G. HADLEY

LINEAR PROGRAMMING

LINEAR PROGRAMMING

by

G. HADLEY
University of Chicago

ADDISON-WESLEY PUBLISHING COMPANY, INC.
READING, MASSACHUSETTS, U.S.A.
LONDON, ENGLAND

Printed in the United States of America

Library of Congress Catalog Card No. 61–5306

PREFACE

Only a little over a decade has passed since George Dantzig formulated the general linear programming problem and developed the simplex method for its solution. In this period, the growth of interest in, and the use of, linear programming has been remarkable. Rarely indeed has a new mathematical technique found such a wide range of practical applications, and simultaneously received so thorough a theoretical development in such a short period of time. The extensive interest in linear programming which has arisen has brought with it the need for texts at different levels of difficulty, suitable for readers of widely varying backgrounds and mathematical maturity. The present work is intended for those who desire to study the subject in some depth and detail. It attempts to provide a fairly rigorous and complete development of the theoretical and computational aspects of linear programming as well as a discussion of a number of practical applications.

Chapter 1 introduces the general linear programming problem and exhibits a series of graphical examples. Chapter 2 covers the mathematical background needed. In Chapter 3 the fundamental theoretical results required for the simplex method are derived. Chapter 4 provides a detailed development of the computational procedure of the simplex method. The two-phase technique is introduced in Chapter 5, which also includes a discussion of the solutions and requirements spaces. Chapter 6 presents Charnes' perturbation technique and the generalized simplex method for resolving the degeneracy problem. The revised simplex method is covered in Chapter 7. Chapter 8 is devoted to duality; included in this chapter are the dual simplex algorithm and the primal-dual algorithm. The solution of transportation problems is the concern of Chapter 9. A novel approach is used to derive the transportation algorithm from the simplex method. Generalized transportation problems are also covered in this chapter. Chapter 10 discusses network flow problems, the primal-dual algorithm for transportation problems, assignment problems, and the transhipment problem. Chapter 11 treats a number of special topics, such as sensitivity analysis, treatment of upper bounds for the general linear programming problem, the primal-dual algorithm for capacitated transportation problems, the decomposition principle, and the relationships between linear programming and zero-sum two-person games. The application of linear programming to practical problems in industry is discussed in Chapter 12, and applications to economic theory are considered in Chapter 13.

The level of presentation in this book assumes that the reader has a familiarity with certain elementary topics in linear algebra (including convex sets). The necessary background material is reviewed in Chapter 2. However, the reader who has had no previous introduction to the material may find this review too abbreviated. It is suggested that such readers study concurrently the author's *Linear Algebra*, which covers in detail the material needed. (Incidentally, the notation used is consistent throughout both volumes.) As an aid to those who are simultaneously attempting to gain some knowledge of linear algebra, the present book develops its major themes in considerable detail, especially in Chapters 3, 4, 5, 6. It might be pointed out that a knowledge of convex sets in n dimensions is not essential for reading the text. The sections dealing with these topics can be omitted without loss of continuity.

Although the text is fairly complete, there are several topics associated with linear programming which are not to be found here. Two such topics, namely, linear programming problems some of whose parameters may be random variables and the solution of the general linear programming problem in integers, are considered by the author to be special cases of nonlinear programming problems and are discussed in a separate volume entitled *Nonlinear and Dynamic Programming*. No account is given of the use of analog computers to solve linear programming problems. Some material in this subject was included in an original version of the manuscript, but was dropped because it was felt to be a diversion from the main theme, of interest only to a small number of readers. General considerations regarding the solution of linear programming problems on digital computers are examined in the text, but no attempt is made to describe coding procedures in detail since the method chosen depends too much on the characteristics of the computer to be employed. Similarly, no detailed description of the linear programming codes available for various computers is given since these would be almost immediately out of date.

The text contains sufficient material for a two-semester course in linear programming although it can easily be used for a one-semester or one-quarter course in the subject. For example, the author has taught a one-semester graduate course at MIT devoted to the theory of linear programming. The students entering the course had no background in linear algebra, and hence about the first seven weeks were devoted to covering selected topics from the first six chapters of the author's *Linear Algebra*. For the remainder of the semester the material covered consisted of Chapters 3, 4, 5 (through Section 5–4), 6 (through Section 6–6), 7, 8 (through Section 8–7), and 9 (through Section 9–12) of the present volume. At the University of Chicago, the author has taught a one-quarter course to graduate students who had a course in linear algebra. After a brief review (about two weeks) of linear algebra and some discussion of convex sets, the above-mentioned material, the first four sections of Chapter 11, and most of Chapter 12 of this book were covered in the remainder of the quarter. It should be pointed out that this quarter course normally required about twelve hours per week of work outside of class.

The present volume should also serve as a supplementary text for courses in mathematical economics, engineering mathematics, operations research, or other courses which attempt to provide a serious treatment of linear programming. Finally, because of its completeness, it should be useful as a reference and as a text for self-study.

At the end of each chapter there will be found a collection of problems for solution. Some of these problems emphasize numerical techniques, while others concentrate on theoretical points. Within each of these two classes, there is a considerable range of difficulty. The author considers the problems to be very important, and anyone studying this work should, at the very least, read all the problems. When linear programming problems are solved by hand, it is a good idea to use a desk calculator if at all possible, or at least a slide rule. In this way, the numbers computed will be obtained as decimals. Of course, if a large-scale digital computer is used, the answers obtained will also be in decimal form. It does not seem sound to get into the habit of solving problems by means of fractions instead of decimal numbers, since fractions become impossibly cumbersome unless the original coefficients are small integers. For this reason, most of the tableaux in the text are presented in decimal form, even though they could have been expressed more simply and accurately as fractions. When the decimal numbers given are not exact, each element in the tableaux is expressed with roughly the same relative error. The same number of decimal places does not appear in each element.

For their helpful suggestions, the author wishes to express his appreciation to Professors H. Houthakker and H. Wagner, who read an early version of the manuscript, and to Professors R. Dorfman and T. M. Whitin, who read a later version. The author is also indebted to one of his students, M. A. Simonnard, whose thesis laid the foundations for the method of development used in the initial sections of Chapter 9. Jackson E. Morris supplied most of the quotations which appear at the beginning of each chapter. The School of Industrial Management at the Massachusetts Institute of Technology generously provided secretarial assistance for typing the manuscript.

G.H.

CONTENTS

CHAPTER 1

INTRODUCTION

". . . Since a crooked figure may
Attest in little place a million;
Then let us, ciphers to this great account,
On your imaginary forces work,"

Shakespeare—Henry V.

1–1 Optimization problems. Problems which seek to maximize or minimize a numerical function of a number of variables (or functions), with the variables (functions) subject to certain constraints, form a general class which may be called *optimization problems.*

Many optimization problems were first encountered in the physical sciences and geometry. The quest for solutions led to applications of the differential calculus and to the development of the calculus of variations. These classical optimization techniques have been known for over 150 years. They have been applied with considerable success to the solution of many problems in the physical sciences and engineering. Later, the differential calculus was found to be very useful in economics, especially in developing the important results in the classical theory of production and consumption

In the last ten or fifteen years, many new and important optimization problems have emerged in the field of economics and have received a great deal of attention. As a class, these problems may be referred to as programming problems. They are of so much interest because of their applicability to practical problems in government, military and industrial operations, as well as to problems in economic theory. In general, classical optimization techniques have been found to be of little assistance in solving these programming problems. Therefore, new methods had to be developed. In this book, we shall treat only a special but very important class of programming problems known as linear programming problems. We shall be concerned with the theory of linear programming, with numerical techniques for solving such problems, and with applications of the theory.

1–2 Programming problems. Broadly speaking, programming problems deal with determining optimal allocations of limited resources to meet given objectives; more specifically, they deal with situations where a number of resources, such as men, materials, machines, and land, are available, and are to be combined to yield one or more products. There are, however, certain restrictions on all or some of the following broad categories, i.e.: on the total amount of each resource available, on the quantity of each product made, on the quality of each product. Even within these restrictions there will exist many feasible allocations. Out of all permissible allocations of resources, it is desired to find the one or ones which maximize or minimize some numerical quantity, such as profit or cost.

The actual conversion of resources to products may be a simple mixing operation, such as mixing raw stock gasolines to form various motor fuels, or a complicated production process involving many types of machines and operations. In certain cases, the resources and products can be identical. For example: We may be interested in finding the cheapest way of transporting a product from a number of origins to a number of destinations.

Linear programming deals with that class of programming problems for which all relations among the variables are linear. The relations must be linear both in the constraints and in the function to be optimized.

1–3 An example. Let us consider a shop with three types of machines, A, B, and C, which can turn out four products, 1, 2, 3, 4. Any one of the products has to undergo some operation on each of the three types of machines (lathes, drills, and milling machines, for example). We shall assume that the production is continuous, and that each product must first go on machine type A, then B, and finally C. Furthermore, we shall assume that the time required for adjusting the setup of each machine to a different operation, when production shifts from one product to another, is negligible. Table 1–1 shows: (1) the hours required on each machine type per unit of each product; (2) the total available machine hours per week; (3) the profit realized on the sale of one unit of any one of the products. It is assumed that the profit is directly proportional to the number of units sold. We wish to determine the weekly output for each product in order to maximize profits.

Examination of Table 1–1 shows that the item with the highest unit profit requires a considerable amount of time on machines A and C; the product with the second-best unit profit requires relatively little time on machine A and slightly less time on machine C than the item with the highest unit profit. The product with the lowest unit profit requires a considerable amount of time on machine B and relatively little time on C. This cursory examination indicates that the maximum profit will not be achieved by restricting production to a single article. It would seem that at least two of them should be made. It is not too obvious, however, what the optimal product mix should be.

Suppose x_j is the number of units of product j produced per week. It is of interest to find the values of x_1, x_2, x_3, x_4 which maximize the total

TABLE 1–1

DATA FOR EXAMPLE

Machine type	Products				Total time available per week
	1	2	3	4	
A	1.5	1	2.4	1	2000
B	1	5	1	3.5	8000
C	1.5	3	3.5	1	5000
Unit profit	5.24	7.30	8.34	4.18	

profit. Since the available machine time is limited, we cannot arbitrarily increase the output of any one product. Production must be allocated among products 1, 2, 3, 4 so that profits will be maximized without exceeding the maximum number of machine hours available on any one of the groups of machines.

Let us first consider the restrictions imposed by the availability of machine time. Machines of type A are in use a total of

$$1.5x_1 + x_2 + 2.4x_3 + x_4 \text{ hours per week,}$$

since 1.5 hours are required for each unit of product 1, and x_1 units of product 1 are produced; and so on for the remaining products. Also, the total time used is the sum of the times required to produce each product. The total amount of time used cannot be greater than 2000 hours. Mathematically, this means that

$$1.5x_1 + x_2 + 2.4x_3 + x_4 \leq 2000. \tag{1–1}$$

It would not be correct to set the total hours used equal to 2000 for type A machines, since there may not be any combination of production rates that would use each of the three groups of machines to capacity. We do not wish to predict which machines will be used to capacity. Instead, we introduce a "less than or equal to" sign; the solution of the problem will indicate which machines will be used at full capacity.

For machines B and C we can write

$$x_1 + 5x_2 + x_3 + 3.5x_4 \leq 8000 \quad \text{(type } B \text{ machines)}, \tag{1–2}$$

$$1.5x_1 + 3x_2 + 3.5x_3 + x_4 \leq 5000 \quad \text{(type } C \text{ machines)}. \tag{1–3}$$

Since no more than the available machine time can be used, the variables x_j must satisfy the above three inequalities. Furthermore, we cannot produce negative quantities; that is, we have either a positive amount of any product or none at all. Thus the additional restrictions

$$x_1 \geq 0, \quad x_2 \geq 0, \quad x_3 \geq 0, \quad x_4 \geq 0 \tag{1–4}$$

require that the variables be non-negative.

We have now determined all the restrictions on the variables. If x_j units of product j are produced, the weekly profit z is

$$z = 5.24x_1 + 7.30x_2 + 8.34x_3 + 4.18x_4. \tag{1–5}$$

We wish to find values of the variables which will satisfy restrictions (1–1) through (1–4) and maximize the profit (1–5).

The above example is clearly a programming problem. Moreover, it is a linear programming problem because the restrictions and the function to be maximized involve only linear relations among the variables.

In practice, it may not be true that the profit derived from the sale of any one product is directly proportional to the number of units sold. More generally, the profit will be some function of the quantities produced, i.e.,

$$z = f(x_1, x_2, x_3, x_4).$$

If this function is not of the form (1–5), we have a nonlinear rather than a linear programming problem. For example, if the profit function were of the form

$$z = 5.24x_1^{1/2} + 7.30x_2^{1/2} + 8.34x_3^{1/2} + 4.15x_4^{1/2}, \tag{1–6}$$

then the determination of the variables which satisfy the constraints (1–1) through (1–4) and maximize Eq. (1–6) would be a special case of a nonlinear programming problem.

1–4 Linear programming. The preceding example illustrated how a linear programming problem and a particular case of a nonlinear programming problem can arise in practice. Linear programming is concerned with solving a very special type of problem—one in which all relations among the variables are linear both in the constraints and the function to be optimized. The general linear programming problem can be described as follows: *Given a set of m linear inequalities or equations in r variables, we wish to find non-negative values of these variables which will satisfy the constraints and maximize or minimize some linear function of the variables.*

Mathematically, this statement means: We have m inequalities or equations in r variables (m can be greater than, less than, or equal to r) of the form:

$$a_{i1}x_1 + a_{i2}x_2 + \cdots + a_{ir}x_r\{\geq, =, \leq\}b_i, \qquad i = 1, \ldots, m, \tag{1–7}$$

where for each constraint one and only one of the signs $\leq$, $=$, $\geq$ holds, but the sign may vary from one constraint to another. We seek values of the variables x_j satisfying (1-7) and

$$x_j \geq 0, \qquad j = 1, \ldots, r, \tag{1-8}$$

which maximize or minimize a linear function

$$z = c_1x_1 + \cdots + c_rx_r. \tag{1-9}$$

The a_{ij}, b_i, c_j are assumed to be known constants.

We have thus formulated the general linear programming problem which, in mathematical terms, can be represented by (1-7) through (1-9). A programming problem is linear if, in the constraints and function to be optimized, the variables appear only as linear forms. A linear form involving n variables x_j is an expression of the type $a_1x_1 + \cdots + a_nx_n + b$, where the a_j and b are constants. It is very important to see what the assumption of linearity implies. Intuitively, linearity implies that products of the variables, such as x_1x_2, powers of variables, such as x_3^2, and combinations of variables, such as $a_1x_1 + a_2 \log x_2$, cannot be allowed.

In more general terms, linearity can be characterized by certain additive and multiplicative properties. In the context of the above example, additivity means: If we use t_1 hours on machine A to make product 1, and t_2 hours to make product 2, the time on machine A devoted to products 1 and 2 is $t_1 + t_2$. In this case, the additivity property seems quite reasonable if the time required to convert from product 1 to 2 is negligible. However, not all physical processes behave in this way. If we mix several liquids of different chemical composition, it is, in general, not true that the total volume of the mixture is the sum of the volumes of the individual constituents. This is an example of a case where additivity may not hold.

The multiplicative property requires: (1) If it takes one hour to make a single item on a given machine, it takes ten hours to make ten parts; this also seems quite reasonable. (2) The total profit from selling a given number of units of a product is the unit profit times the number of units sold; this is not always true. In general, the profit is not directly proportional to the number of units sold even if the selling price is constant, since manufacturing costs per unit may vary with the number of units made. Thus the linearity implied in a linear programming problem is not always expected to be an absolutely accurate representation of the real world. Fortunately, the assumed linearity is often a close enough approximation of actual conditions so that it can provide very useful answers.

One other important restriction is inherent in a linear programming problem: It is assumed that the variables x_j can take on any values allowed by the restrictions (1-7) and (1-8); in other words, we cannot, for example, require that the variables assume only integral values. If the additional restriction is imposed that the variables must be integers, then, in general, we do not have any longer a linear programming problem. Actual situations often require that the variables be integers; such problems are frequently solved by linear programming, and the answers are rounded off to the nearest integers which satisfy the constraints. This may or may not be a valid approximation. In general, the approximation is good if the solution requires that a large number of units of each variable be used. Since, in the example discussed in Section 1–3, the production is assumed to be continuous, it is not necessary that the optimal x_j be integers. Of course, in practice no attempt would be made to schedule weekly production down to a fraction of a unit.

The assumption that the variables can vary continuously goes somewhat deeper than indicated in the previous paragraph. Fundamentally, everything in the real world comes in discrete units, and nothing is infinitely divisible. However, the basic building blocks (molecules, photons, etc.) are often so small in comparison with the quantities under consideration that for all practical purposes (i.e., to ten or fifteen decimal places), it can be assumed that the physical quantity is continuously variable. The real difficulty appears when the discrete units are not small in comparison with the magnitudes of the variables. In situations of this kind, one must be concerned about the discreteness of the variables.

The function to be optimized, (1–9), is called the *objective function*. Note that no constant term appears in the objective function, i.e., we do not write $z = \sum_{j=1}^{r} c_j x_j + c$. The reason for this is simple. The values of the x_j which optimize z are completely independent of any additive constant c. Hence, if there is such a constant, it can be ignored during the process of determining the best x_j, and added to z after the problem has been solved.

Mathematically, the constraints (1–8) which require that the variables x_j be non-negative do not differ from the constraints (1–7). However, when solving a linear programming problem, the non-negativity constraints are handled differently from the other constraints. For this reason, we shall refer to the non-negativity constraints as *non-negativity restrictions*, while the term *constraint* will be used to denote constraints other than the non-negativity restrictions. Thus when we say that there are m constraints on the problem, we mean that there are m constraints of the form (1–7). Then in addition, there are the non-negativity restrictions. This terminology will save some confusion later.

Any set of x_j which satisfies the constraints (1–7) will be called a *solution* to the linear programming problem. Any solution which satisfies the non-negativity restrictions is called a *feasible solution*. Any feasible solution which optimizes the objective function is called an *optimal feasible solution*. The task of solving a linear programming problem consists in finding an optimal feasible solution. Normally, there will be an infinite number of feasible solutions to a linear programming problem. Out of all these solutions, we must find one which optimizes the objective function.

1-5 The transportation problem; another example. In practice, linear programming has been of particular significance in its application to so-called transportation problems. To provide the reader with a little more feeling for the nature of linear programming, we shall discuss these briefly. Later two chapters will be devoted to solving transportation problems.

A typical transportation problem can be described as follows: Given amounts of a uniform product are available at each of a number of different origins (e.g., warehouses). We wish to send specified amounts of the product to each of a number of different destinations (e.g., retail outlets). The cost of shipping one unit amount from any one origin to any one destination is known. Assuming that it is possible to ship from any one warehouse to any one retail outlet, we are interested in determining the minimum-cost routing from the warehouses to the retail outlets.

We shall suppose that there are m origins and n outlets. Take x_{ij} to be the number of units shipped from origin i to destination j. Note that here it is convenient to use a double subscript since it simplifies the notation. For a given i (warehouse), there are n possible j-values (retail outlets to which units can be shipped). Hence we have a total of mn different x_{ij}. Since negative amounts cannot be shipped, we must have $x_{ij} \geq 0$ for all i, j.

Let a_i be the number of units of the product available at origin i and b_j the number of units required at destination j. We cannot ship more goods from any one origin than are available at that origin. Hence summing over all destinations, we have

$$\sum_{j=1}^{n} x_{ij} = x_{i1} + x_{i2} + \cdots + x_{in} \leq a_i, \qquad i = 1, \ldots, m. \tag{1-10}$$

There are m such constraints, one for each origin. We must supply each destination with the number of units desired; thus

$$\sum_{i=1}^{m} x_{ij} = x_{1j} + \cdots + x_{mj} = b_j, \qquad j = 1, \ldots, n. \tag{1-11}$$

The total amount received at any destination is the sum over the amounts received from each origin. The needs of the outlets can be satisfied if and only if

$$\sum_{i=1}^{m} a_i \geq \sum_{j=1}^{n} b_j. \tag{1-12}$$

We assume that this is the case.