_{\Pi}(i,j) + N_{\Pi'''}(i,j) - N_{\Pi'}(i,j) - N_{\Pi''}(i,j).$$

However, by Lemma 2,

$$\mathcal{F}_{\Pi}(p,q,t) + \mathcal{F}_{\Pi'''}(p,q,t) - \mathcal{F}_{\Pi'}(p,q,t) - \mathcal{F}_{\Pi''}(p,q,t) = 0,$$

so by Theorem 2 we have

$$N_{\Pi}(i,j) + N_{\Pi'''}(i,j) - N_{\Pi'}(i,j) - N_{\Pi''}(i,j) = 0,$$

or equivalently,

$$N_{\Pi}(i,j) + N_{\Pi}(r-i,s-j) = N_{\Pi}(r-i,j) + N_{\Pi}(i,s-j).$$

□

4 Notes

An interesting open problem is that of characterizing the counting matrices N. Certainly, these matrices satisfy the following conditions: The entries are nonnegative integers; the sum of the entries is $\binom{n}{s}$; and, for the matrix Q having $Q(i,j) = N(i,j)+N(r-i,s-j)$, one has $Q(i,j) = Q(r-i,j) = Q(i,s-j) = Q(r-i,s-j)$. These conditions are far from sufficient to characterize the possible matrices.

The results in this paper hold more generally for uniform oriented matroids. With minor changes, the proofs hold in the more general setting. It is unknown whether or not one obtains additional matrices N from the "linear programming problems" derived in the setting of uniform oriented matroids.

The problem of characterizing the analogous matrices N, when total nondegeneracy is not assumed, is also open.

References

1. L. J. Billera and C. W. Lee, A proof of the sufficiency of McMullen's conditions for f-vectors of simplicial convex polytopes. *J. Combinat. Theory, Ser. A* **31** (1981), 237–255.
2. S. Gass, *Linear Programming Methods and Applications*. McGraw-Hill, 1958. (Fourth edition, 1975.)
3. A. J. Goldman and A. W. Tucker, Theory of linear programming. In *Linear Inequalities and Related Systems*, *Annals of Mathematics Studies* **38**, eds. H. W. Kuhn and A. W. Tucker, Princeton University Press, 1956.
4. J. Lawrence, Total polynomials of uniform oriented matroids. *European Journal of Combinatorics*, **21** (2000), 3–12.
5. J. Lawrence, Mutation polynomials and oriented matroids. *Discrete and Comput. Geom.*, **24** (2000), 365–389.
6. P. McMullen, The maximum numbers of faces of a convex polytope. *Mathematika* **17** (1970), 179-184.
7. P. McMullen, The numbers of faces of simplicial polytopes. *Israel J. Math.* **9** (1971), 559–570.
8. K. Mulmuley, Dehn–Sommerville relations, upper bound theorem, and levels in arrangements. *Proceedings of Ninth Annual Symposium on Computational Geometry: San Diego, California, May 12–21, 1993*, 240–246.
9. R. P. Stanley, The numbers of faces of a simplicial convex polytope. *Adv. Math.* **35** (1980), 236–238.