An observation on the structure of production sets with indivisibilities

(integer programming/complexity/fixed point methods)

HERBERT E. SCARF

Cowles Foundation for Research in Economics at Yale University, Box 2125, Yale Station, New Haven, Connecticut 06520

Contributed by Herbert E. Scarf, May 25, 1977

ABSTRACT   A subset of the constraints of an integer programming problem is said to be binding if, when the remaining constraints are eliminated, the smaller problem has the same optimal solution as the original problem. It is shown that an integer programming problem with n variables has a set of binding constraints of cardinality less than or equal to $2^n - 1$. The bound is sharp.

1. Introduction

The purpose of the present paper is to illustrate a new point of view in the discussion of discrete programming problems by demonstrating a theorem concerning the maximum number of binding constraints in an integer programming problem with n variables. The arguments are a blend of those used in fixed point computations and in the geometry of numbers, without the symmetry assumptions which are characteristic of the latter field. Subsequent papers will discuss the ramifications of this point of view in greater detail.

Consider an integer programming problem of the following form:

1.1
$$max\ a_{01}h_1 + \ldots + a_{0n}h_n$$

$$a_{11}h_1 + \ldots + a_{1n}h_n \geq b_1$$

$$\vdots$$

$$a_{m1}h_1 + \ldots + a_{mn}h_n \geq b_m,$$

with the variables $h_1, \ldots, h_n$ restricted to integral values. Any requirement that some or all of the variables be non-negative will be incorporated in the constraints, so that typically the number of inequalities, m, will be greater than or equal to the number of variables, n. The following assumptions will be made throughout the discussion.

1.2. ASSUMPTIONS. We assume that for any vector c in R^{m+1} the set of integral h satisfying $Ah \geq c$ is finite. In addition the constraints are assumed to have at least one feasible integral solution for the given right hand side b.

It is an immediate consequence of these assumptions that problem 1.1 has a finite maximum.

1.3. DEFINITION. Let S be a subset of $(1, 2, \ldots, m)$. The constraints $\Sigma_j a_{ij}h_j \geq b_i$, for $i \in S$, are said to be *binding* if the integer programming problem

$$max \sum_j a_{0j}h_j$$

$$\sum_j a_{ij}h_j \geq b_i \quad \text{for } i \text{ in } S,$$

has an optimal solution in which the remaining constraints are all satisfied, and which is therefore optimal for the original problem.

The primary result of the present paper is the following theorem.

1.4. THEOREM. The integer programming problem 1.1 has a set of binding constraints whose cardinality is less than or equal to $2^n - 1$. Moreover, this bound is sharp in the sense that there exist integer programming problems with n variables that have no sets of binding constraints of cardinality less than $2^n - 1$.

The situation described in Theorem 1.4 is distressing when compared to that arising in ordinary linear programming in which a set of binding constraints of cardinality n can always be found. The simplex method may, in fact, be viewed as a systematic algorithm that searches through sets of n constraints in order to determine whether or not they are binding and concludes by exhibiting a specific set of such constraints. Theorem 1.4 will also be demonstrated by means of an algorithm— whose properties will be explored in subsequent papers—that terminates with the appropriate set of binding constraints. The theorem indicates the intrinsic complexity of some programming problems imposed by the requirement that the variables take on integral values.

2. Production set associated with the integer programming problem

The study of linear programming problems is greatly facilitated by introducing an associated production possibility set, i.e., the convex cone of production plans obtained by varying the activity levels in an arbitrary fashion. We shall introduce a similar construction for integer programming.

2.1. DEFINITION. We define X to consist of the set of vectors $(x_0, \ldots, x_m)$ given by

$$x_0 = \sum_j a_{0j}h_j$$

$$\vdots$$

$$x_m = \sum_j a_{mj}h_j,$$

as $h_1, \ldots, h_n$ vary over the integers.

In order to avoid certain technical difficulties the following regularity assumption will be made about the matrix A and then removed in *Section 6* of the paper.

2.2. ASSUMPTION. The entries in A are independent over the integers in the sense that if h is an integral vector other than the zero vector, $\Sigma_j a_{ij}h_j \neq 0$ for each row i.

The set X may be shown to lie on an m-dimensional hyperplane with a non-negative normal. Any bounded subset of the hyperplane contains a finite number of vectors in X, and no two

FIG. 1.   A finite subset of the vectors in X for $m = n = 2$.

FIG. 2.   A finite subset of vectors in X in which n is small compared to m.

vectors in X have their ith coordinates equal for any coordinate.

Fig. 1 represents a finite subset of the vectors in X for $m = n = 2$. The set consists of seven vectors $x^0, \ldots, x^6$. I have drawn through each of these vectors that translate of the *negative* orthant in R^{m+1} having its vertex at the particular vector in question. This provides us with an intuitive picture of the "upper surface" of this portion of the production possibility set X.

We shall now introduce an infinite simplicial complex C, whose vertices are the vectors in X. The first step is to define those collections of $m + 1$ vertices $(x^{j0}, x^{j1}, \ldots, x^{jm})$ forming the m-simplices of the complex. We begin by translating the *positive* orthant in R^{m+1} parallel to itself until its vertex lies above the plane containing X. We then translate the orthant downwards, passing through no vectors in X, until no further reduction in any of the coordinates of this vertex is possible. The orthant will typically be stopped by a collection of $(m + 1)$ vectors $x^{j0}, x^{j1}, \ldots, x^{jm}$ one on each coordinate hyperplane of the translated orthant. (Assumption 2.2 implies that a coordinate hyperplane of the translated orthant cannot contain more than one vector in X.) These collections of $(m + 1)$ vectors, which have elsewhere been given the name of *primitive* sets, will be the m-simplices in C.

In Fig. 1 the vectors (x^0, x^2, x^3) form a primitive set with the corresponding vertex at α; (x^0, x^1, x^3) a primitive set with the vertex at β. Of course $\alpha = min[x^0, x^2, x^3]$, in each coordinate, and as a consequence of the construction there is no vector x in X with $x > \alpha$.

We have the following general definition.

2.3. DEFINITION. A set of vectors $(x^{j0}, x^{j1}, \ldots, x^{jm})$ in X is defined to be a *primitive* set if there is no vector x in X with

$$x > min[x^{j0}, x^{j1}, \ldots, x^{jm}].$$

It should be remarked that the question of whether a given set of $(m + 1)$ vectors $x^{j0}, x^{j1}, \ldots, x^{jm}$ forms a primitive set is a local question in the sense that there are a finite number of vectors in X contained in any specific translate of the positive orthant.

If n is sufficiently small compared to m, then it is possible that no collections of $(m + 1)$ vectors satisfy the Definition 2.3. In Fig. 2 either the first or second coordinates of the vertex of the translated positive orthant can be decreased indefinitely without passing through a vector in X.

In order to overcome this difficulty we shall adjoin to the set X, $(m + 1)$ "ideal" vectors $\xi^0, \xi^1, \ldots, \xi^m$, which we shall refer

to as *slack vectors* because of an analogy with linear programming. The ith of these vectors, ξ^i, will be assumed to have its ith coordinate very negative, and its remaining coordinates very positive. Definition 2.3 is applied to this extended set and primitive sets will now consist of $(m + 1)$ vectors, some of which come from the original set X, and the remainder of which are slack vectors. In Fig. 2, x^1, x^2, and ξ^1 form a primitive set, as do x^1, x^2, and ξ^2.

Consider a set of vectors $x^{j0}, x^{j1}, \ldots, x^{jm}$ arranged in such a way that

$$\alpha_i = min[x_i^{j0}, x_i^{j1}, \ldots, x_i^{jm}] = x_i^{ji}$$

for $i = 0, 1, \ldots, m$. Let $\{x^{ji}\}$ for $i \in T$ be the *non-slack* vectors in this set. In order to test whether this set of vectors forms a primitive set it is sufficient to verify that there is no vector in X with

$$x_i > \alpha_i \text{ for } i \in T.$$

The simplicial complex C obtained by taking the collection of primitive sets as m-simplices and all proper subsets as lower dimensional simplices is a basic tool in the analysis of production sets with indivisibilities and will be discussed in detail in subsequent publications. As an example of the relevance of this construction I quote the following theorem, whose proof will not be given here, because the theorem will not be used in developing the arguments of the present paper.

2.4. THEOREM. Let $x \in X$ and define the neighborhood of x to be the set of vectors in X that are contained in some primitive set that also contains x. Let x satisfy the constraints of the integer program

$$x_i \geq b_i \text{ for } i = 1, \ldots, m.$$

Then a necessary and sufficient condition that x be an optimal solution to the programming problem is that x be a local optimum, in the sense that every neighbor of x either violates one of the constraints or yields a smaller value of the objective.

3. Sperner's lemma

In this section we consider an arbitrary finite subset of X, say X', augmented as before by the $(m + 1)$ slack vectors $\xi^0, \xi^1, \ldots, \xi^m$. Primitive sets are defined to be those collections of $(m + 1)$ vectors $x^{j0}, \ldots, x^{jm}$ (either slack vectors or members of X') such that there is no vector in X' satisfying

$$x > min[x^{j0}, x^{j1}, \ldots, x^{jm}].$$

If, for example, X' consists of the seven vectors illustrated by Fig. 1, then the following triples are included among the

Mathematics and Economics: Scarf

Proc. Natl. Acad. Sci. USA 74 (1977)   3639

FIG. 3.   Finding the replacement for x^1.

primitive sets: (x^0, x^2, x^3), (x^0, x^1, x^3), (x^0, x^1, ξ^1), (x^0, ξ^1, ξ^2), (x^1, x^4, ξ^1).

The following theorem, whose proof may be found in ref. 1, is a version of Sperner's lemma appropriate to the present setting.

3.1.   THEOREM. *For each $x \in X'$ let $l(x)$ be an integer in the set $(0, 1, \ldots, m)$, and let $l(\xi^i) = i$ for $i = 0, 1, \ldots, m$. Then there exists a primitive set $x^{j_0}, \ldots, x^{j_m}$ (including possibly some slack vectors) whose labels $l(x^{j_0}), \ldots, l(x^{j_m})$ are all different.*

Theorem 3.1 is a fundamental tool in numerical techniques for the approximation of fixed points of a continuous mapping, when the underlying combinatorial objects are given by primitive sets rather than a simplicial decomposition of the simplex. The proof of Theorem 3.1 depends on the following lemma, which is illustrated in Fig. 3.

3.2.   LEMMA. *Let $x^{j_0}, \ldots, x^{j_m}$ be a primitive set (including possibly some slack vectors). Consider an $(m-1)$-dimensional face of this simplex obtained by deleting a specific vector in the primitive set. Then, aside from the special case in which the m vectors on the face are all slack vectors, there is a* unique *other primitive set with the same $(m-1)$-dimensional face. In the special case there is* no *other primitive set with this property.*

In Fig. 3 (x^0, x^1, x^2) form a primitive set with vertex α. We have named the vectors so that $\alpha_i = x_i^i$ for $i = 0, 1, 2$. To remove x^1, we locate the vector in the primitive set with the *second* smallest first coordinate; in this case x^2. We then locate that unique vector in X' (or possibly a slack vector) with the largest second coordinate, subject to $x_0 > x_0^0$, $x_1 > x_1^2$; in this case x^3.

The special case referred to in Lemma 3.2 occurs when the primitive set consists of m slack vectors—say all of the slack vectors except ξ^{i*}—and that particular vector in X' with the largest $i*$th coordinate. The lemma may be rephrased as saying that every $m - 1$ face in the complex appears in precisely *two* m-simplices, excepting those $m - 1$ faces consisting entirely of slacks; these appear in precisely one m-simplex.

The algorithm for determining a simplex with distinct labels begins with one of the special simplices—say the m-simplex $(x^*, \xi^1, \ldots, \xi^m)$ with x^* that vector in X' with largest 0th coordinate. If $l(x^*) = 0$, the algorithm terminates, because $l(\xi^i) = i$ for $i = 1, \ldots, m$. If $l(x^*) \neq 0$, then we remove that unique vector in the primitive set, other than x^*, whose label agrees with $l(x^*)$, and continue.

At each iteration of the algorithm, prior to termination, we are faced with a primitive set whose vectors contain all of the labels in the set $(1, 2, \ldots, m)$. A single pair of vectors, one of which has just been intorduced into the primitive set, have the same labels. The algorithm continues by removing the other

member of this pair. A simple graph theoretic argument demonstrates that the algorithm never returns to a primitive set previously encountered. Because X' is a finite set, the algorithm must terminate after a finite number of iterations by introducing a vector with the label 0 resulting in an m-simplex with distinct labels.

4.   Labeling procedure for integer programming

In this section we use a particular labeling procedure for the vectors in X'. When a completely labeled primitive set is found by the arguments of the preceding section, that vector in the primitive set with the label 0 will solve the problem

$$max \sum a_{0j} h_j$$
$$\sum a_{1j} h_j \geq b_1$$
$$\vdots$$
$$\sum a_{mj} h_j \geq b_m,$$

and also $x \in X'$. If the vector with the label 0 is ξ^0, then there will be no vectors in X' satisfying the constraints.

4.1.   LABELING RULE. Let $x \in X'$. We define $l(x) = i$ ($i = 1, 2, \ldots, m$), if i is the first coordinate for which $x_i < b_i$. If $x_i \geq b_i$ for all $i = 1, 2, \ldots, m$, then $l(x) = 0$.

Let $x^0, x^1, \ldots, x^m$ be the vectors in a completely labeled primitive set (including possibly some slack vectors) arranged in such a way that $min[x_i^0, x_i^1, \ldots, x_i^m] = x_i^i$. I claim that because of the particular labeling rule being used we must have $l(x^i) = i$.

This is certainly true for $i = 1$, because $x_1^1 \leq x_1^0, x_1^2, \ldots, x_1^m$. If the vector x^1 did not receive the label 1, we would have $x_1^1 \geq b_1$ and therefore $x_1^j \geq b_1$ for all j, implying that no vector receives the label 1. We see, therefore, that $x_1^1 < b_1$. Moreover, $x_1^j \geq b_1$ for all $j \neq 1$, because otherwise two vectors would have the label 1. But then x^2 must receive the label 2, because if it did not we would have $x_2^2 \geq b_2$ and therefore $x_2^j \geq b_2$ for all j, implying that no vector receives the label 2. We see, as before, that $x_2^2 < b_2$ and $x_2^j \geq b_2$ for $j \neq 1, 2$.

The argument, when continued, verifies that $l(x^j) = j$ for $j = 1, \ldots, m$, and therefore $l(x^0) = 0$. It follows that x^0, if it is not the 0th slack vector, must satisfy the constraints $x_i^0 \geq b_i$ for $i = 1, \ldots, m$. But there can be no other vector $x \in X'$ that satisfies the constraints and whose 0th coordinate is larger than that of x^0, for such a vector would have the property that

$$x_0 > x_0^0 = min[x_0^0, x_0^1, \ldots, x_0^m],$$

and

$$x_i \geq b_i > min[x_i^0, x_i^1, \ldots, x_i^m] \quad \text{for } i = 1, \ldots, m.$$

violating the definition of a primitive set. It is also easy to verify that if x^0 is the 0th slack vector, then no member of X' satisifes the constraints $x_i \geq b_i$.

Assuming that the constraints are feasible for some x in X', we define S to be the set of indices i (other than 0) for which x^i is *not* a slack vector. It follows from the definition of primitive sets with slack vectors that x^0 is the vector in X' with the largest 0th coordinate subject only to the constraints

$$x_i \geq b_i \quad \text{for } i \text{ in } S.$$

In other words, the constraints associated with slack vectors are not binding and can be discarded without introducing a new optimal solution.

The arguments of this section have been conducted in terms of an arbitrary finite subset of X. We now take a more specific form by selecting a large positive number M, and letting X' be

the set of vectors in X with $x_i \geq -M$ for $i = 0, 1, \ldots, m$. If M is sufficiently large, X' will contain the optimal solution, say x^*, of our original integer programming problem 1.1. The primitive set exhibiting this fact will have a set of indices S, corresponding to the non-slack vectors in this primitive set (other than x^*). The activity levels associated with x^* will therefore be the optimal solution to

$$max \ \sum a_{0j}h_j$$
$$\sum a_{ij}h_j \geq b_i \quad \text{for } i \in S$$
$$\sum a_{ij}h_j \geq -M \quad \text{for } i \in S^c.$$

If M tends to infinity through a sequence of values for which the index set S remains the same, we conclude that the constraints associated with S form a binding set of constraints for 1.1. Our main theorem will therefore be demonstrated if we can show that there are never more than 2^n non-slack vectors (this includes x^*) in a primitive set when X' is given by those vectors in X all of whose coordinates are greater than or equal to $-M$.

5. Maximum number of non-slack vectors in a primitive set

Let $x^0, x^1, \ldots, x^m$ be a primitive set based on vectors $x \in X$ with $x_i \geq -M$ for $i = 0, 1, \ldots, m$, and arranged so that

$$x_i{}^i = min[x_i{}^0, \ldots, x_i{}^m]$$

for $i = 0, 1, \ldots, m$. Because of Assumption 2.2 we have $x_i{}^j > x_i{}^i$ for $j \neq i$. Let T be that set of indices i in $(0, 1, \ldots, m)$ for which x^i is *not* a slack vector.

For each $i \in T$ we have $x^i = Ah^i$, with h^i a vector in R^n with integral coordinates. We define H to be the convex hull of h^i, for $i \in T$.

We shall demonstrate that each h^i is an extreme point of the convex hull H, and moreover that there are no other lattice points in H. If either of these situations were to occur there would be a lattice point h (possibly one of the h^i themselves) with

$$h = \sum \alpha_j h^j, \ \sum \alpha_j = 1, \ \alpha_j \geq 0,$$

and at least two α_j strictly positive. But then $x = Ah$ is in X' and satisfies $x_i = \sum \alpha_j x_i{}^j > x_i{}^i$ for $i \in T$ and trivially $x_i > x_i{}^i$ for $i \in T^c$. This contradicts the definition of a primitive set.

The major conclusion of this paper—that part of Theorem 1.3 asserting the existence of a set of binding constraints of cardinality $(2^n - 1)$—is an immediate consequence of the following result, which appears in a letter to Roger Howe from J. W. S. Cassels.

5.1. THEOREM. Let $\{h^j\}$ be a set of distinct lattice points in R^n, which are extreme points in their convex hull H, and such that H contains no other lattice points. Then the number of points in $\{h^j\}$ is less than or equal to 2^n.

To prove 5.1, Cassels demonstrates by induction on n, that any set of more than 2^n distinct lattice points in R^n must contain a pair of lattice points h and h' with

$$h_i \equiv h_i' \ mod(2).$$

But then $(h + h')/2$ is a lattice point in H which differs from the vectors generating H. This demonstrates the existence of a set of binding constraints of cardinality $2^n - 1$.

Exhibiting an integer program with no set of binding constraints of cardinality less than $2^n - 1$ is, of course, trivial. We simply take the unit cube in R^n with vertices $v^1, \ldots, v^{2^n}$ (see Fig. 4). Through each such vertex we pass a hyperplane with


FIG. 4. A problem for which all constraints are binding.

normal a^i that strictly separates that vertex from the rest of the cube, i.e.,

$$a^i \cdot v^k > a^i \cdot v^i \quad \text{for } i \neq k.$$

The integer program

$$max \ a^1 \cdot h$$
$$a^i \cdot h > a^i \cdot v^i \quad \text{for } i \neq 1,$$

will have v^1 as its solution, but if any particular constraint is dropped, a superior solution will be obtained.

6. Removing Assumption 2.2

Assumption 2.2 guarantees that no two vectors in X have an identical first, second, $\ldots$, or mth coordinate. This is a nondegeneracy assumption which plays two distinct roles in our arguments. First of all, there is no ambiguity about which set of $(m + 1)$ vectors impedes the downward translation of the positive orthant in the definition of primitive sets because no two vectors in X lie on the same translated coordinate hyperplane. Moreover, there will always be a unique replacement for a given vector in a primitive set because this involves finding the vector in X that maximizes a specific coordinate subject to inequalities on the remaining coordinates.

This technical difficulty can easily be overcome if Assumption 2.2 does not hold—as will be the case when the matrix A is composed entirely of integers. We simply introduce some systematic rule that permits us to decide, when $x_i = x_i'$ for two different vectors, which of these two numbers is to be considered the larger. For example, we can say that in the event of a tie x_i will be considered to be larger than x_i' if and only if the vector x is lexicographically larger than x'.

If we adopt this rule, the only argument that requires some reexamination is that of the last section, i.e., that no strictly convex combination of two or more vectors h^j will be in H. But this follows from the observation that if x is lexicographically larger than x', then $\alpha x + (1 - \alpha)x'$ is also lexicographically larger than x' for $0 < \alpha \leq 1$. This demonstrates the validity of our main theorem without Assumption 2.2.

7. Concluding observations

The proof of Theorem 1.4 is based on an algorithm for integer programming problems. The algorithm may be shown to have sufficient monotonicity properties (under the labeling rule 4.1) so that finite convergence to a unique completely labeled primitive set can be demonstrated without the topological arguments of Sperner's lemma.

The major difficulty in implementing the algorithm is in the replacement operation described in *Section 3*. This would seem,

Mathematics and Economics: Scarf

Proc. Natl. Acad. Sci. USA 74 (1977) 3641

at first glance, to involve a search that is of the same order of difficulty as the original integer programming problem itself. There is, however, considerable structure in the problem, which may be capable of exploitation. For example, there will typically be a finite number of distinct simplices arising from primitive sets that form a fundamental region in the sense that every primitive set is equivalent, by translation, to one of these (see Fig. 1). Each simplex in lattice space arising from a primitive set will have an index $(+1, 0, -1)$ associated with it in such a way that the index sum, over the simplices covering any particular point in lattice space, is unity. Moreover, the sum of the algebraic volumes of the distinct simplices in a fundamental region is also unity. These and related observations permit us to describe the complete structure of primitive sets for certain classes of problems, which can then be readily solved by our algorithm.

The complexity of the integer programming problem seems to relate not only to the number of binding constraints but also to the singularities of the embedding of the associated simplicial complex C in R^n. A step in understanding this complexity would be a complete description of the integral polyhedra with 2^n vertices that appear in the argument in *Section 5*. At the present moment such a description is known only for $n \leq 3$.

1. Scarf, H. E. (1973) *The Computation of Economic Equilibria* (with the collaboration of T. Hansen) (Yale University Press, New Haven, CT).