Linear Programming—Aluminum Alloy Blending

IBM

Data Processing Application

Linear Programming–
Aluminum Alloy Blending

IBM

Data Processing Application

CONTENTS

Copies of this and other IBM publications can be obtained through IBM branch offices. Address comments concerning the contents of this publication to IBM, Technical Publications Department, 112 East Post Road, White Plains, N.Y. 10601

INTRODUCTION

This manual explores the use of linear programming (LP) in the aluminum alloying industry. It discusses the nature of the industry, the decisions involved in producing an alloy, the mechanics of alloying, and the techniques and advantages of applying LP to the industry. Particular emphasis is placed on explaining how linear programming can be used to:

- Optimize furnace scheduling
- Minimize the cost of an alloy blend
- Minimize off-compositions
- Maximize the use of low-cost scrap
- Reduce scrap holding
- Provide a more accurate scrap inventory
- Influence purchasing decisions

Linear programming is not new, nor are its economic benefits strange to many business and industrial areas. As a proven mathematical technique for allocating resources to minimize cost or maximize profit, it is today an indispensable decision-making tool in many companies.

Perhaps the oldest and most extensive user is the petroleum industry, where LP routinely aids in such areas as refinery scheduling, gasoline blending, and refinery expansion studies. There are few, if any, refineries which do not rely in some way upon linear programming problem solutions.

In other areas, LP is used to determine optimal feed mixes, chemical compositions, food mixes, transportation routes, steel compositions, etc.

It is significant to note that the problems solved by LP in several of these areas closely parallel the basic problem of aluminum alloy blending: how to produce at least cost a product of known composition using available raw materials and physical resources.

Despite a growing and more diverse use of LP, and the often astounding economical advantages it imparts, the technique as an everyday management aid remains in its infancy. There are two reasons for this.

First, there has been both ignorance and fear regarding the mathematical aspects of LP problem formulation. Many people have felt that a staff of highly trained specialists in LP problem formulation is a necessary bridge between those operating personnel who know the problem and the use of linear programming for decision-making. Ordinarily, the user does not need to know a great deal about how the computer finds the optimal solution or how it arrives at the effect of changes. He does, however, need to know the elements of these methods in order to formulate his problem most effectively and to interpret the results intelligently. Actually, the basic mathematics involved in expressing a problem as an LP model is readily learned and easily used. The main ingredient of successful LP application is a firm grasp of the problem to be solved, and not mathematical dexterity.

Second, the solution of an LP problem of any practical size requires a computer and an associated LP code. In the past, computers were far more abundant than codes. Thus, to make use of the LP method, a computer user frequently resorted to building his own LP code — rather costly practice.

However, LP codes are now available as application programming packages for IBM computing systems. They offer another unique opportunity to expand the application of IBM computers for better day-to-day management and decision-making in the aluminum industry.

PROBLEM PROFILE

The alloy blending problem can be expressed as deciding how to satisfy the requirements of an alloy specification most economically.

PROBLEM ECONOMICS

The economics of the problem are manifested in many areas of company operation and vary from company to company — an overhead expense in one company might be a line production cost in another; a company which produces its own raw materials may have a different concept of costs than a firm that purchases these materials.

Similarly, a company that produces and sells scrap and pure metals, as well as consuming them internally, must often weigh the profits of sale against the cost of internal consumption and the profit from the sale of product alloy. Also, there is the need to consider the availability and cost of purchased materials.

The complexity of the most economical blending of an alloy is readily apparent, even when viewed only in terms of how to determine the costs of raw materials.

For purposes of this discussion, it will be assumed that raw materials can be accurately priced and that the prices encompass the following costs:

1. Purchase price or manufacturing cost
2. Transportation cost
3. Handling cost
4. Furnace cost
5. Inventory cost
6. Sales/purchasing costs

For the moment, the problem is narrowed to deciding how to make the most economical use of raw materials of known cost and availability in computing a furnace charge. This, of course, is the heart of the solutions to be derived from the linear programming technique. The technique can also contribute substantially to arriving at more precise raw materials costs, raw materials inventory, and selling price of the manufactured alloy or purchase price of raw materials.

HISTORICAL CONDITIONS

Historically, the implements for calculating furnace charges have been desk calculators, the alloy specifications, an often outdated raw materials inventory, experience, and intuition.

Since an alloy specification usually provides minimum and maximum bounds on the amounts of each metallic element to be used, it is not an extremely difficult task to compute a workable furnace charge.

Although economic conditions are always taken into account, it simply is not possible to explore all or even a large part of the possible charges and their comparative benefits by manual methods.

Problem Complexity

Assume, for example, the need to produce 1000 pounds of alloy 7000, whose specification calls for elements 1 and 2 in the following proportions:

Alloy 7000 (1000 pounds)

Ingredients	Requirements (pounds)
Element 1 minimum	150
Element 1 maximum	200
Element 2 minimum	30
Element 2 maximum	50
Other elements minimum	750
Other elements maximum	820

Assume also the following availability of two scraps having a relatively high content of elements 1 and 2:

	Scrap 1	Scrap 2
Element 1	30%	50%
Element 2	5%	10%
Availability	500 pounds	300 pounds
Cost	.30/pound	.75/pound

A natural impulse would be to use all 500 pounds of the cheaper scrap 1 since the minimum requirements for element 1 would be immediately satisfied (500 X 30% = 150 pounds). This would, however, not satisfy the element 2 requirements since 500 X 5% = 25 pounds, which is five pounds short of the minimum requirement of 30 pounds.

There are several alternatives to be considered at this point:

1. Use all of scrap 1 as originally planned, and add five pounds of pure element 2. The cost would be $150 plus the cost of the pure metal.
2. Use all of scrap 2, which will satisfy the requirements for both element 1 and element 2, but at a cost of $225.
3. Use all of scrap 1 and ten pounds of scrap 2, or some other combination of the two scraps that satisfies the alloy specification.

Taking into account only elements 1 and 2, either alternative 1 or 3 would probably prove economically attractive. But if there were several additional elements to be considered and ten to 20 available grades of scrap, as is usually the case, the number of alternative solutions would increase tremendously.

The problem becomes even more complex because of at least two other factors which often enter into the computations: (1) the frequent need to control the amount of one element used in proportion to the amounts of other elements, and (2) the desirability of using raw materials in units of ingots rather than pounds.

The Manual Approach

Almost the only practical means of manually arriving at a furnace charge is to first select the raw material that comes closest to matching the alloy specification requirements, then add pure metals and/or scrap to satisfy one or more additional requirements without disrupting those satisfied earlier (that is, without exceeding the maximum of an ingredient whose minimum has been equaled or exceeded).

Once a charge has been formulated for an alloy specification, it is used as the basis for specifying charges for similar alloys; ingredients are added or removed to compensate for the specification differences, and, of course, amounts are adjusted as required.

While this is a most practical approach, it perpetuates the use of pure metals and popular scrap — to the exclusion of the off-specification, low-cost grades of scrap. The result is that scrap not easily used is accumulated in inventory, while pure metals and on-specification scrap are purchased.

Linear programming provides a means of examining all existing possible combinations and quickly arriving at the most economical furnace charge. It is also possible to force the use of ingots, low-cost scrap, or any scrap in large supply while at the same time minimizing the cost of the charge.

LP MODEL FORMULATION

A linear programming model for aluminum alloy blending is a mathematical representation of all known and estimated factors which influence the calculation of furnace charges. A single-furnace model encompasses one alloy and one furnace; multifurnace models may represent an entire alloying shop and several different alloy types.

The following discussion deals with (1) the types and sources of data that must go into an LP model, and (2) procedures for formulating single-furnace and multifurnace problems.

INPUT DATA REQUIREMENTS

The basic information required to formulate the LP model includes the following:

1. Alloy specifications showing the composition of each alloy to be blended
2. The inventory level of all raw materials that might be used in the production of one or more alloys
3. The composition of the raw materials
4. The per-pound cost of raw materials
5. The state of the raw materials, that is, ingot form or loose scrap
6. Furnace capacities
7. The number of pounds of each alloy to be produced

Most of this information is readily available from purchasing, cost accounting, inventory accounting, or other sources and is probably used in existing systems for computing furnace charges. Where exact information cannot be readily obtained, estimates should be made since it is an easy matter to change the input data and re-solve the problem once an optimal solution has been obtained. Indeed, the easy calculation of the effect of changes in the input is a prime advantage of the linear programming approach.

SINGLE-FURNACE PROBLEM

Figure 1 states the specifications for an alloy, which will be called alloy 7000. Note that each ingredient has a minimum and maximum limit — pounds of element in the product alloy — and that provision is made for general impurities.

In order to blend 10,000 pounds of alloy 7000, the required composition is 550 to 590 pounds of zinc, 140 to 190 pounds of copper, 245 to 275 pounds of magnesium, and so on for each of the listed ingredient metals.

Element	Minimum (lbs.)	Maximum (lbs.)
Zinc (Zn)	555.0	590.0
Copper (Cu)	140.0	190.0
Magnesium (Mg)	245.0	275.0
Chromium (Cr)	19.0	22.0
Beryllium (Be)	2.0	4.0
Iron (Fe)	0.0	15.0
Silicon (Si)	0.0	10.0
Manganese (Mn)	0.0	3.0
Nickel (Ni)	0.0	2.0
Titanium (Ti)	0.0	2.0
Lead (Pb)	0.0	2.0
Tin (Sn)	0.0	2.0
Bismuth (Bi)	0.0	8.0
Aluminum (Al)	8867.0	9049.0
General Impurities	0.0	8.0

Figure 1. Specifications for alloy 7000

Available raw materials, each of which contains some percentage of one or more of the required ingredients, are as follows:

1. Four different grades of aluminum, AL1 through AL4, ranging in cost from 23¢ to 28¢ per pound
2. Pure copper at 31¢ per pound, pure magnesium at 38¢ per pound, and pure zinc at 22¢
3. A chromium-aluminum alloy at 27¢ per pound and a beryllium-aluminum alloy at $3.60 per pound
4. Eleven scrap aluminum alloys, SC1 through SC11, at either 20¢ per pound or 21¢ per pound

This provides a total of 20 different sources for the metals required by alloy 7000. It will be assumed that the supply of only one raw material, SC1, is limited (to 900 pounds); an unlimited supply of all other materials is available for allocation. It will also be assumed that the state of the raw materials (either ingot or loose form), is not a problem factor.

An LP matrix would be formulated as shown in Figure 2 for the problem as stated. Each of the available raw materials is expressed as a matrix column, or activity; the information from the alloy specification forms the row names and the elements of the right-hand-side (RHS) column. The coefficients of the row inequations, rows ZN through GX, express the proportion of each metal contained in one pound of each raw material.

The cost row coefficients are the per-pound costs of each raw material; the last two rows in the model constrain the amount of SC1 which may be used and the total amount of alloy to be produced (furnace load). The significance and derivation of each matrix component will be discussed in detail in the following paragraphs. All data in the model is expressed in terms of pounds.

Although the recording of all data in pounds requires considerable conversion of the problem data and necessitates the use of rather large fractional values, it assures a uniformly scaled matrix — a basic requirement of any LP system. It also means that the problem solution will be expressed in terms of pounds.

Activities (Matrix Columns)

The sources of the various alloy elements are expressed as matrix columns, or problem activities. The mechanics of preparing the matrix columns are relatively simple. Once the constraint rows have been established and listed, each raw material is listed and its composition is recorded vertically according to row.

The problem activities are invariably the most critical parts of the model because they offer the widest margin for error. Where the problem rows contain known values as dictated by the alloy specification, furnace capacity, inventory level, etc., raw materials composition cannot be as closely ascertained.

There are exceptions, as is the case with pure and near-pure metals and standard alloy scrap. However, in many cases the composition of a scrap alloy is derived from estimates based on experience or blending records. Although the accuracy of these estimates does not affect the computation of a problem solution, it does affect its validity. The validity of computed charges is measured by the number of off-compositions that occur during the actual blending operations.

Constraints (Matrix Rows)

Problem constraints (rows) appear in the matrix as a set of simultaneous linear equations or inequations. In the problem shown in Figure 2, there are the following types of constraint rows:

1. Cost constraint
2. Element specifications or restrictions (material balances)
3. Raw materials availability constraints
4. Physical constraints (furnace load)

An LP model invariably contains at least one of each constraint row type. It may contain others which place special restrictions on the problem solution. (This is discussed later.)

Cost Constraint

The cost row contains the per-pound costs of each raw material and constrains the solution to that combination of activities which satisfies all other constraints at the minimum possible cost.

The cost of any feasible solution can be determined by solving the cost equation using selected activity levels. For example, if 9,000 pounds of A4, 600 pounds of SC1, and 400 pounds of SC11 represented a feasible solution, its cost (see Figure 2) would be computed as follows:

	A4		SC1		SC11		RHS
COST	.23(9000)	+	.21(600)	+	.21(400)	=	$2,280.00

For each feasible problem solution, the LP system performs essentially the same basic arithmetic to determine which of the solutions minimizes the cost of the alloy.

Element Specifications

Constraints on the amount of each ingredient metal that may be charged are represented in the model (Figure 2) by rows ZN through GX. The coefficients

Column Names		Pure Aluminum 1	2	3	4	Pure Copper	Pure Magnesium	Beryllium Aluminum Alloy	Pure Zinc	Chromium Aluminum Alloy	Scrap 1	2	3	4	5	6	7	8	9	10	11			
Row Names		A1	A2	A3	A4	C	M	B/A	Z	C/A	SC1	SC2	SC3	SC4	SC5	SC6	SC7	SC8	SC9	SC10	SC11		RHS	
	Cost	.28	.26	.25	.23	.31	.38	3.60	.22	.27	.21	.20	.21	.20	.21	.20	.21	.20	.21	.20	.21	=	Min	$
Zinc Minimum	ZN								.95		.0009	.0012	.0568	.0563	.0460	.0455	.0009	.0006	.0009	.0008	.0675	≧	555.	lbs
Zinc Maximum	ZX								.95		.0009	.0012	.0568	.0563	.0460	.0455	.0009	.0006	.0009	.0008	.0675	≦	590.	lbs
Copper Minimum	CN					1.00					.0444	.0026	.0152	.0149	.0071	.0071	.0447	.0623	.0034	.0003	.0195	≧	140.0	lbs
Copper Maximum	CX					1.00					.0444	.0026	.0152	.0149	.0071	.0071	.0447	.0623	.0034	.0003	.0195	≦	190.0	lbs
Magnesium Minimum	MN						1.00				.0042	.0060	.0248	.0238	.0343	.0343	.0143		.0093	.0249	.0265	≧	245.0	lbs
Magnesium Maximum	MX						1.00				.0042	.0060	.0248	.0238	.0343	.0343	.0143		.0093	.0249	.0265	≦	275.0	lbs
Chromium Minimum	CHN									.0300	.0001	.0018	.0020	.0019	.0013				.0019	.0016	.0020	≧	19.0	lbs
Chromium Maximum	CHX									.0300	.0001	.0018	.0020	.0019	.0013				.0019	.0016	.0020	≦	22.0	lbs
Beryllium Minimum	BN							.0600														≧	2.0	lbs
Beryllium Maximum	BX							.0600														≦	4.0	lbs
Iron Maximum	IX	0004	.0006	.0011	.0026						.0024	.0026	.0016	.0019	.0017	.0016	.0026	.0017	.0030	.0015	.0014	≦	15.0	lbs
Silicon Maximum	SX	.0005	.0006	.0007	.0012						.0101	.0106	.0013	.0011	.0013	.0011	.0013	.0010	.0062	.0011	.0008	≦	10.0	lbs
Manganese Maximum	MGX										.0079	.0003	.0005	.0004	.0018	.0017	.0052	.0025	.0002	.0002	.0002	≦	3.0	lbs
Nickel Maximum	NX										.0001	.0002										≦	2.0	lbs
Titanium Maximum	TX										.0004	.0004	.0004	.0004	.0002	.0002	.0003	.0005	.0003			≦	2.0	lbs
Lead Maximum	LX										.0001	.0001	.0003	.0003	.0002	.0002	.0001	.0001				≦	2.0	lbs
Tin Maximum	TNX										.0001	.0001	.0003	.0003	.0002	.0002	.0001	.0001				≦	2.0	lbs
Bismuth Maximum	BIX																		.0005			≦	8.0	lbs
General Impurities	GX										.0001	.0002						.0025				≦	8.0	lbs
Scrap 1 Limit	SCX										1.00											≦	900.0	lbs
Furnace Load	FL	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	≧	10000.	lbs

Figure 2. Alloy 7000 matrix model

in each row are the amounts of the element (represented by the row) in one pound of each activity. The unknowns are the various activity levels, or the number of pounds of each activity to be included in the charge.

Using only the Z, SC1, SC2, and SC3 activities in the matrix, the constraint for ZX (zinc maximum) can be written explicitly as the inequality:

	Z	SC1	SC2	SC3	RHS
ZX	.95(Z) +	.0009(SC1) +	.0012(SC2) +	.0568(SC3) ≤	590.0

This states that the aggregate optimal levels of these activites (number of pounds to be charged) must not exceed 590 pounds.

Using 500 pounds of Z, 1,500 pounds of SC3, and none of SC1 or SC2, this constraint (ZX ≤ 590.0) could be satisfied as follows:

	Z	SC1	SC2	SC3	RHS
ZX	.95(500) +	.0009(0) +	.0012(0) +	.0568(1500) =	560.2

The same activity levels would also satisfy the constraint ZN ≥ 555.0, but would fall considerably short of satisfying the next constraint in the matrix, CN ≥ 140.0 (copper minimum), and possibly many of the other constraints on the problem, as illustrated below:

	Z	SC1	SC2	SC3	RHS
ZN	.95(500) +	.0009(0) +	.0012(0) +	.0568(1500) =	560.2
ZX	.95(500) +	.0009(0) +	.0012(0) +	.0568(1500) =	560.2
CN	.00(500) +	.0444(0) +	.0026(0) +	.0152(1500) =	22.8

From the mathematics used, we learn that there will never be more nonzero activities in an optimal solution than there are constraints; there may be fewer. This means, for instance, that in a problem with two constraints and ten possible activities, at least eight of the possible activities will be at the zero level in the optimal solution. And, further, the activity levels of these activities must satisfy simultaneously all the problem constraints.

It should be noted that only 14 of the 15 elements in the alloy 7000 specification (see Figure 1) are represented in the model. The missing element, aluminum, need not be controlled directly since it will comprise the difference between the 10,000 pounds of the product alloy and the total amount of the 14 other controlled elements.

It should also be noted that it is not necessary to include minimum constraint rows for elements whose minimum content in the product alloy may equal zero. The reason for this is that linear programming systems will not allocate negative amounts of any ingredient (since it is not possible to take away an element which was never allocated to begin with).

Raw Materials Availability Constraints

Constraints on the permissible level may be included for any or all activities in the problem. These are conventionally placed at the bottom of the matrix following the last ingredient row.

Constraints of this type consist of a coefficient of 1 in the appropriate activity column and the maximum or minimum permissible level in the right-hand-side vector.

The matrix for alloy 7000 contains only one such constraint row, SCX, which limits the total allocation of SC1 to less than or equal to 900 pounds (see Figure 2). Using the same simple technique, constraints on any raw material may be imposed according to current inventory levels or general availability of the raw material. For example:

... SC1	SC2	SC3	SC4	SC5 ...		RHS
1					≤	900.0
	1				≤	500.0
		1			≤	900.0
			1		≤	50.0
				1	≤	1500.0

Similarly, the same type of constraint may also be used to indicate a desired minimum allocation of any activity. For example, if the supply of SC2 were abnormally large, it might be desirable to attempt to force the allocation of a specified minimum number of pounds (SC2N):

	SC1	SC2	...	RHS
SCX	1		≤	900.0
SC2N		1	≥	1800.0

Such a constraint does not guarantee that any of SC2 will be allocated (methods of forcing allocation are discussed later), but it does guarantee that if SC2 is allocated, no less than 1,800 pounds will be used.

There is another consideration in the use of minimum allocation constraints: the minimum amount specified (RHS value) must be feasible. In the problem as stated, the constraint on SC2 of 1,800 pounds is completely infeasible because of the constraint on silicon and the silicon content of SC2. By inspecting the sample matrix (Figure 2), it will be noted that one pound of SC2 contains .0106 pounds of silicon and, further, the total amount of silicon in alloy 7000 cannot exceed ten pounds (see specification in Figure 1). Since .0106 X 1800 = 19.08 pounds of silicon (exceeding the minimum limit), none of SC2 would be allocated. Infeasibilities of this sort can usually be discovered by visual inspection of the matrix.

Physical Constraints

A model <u>must</u> contain a constraint which specifies the amount of alloy to be blended, which is often equal to the capacity of the furnace to be charged.

In the matrix for alloy 7000 (see Figure 2), this constraint is imposed by the furnace load (FL) row. Its form is a coefficient of 1 in each column and the weight of the desired charge in the RHS column.

Summary

Construction of the basic LP model entails little more than organizing, in a special format, the data historically used in calculating furnace charges. Once constructed initially and converted to an input media for computer processing, the model becomes a master record. It will be updated regularly to account for new conditions such as the addition or deletion of activities, new inventory constraints or changes to existing ones, new costs, etc.

In addition to meeting the requirements of the LP system, the matrix format offers an excellent graphic display of the problem. In fact, with experience, it will become possible to visually inspect matrices and predict which materials will most likely be used and, in some cases, the amounts that will be used. It may be often desirable to make changes based on these observations before the problem is solved.

MULTIFURNACE MODEL

A multifurnace model contains data for computing more than one furnace charge during the same computer run. Depending upon its contents, the model may be used for computing consecutive charges for one furnace, or charges for several furnaces to be charged simultaneously. In addition, the charges may be for the same alloy or for several different alloys.

Figure 3 is an outline of a multifurnace model that contains the data for alloy 7000 and two other alloys, 6000 and 5000, which are to be blended in different furnaces.

Note that the matrix is divided into three submatrices, each of which contains the material balance constraints for one of the alloys. At the bottom of the submatrices are the constraints for furnace loads and inventory availability (these might just as well have been placed at the bottom of the appropriate submatrices). A common cost row is at the top of the matrix.

The objective is to allocate the same raw materials in the production of three different alloys. Note that the activities are the same for each submatrix, but are prefixed by either 1, 2, or 3 to identify them with one of the furnaces. For example, 1SC1, 2SC1, and 3SC1 are the same scrap alloy (SC1 in the single-furnace model shown in Figure 2). The solution might possibly call for 100 pounds of 1SC1, 200

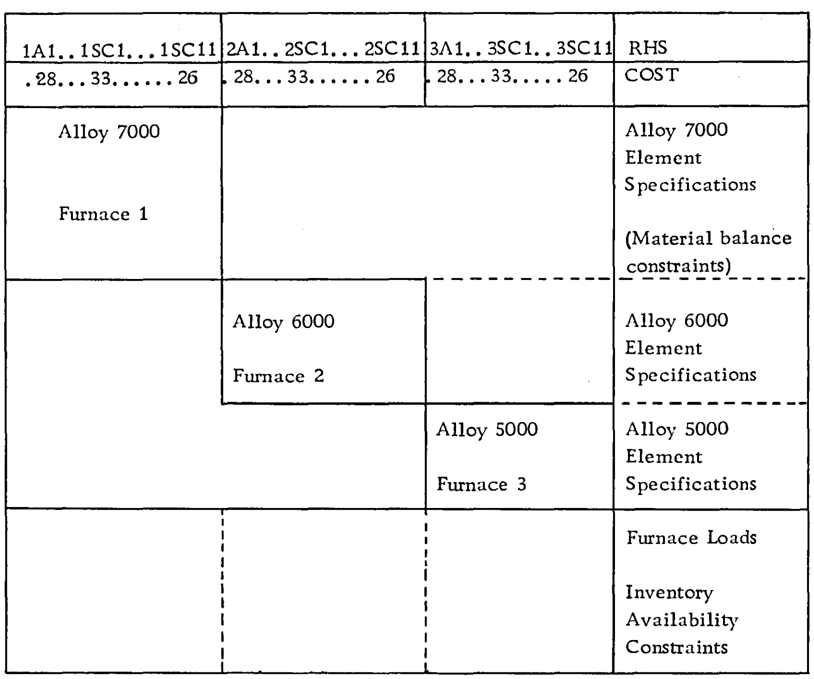

Figure 3. Outline of a multifurnace model

pounds of 2SC1, and 700 pounds of 3SC1 for a total allocation of 1,000 pounds of SC1.

Availability constraints on SC1 could limit the amount of this scrap available for use in each furnace. For example:

$$1SC1 \leq 400 \text{ pounds}$$
$$2SC1 \leq 400 \text{ pounds}$$
$$3SC1 \leq 400 \text{ pounds}$$

These constraints ensure that no more than 400 pounds of SC1 could be used in each furnace.

It is not necessary to preallocate limited-supply scrap in this fashion. In fact, it is probably more desirable to allow the computer to determine the most economical division of the scrap between the furnaces while obtaining the linear programming solution. This is done simply by limiting the total amount of SC1 available to the three furnaces. The same activities are not necessarily associated with each submatrix. There may be activities common to all submatrices and others peculiar to only one or two submatrices.

Hence, the formulation techniques used in building multifurnace models are the same as for single-furnace models. The major difference is that there are several sets of activities, material balance constraints, furnace load constraints, and activity level constraints.

There is, however, a considerable difference in the concept of applying multifurnace models. The biggest advantage, of course, is the simultaneous computation of several furnace charges which optimize the consumption of raw materials. Using this concept, it becomes possible to allocate all raw materials to all furnaces in the most advantageous manner.

SPECIAL CONSTRAINTS AND FORMULATIONS

In addition to the usual problem constraints discussed in the single-furnace problem, it is often necessary to express (1) complex process or technological constraints, or (2) special material balances in the model. Examples of two such constraints are discussed to illustrate the formulation techniques involved.

The first can be classified as an equivalence constraint where relationships are to be established between the elements. The equivalence example illustrates the formulation of constraints to control the allocation ratio between elements. Such constraints are needed in several situations, for example, to ensure the solubility of certain elements or to ensure a required conductivity in the alloy.

The second can be classified as a substitution constraint, and it illustrates how to allocate groups of lower cost elements in place of all or part of a more expensive element, while ensuring that the alloy produced still retains the required special properties readily achieved by the more expensive element alone.

The formulation techniques that follow relate to the basic matrix in Figure 4, which shows only part of an LP model matrix. For simplification, in the following examples E1 will represent the accumulative amounts of the element E1 contained in materials A1 through SC13, as depicted by row 1 in the basic matrix (see Figure 4). Similarly, E2, E3, and E4 will represent the respective elements depicted by rows 2, 3, and 4 in Figure 4. For example, E3 represents .0001 A1 + ... + .1060 SC12 + .0002 SC13.

	A1	SC12	SC13		RHS
Cost	.23	.25	.35	=	Min.
(1)	.0002	.0802		≥	4
(2)		.0600	.0160	≥	6
(3)	.0001	.1060	.0002	≥	16
(4)			.0150	≥	20

Figure 4. Basic matrix

Example 1 — Equivalence

Equivalence constraints often involve controlling the allocation ratio of one element to another. This is quite common in aluminum alloy blending with the use of the two elements iron and silicon.

In linear programming, the amount of an element (constraint row) contained in a solution is expressed in terms of the variables (columns) available. In order to express a direct relationship between elements, variables must be added to the matrix that will represent the element.

One method of showing this relationship is to add slack variables representing the amount of an element in a solution that is greater than the minimum amount required, or the amount that is less than the maximum amount required.

Referring to the basic matrix in Figure 4, such variables can be added in the following manner:

Given:

$$E1 \geq 4$$

So:

$$E1 - S1 = 4$$

Where:

E1 is the accumulative amounts of the element E1 in the solution contained in material A1 through SC13.

S1 is the slack variable representing the amount of E1 in the solution that exceeds the minimum required.

4 is the minimum amount of E1 required in the solution.

In a similar fashion, slack variables are established for the other elements. The addition of the slack variables changes the constraint rows from inequalities to equalities as shown in Figure 5.

If the alloy is to contain equal amounts of element E1 and element E2, this can be expressed as E1 = E2, or E1 - E2 = 0.

With the addition of slack variables in Figure 5, the following relationships can be noted:

E1 - S1 = 4 or, E1 = 4 + S1
E2 - S2 = 6 or, E2 = 6 + S2
E3 - S3 = 16 or, E3 = 16 + S3
E4 - S4 = 20 or, E4 = 20 + S4

Now it is possible to express equal allocation of the elements E1 and E2 in terms of the slack variables:

$$E1 = E2$$
$$4 + S1 = 6 + S2$$
$$S1 - S2 = 6 - 4$$
so $S1 - S2 = 2$

This equation ensures that E1 = E2, and, as it can be expressed in terms of the slack variables, it can now be added to the matrix (Figure 5) as shown in Figure 6 by the row labeled EQ1.

Other allocation ratios can be expressed in a similar fashion. For example, referring to the matrix in Figure 5, the ratio 2E2 = E3 can be formulated as follows:

$$2E2 = E3$$
$$2(6 + S2) = 16 + S3$$
$$12 + 2S2 = 16 + S3$$
$$2S2 - S3 = 16 - 12$$
or $2S2 - S3 = 4$

This equivalence equation can now be added to the matrix (Figure 6) as shown in Figure 7 by the row labeled EQ2.

Instead of adding slack variables that represent the amount of an element in a solution over the minimum requirement or under the maximum requirement, another method is to add variables (columns) to the basic matrix to represent the actual amounts of elements used in the solution.

The addition of such variables will leave the original constraint row unconstrained, and the constraint is transferred to an additional row in the matrix. In the new row the constraint is applied against the variable representing the actual amount of the element used in the solution.

	A1	SC12	SC13	S1	S2	S3	S4		RHS
COST	.23	.25	.35	.00	.00	.00	.00	=	Min.
(1)	.0002	.0802		-1				=	4
(2)		.0600	.0160		-1			=	6
(3)	.0001	.1060	.0002			-1		=	16
(4)			.0150				-1	=	20

Figure 5. Adding slack variables

	A1	SC12	SC13	S1	S2	S3	S4		RHS
COST	.23	.25	.35	.00	.00	.00	.00	=	Min.
(1)	.0002	.0802		-1				=	4
(2)		.0600	.0160		-1			=	6
(3)	.0001	.1060	.0002			-1		=	16
(4)			.0150				-1	=	20
EQ1				+1	-1			=	2

Figure 6. Adding equivalence constraint EQ1

	A1	SC12	SC13	S1	S2	S3	S4		RHS
COST	.23	.25	.35	.00	.00	.00	.00	=	Min.
(1)	.0002	.0802		-1				=	4
(2)		.0600	.0160		-1			=	6
(3)	.0001	.1060	.0002			-1		=	16
(4)			.0150				-1	=	20
EQ1				+1	-1			=	2
EQ2					+2	-1		=	4

Figure 7. Adding equivalence constraint EQ2

Referring to the basic matrix in Figure 4, equations would be formulated in the following manner:

Given:

$E1 \geq 4$

Let:

$E1 = AE1$

Therefore:

$E1 - AE1 = 0$

And:

$AE1 \geq 4$

where:

E1 still represents the accumulative amount of the element E1 in the solution contained in materials A1 through SC12.

AE1 represents the actual amount of the element E1 used in the solution.

4 is the minimum amount of element E1 required in the solution.

In a similar fashion, equations are formulated for the other elements, and columns and rows are added, resulting in the matrix shown in Figure 8.

One advantage of this method is that any change to the RHS value of a constraint row can be readily handled without going through the arithmetic procedures that are necessary when using slack variables.

Another advantage is the ease of expressing ratios between elements. For example, if the alloy is to contain equal amounts of element E1 and E2, an equivalence equation depicting the relationship is easily formulated.

Given:

$E1 = E2$

Also:

$E1 = AE1$, and $E2 = AE2$

Therefore:

$AE1 = AE2$

Or: $AE1 - AE2 = 0$

Similarly, the ratio $2E2 = E3$ can be readily formulated as follows:

Given:

$2E2 = E3$

Also:

$E2 = AE2$, and $E3 = AE3$

Therefore:

$2(AE2) = AE3$

Or:

$2AE2 - AE3 = 0$

These two equivalence equations are shown as rows EQ1 and EQ2, respectively, in Figure 9.

The advantages of adding variables that represent the actual amount of an element used instead of slack variables that represent the excess over or under the constraint value are not as apparent in the simple examples illustrated as they would be in a more complex situation.

To control the (proportional) allocation of a group of elements to achieve a desired conductivity, consider a conductivity equation of the form:

$$b = a_1E1 + a_2E2 + a_3E3 + a_4E4$$

Where b is the desired conductivity level, and the a's are metallurgical coefficients reflecting the conductivity of the 4 elements E1, E2, E3, and E4. This equation can be incorporated into the LP matrix in the same manner as the previous allocation ratio constraint.

	A1	SC12	SC13	AE1	AE2	AE3	AE4		RHS
COST	.23	.25	.35	.00	.00	.00	.00	=	Min.
(1)	.0002	.0802		-1				=	0
(2)		.0600	.0160		-1			=	0
(3)	.0001	.1060	.0002			-1		=	0
(4)			.0150				-1	=	0
(5)				+1				≥	4
(6)					+1			≥	6
(7)						+1		≥	16
(8)							+1	≥	20

Figure 8. Transferring constraints to columns that represent actual amounts of the elements used

	A1	SC12	SC13	AE1	AE2	AE3	AE4		RHS
COST	.23	.25	.35	.00	.00	.00	.00	=	Min.
(1)	.0002	.0802		-1				=	0
(2)		.0600	.0160		-1			=	0
(3)	.0001	.1060	.0002			-1		=	0
(4)			.0150				-1	=	0
(5)				+1				≥	4
(6)					+1			≥	6
(7)						+1		≥	16
(8)							+1	≥	20
EQ1				+1	-1			=	0
EQ2					+2	-1		=	0

Figure 9. Depicting the addition of equivalence constraints

Example 2 — Substitution

The practice of exchanging or substituting for all or a portion of a higher-cost alloying element with specified groups of lower-cost elements is becoming quite widespread. This can be an important factor in the overall alloy cost.

The cost of an element is thought of as an indirect evaluation of the element, implied by the cost of the input materials in which the element occurs. For example, in the LP model matrix illustrated in Figure 2, the element zinc is available in a pure form (95% Z at 22¢ per pound, as well as in small percentages (.06% - 5.68%) in the scrap materials. However, chromium occurs in much smaller proportions in the scrap (.01% - .2%) and the chromium aluminum alloy material contains only 3% chromium; at 27¢ a pound it is more expensive than zinc. Under these conditions, it is less expensive to have zinc in an alloy than it is to have chromium.

Consider the problem where an equally allocated amount of three elements (E1 + E2 + E3) is allowed to be allocated in place of an equal amount of a fourth element, E4. The element E4 has a minimum specification, as do the three elements E1, E2, and E3 (refer to Figure 4).

Once again, the only way direct relationships between elements can be expressed in linear programming is to add variables to the matrix that will represent the elements. This example illustrates the advantage of adding variables to represent the actual amount of elements used and the transferring of the constraint value to these added variables.

In the example, the problem is to allocate to the elements E1, E2, and E3, where E1 = E2 = E3, all or part of the amount of element E4 required in the solution.

Starting with the basic matrix as shown in Figure 4, variables are added to represent the amount of each element used in the solution, and rows are added to permit the transferring of the constraint values to the new variables, leaving the original rows unconstrained.

The formulation of the equations is as follows:

Given:

$$E1 \geq 4$$

Let:

$$E1 = UE1 + SE1$$

Therefore:

$$E1 - UE1 - SE1 = 0$$

And:

$$UE1 + SE1 \geq 4$$

Where:

E1 represents the accumulative amount of the element E1 in the solution contained in materials A1 through SC13.

UE1 represents the actual amount of element E1 used, excluding the amount substituted for E4.

SE1 represents the actual amount of element E1 used to substitute for element E4.

4 is the minimum amount of element E1 required in the solution.

In a similar fashion, equations are formulated for elements E2 and E3. In the case of element E4, the equations are formulated as follows:

Given:

$$E4 \geq 20$$

Let:

$$E4 = AE4 + SE4$$

Therefore:

$$E4 - AE4 - SE4 = 0$$

And:

$$AE4 + SE4 \geq 20$$

Where:

E4 represents the accumulative amount of the element E4 in the solution contained in materials A1 through SC13.

AE4 represents the actual amount of the element E4 used in the solution.

SE4 represents the accumulative amounts of elements E1, E2 and E3 which have been substituted for element E4.

20 is the minimum (actual or substituted for) element E4 required in the solution.

These equations are shown in the matrix in Figure 10.

To ensure that the amounts of the elements E1, E2 and E3 being substituted for E4 are equal, the respective substitute amounts must be aligned in one column. Referring to Figure 10, this is accomplished by merging the columns SE1, SE2 and SE3 into one column SEQ. To reflect the fact that SE1 + SE2 + SE3 = SE4, or 3SEQ = SE4, the coefficient + 3 is added to the SEQ column in row 8, which is the constraint row for element E4.

This completes the constraints and ensures that the amount of elements E1, E2 and E3 are equal and, that all four elements have met their minimum requirements. These constraints now complete the model matrix as shown in Figure 11.

Elements will often have maximum constraints alone or in addition to minimum constraints. These are incorporated through a simple extension of the model.

Element substitution constraints will possibly be more complicated than the one just modeled. The substitution example culminating in Figure 11 should serve as an adequate basis for other substitution constraints.

	A1	SC12	SC13	UE1	UE2	UE3	AE4	SE1	SE2	SE3	SE4		RHS
COST	.23	.25	.35	.00	.00	.00	.00	.00	.00	.00	.00	=	Min.
(1)	.0002	.0802		-1				-1				=	0
(2)		.0600	.0160		-1				-1			=	0
(3)	.0001	.1060	.0002			-1				-1		=	0
(4)			.0150				-1					=	0
(5)				+1				+1				=	4
(6)					+1				+1			=	6
(7)						+1				+1		=	16
(8)							+1				+1	=	20

Figure 10. Adding variables for used and substituted values as well as transferring constraints to new rows

	A1	SC12	SC13	UE1	UE2	UE3	AE4	SEQ		RHS
COST	.23	.25	.35	.00	.00	.00	.00	.00	=	Min.
(1)	.0002	.0802		-1				-1	=	0
(2)		.0600	.0160		-1			-1	=	0
(3)	.0001	.1060	.0002			-1		-1	=	0
(4)			.0150				-1		=	0
(5)				+1				+1	≥	4
(6)					+1			+1	≥	6
(7)						+1		+1	≥	16
(8)							+1	+3	≥	20

Figure 11. Completed matrix model

PROBLEM ANSWERS

The answers produced by an LP system provide information to be acted upon or analyzed in many different company operating areas. This information ranges from the physical problem answers — the amounts of each raw material to be charged into the furnace — to the economic data for analyzing the criticality of raw materials costs and demand.

OPTIMAL SOLUTION

The solution to the problem previously formulated in Figure 2 is presented and discussed to illustrate the information available in the optimal solution of a blending model. The optimal solution output using the IBM 1620/1311 Linear Programming System is shown in Figure 12.

NAME	ACTIVITY LEVEL
C	66.561
M	19.959
B/A	33.333
Z	404.793
C/A	111.724
SC4	2476.077
SC8	274.808
SC10	5704.371
SC11	908.374

Figure 12. Optimal output solution — BASIS VARBLS

In the optimal solution output labeled BASIS VARBLS (Figure 12), the names of all raw materials that are to go into the furnace charge are listed under the NAME column. These are the optimal activities. Under the ACTIVITY LEVEL column are listed the optimal level of each optimal activity or the number of pounds of each raw material to be charged.

Assuming the solution is to be implemented without change, this information can be immediately disseminated to two operating centers: inventory accounting and the alloying shop.

In the inventory accounting department, the solution serves as a record of raw materials consumed in producing the alloy and may be used directly to update inventory records. In the blending shop, it serves as a sort of work order to be followed in charging the furnace.

Other Optimal Output

The slack variables for each constraint row and the cost of the solution — that is, the minimum cost of producing the furnace load required — are listed in Figure 13.

NAME	ACTIVITY LEVEL	SIMPLEX MULT.
COST	2149.248	
ZN	35.000	
ZX		.006
CN		.084-
CX	50.000	
MN		.154-
MX	30.000	
CHN		1.462-
CHX	3.000	
BN		56.231-
BX	2.000	
IX		2.599
SX		25.771
MGX		.473
NX	2.000	
TX	.872	
LX	1.230	
TNX	1.230	
BIX	8.000	
GX	7.313	
SCX	900.000	
FL	.000	.226-

Figure 13. Other output solution — SLACKS

Under the column heading NAME are listed the names of the slack variables formed to make equalities from inequalities. The linear programming system has given these variables the names of the inequalities with which they are associated. For example, ZN is the slack variable associated with the minimum zinc requirement, ZX is the slack associated with the maximum zinc constraint. Under the column heading ACTIVITY LEVEL are listed the actual slack activity levels computed. (These levels indicate how the solution differs from the maximum or minimum levels given in the constraints.) For example, ZN = 35. 0 means that the total zinc in the alloy mix exceeds the minimum zinc constraint by 35 pounds. ZX is blank which indicates that the total zinc in the mix is equal to the maximum constraint (590. 0 lbs).

The marginal costs of introducing nonoptimal slacks into the solution are listed under the heading SIMPLEX MULT. The marginal cost is the cost of introducing one unit of the slack into the solution and, at the same time reducing some other variable(s) from the solution so that the constraint corresponding to the slack does not exactly meet its bound (maximum or minimum). An example is the simplex multiplier . 006 associated with ZX.

This means that if the zinc maximum was 589 pounds (one pound less than its current 590-pound level), the resulting total alloy cost would increase . 6¢.

The next simplex multiplier, -.084, is associated with CN, the copper minimum level. This indicates that the cost of the alloy would decrease by 8.4¢ if the copper minimum was lowered by one pound.

This data provides an indication of the cost of raising or lowering RHS values and shows how relaxed alloy specifications can raise or lower the end metal cost.

POST-OPTIMAL OUTPUT

DO. D/J

The data output in the DO. D/J report (see Figure 14) indicates the amounts by which costs of nonbasis variables would have to be reduced, before these variables would tie for entry into the optimum solution basis. For example, if the cost of A1 (the first high-grade aluminum input in the matrix) was reduced by more than 6.8¢ to less than 21.2¢ per pound, this activity would then be included in the furnace charge. The levels of one or more optimal activities would be adjusted to account for the usage of this aluminum. One activity would also leave the solution at this stage and would no longer be in the optimal furnace charge.

The problem would have to be resolved in order to determine which activity leaves the solution and in order to find the new activity levels of the other variables which changed in the solution.

Inactive Materials Analysis

The DO. D/J report (Figure 14) contains the names and prices of raw materials that were not used in the alloy charge. This information, when combined with the other data in the report, makes it possible to determine why these activities were not included in the solution. It is also possible to evaluate the increased costs to be incurred, should it become necessary to replace an optimal material with one that is nonoptimal.

From a long-range viewpoint, however, this list becomes even more significant. Over a period of time, it becomes possible to determine which raw materials remain largely unused in charges for many different alloys. Trends and statistics for each scrap can be developed and then put to use in achieving more economical operation in many areas. First, such information informs the purchasing department that these slow-moving materials should be purchased carefully, if at all. Second, decisions regarding the sale of these materials can be made authoritatively according to in-house utilization rates, costs, and inventory level. Third, it allows the inventory level of these materials to be reduced to an optimal level.

All the above can be readily evaluated in dollar savings. A quicker turn of profit, however, can come from building a heat around off-specification scrap. Using linear programming, experimenting with such scraps in the production of various alloys can become an everyday practice — and the savings can amount to many thousands of dollars per year in reduced raw materials costs.

COST. R

The COST. R output report (see Figure 15) indicates the cost range over which the costs of the variables in the current solution can be varied without changing the optimal charge. For example, the cost of C (pure copper) can be varied from its present level of 31¢ per pound to as high as 43.996¢ or as low as 28.564¢ per pound without this variable (pure copper) being excluded from the optimal charge (solution).

This report also indicates which activities would enter the solution (optimal charge) if these cost bounds were to be exceeded. For example, if scrap 10 (SC10) was to be raised above 20.231¢ per pound (its current level being 20¢) then scrap 10 would no longer help form a least-cost charge and scrap 6 (SC6) would enter the charge.

However, the costs of the optimal activities can be varied between their upper and lower limits without affecting the allocation levels of any of the optimal activities.

The affect on the optimal solution by violating the cost range can sometimes be manually computed; normally, however, it is necessary to make the desired cost change and resolve the problem.

In addition to providing information for experimentation or manual adjustments of the solution, the cost ranges guide the purchasing department in buying or selling raw materials. They provide a reference for evaluating a vendor's quotations, and to ensure that raw materials costs will not raise the product alloy production costs. Moreover, the same information indicates where lower materials costs can contribute most to production savings.

APPLYING THE PROBLEM ANSWERS

The answers produced by the LP method pervade virtually every major operation related to the blending of an alloy. The actual furnace charges and much of the supporting analytical data are directly applicable to current shop operations and can be put to immediate use.

This data can become the basis for integrating and improving systems for purchasing, inventory control, quality control, pricing, cost analysis, and furnace scheduling.

NAME	CURRENT COST	REDUCED COST	BASIS VALUE
A1	.280	.068	.212
A2	.260	.051	.209
A3	.250	.045	.205
A4	.230	.042	.188
SC1	.210	.250	.040-
SC2	.200	.476	.276-
SC3	.210	.014	.196
SC5	.210	.015	.195
SC6	.200	.002	.198
SC7	.210	.021	.189
SC9	.210	.147	.063

Figure 14. Postoptimal output — DO. D/J

NAME	CURRENT COST	HIGHEST COST	HI-VAR	LO-VAR	LOWEST COST
C	.31000	.43996	ZX	MGX	.28564
M	.38000	.41619	MGX	MN	.23175
B/A	3.60000	INFINITY		BN	.22615
Z	.22000	.22556	ZX	IX	.20582
C/A	.27000	.28443	MGX	CHN	.22941
SC4	.20000	.20145	IX	MGX	.19518
SC8	.2000	.20105	MGX	IX	.18560
SC10	.2000	.20231	SC6	IX	.19893
SC11	.2100	.21232	IX	MGX	.20869

Figure 15. Postoptimal output — COST. R

Suggestions of how LP answers can be used in several areas were made in the preceding descriptions of the output reports. The following summarizes these ideas and amplifies them as they apply to each operating center.

FURNACE CHARGES

The furnace charges computed through the use of LP are optimum — based on the content of the alloy model, there is no other combination of raw materials or proportional allocation of these materials that will allow the desired alloy to be produced at less cost.

Multifurnace solutions or series of single-furnace solutions can be used to effectively schedule daily production within an entire blending shop. The solutions can be grouped so that alloys that are metallurgically similar are charged through the same furnace. The lists of raw materials by furnace for each charge not only speed the movement of materials from inventory to the blending shop but reduce materials handling within the shop.

QUALITY CONTROL

The accuracy of the charges computed with linear programming techniques should contribute substantially to reducing the number of off-composition heats. The loss of certain elements during blending of certain alloys is apparent to the experienced observer. By simply modifying the constraints for these elements, the probability of off-composition can be reduced even more. When off-composition does occur, remedial action can be quickly determined through re-solution of the problem after the required adjustments have been made to the model.

The question of "what price for quality" can be answered hypothetically by tightening the minimum and maximum constraints for an alloy and obtaining solutions for each change.

INVENTORY

Reports produced by LP show exactly how much of each raw material is consumed in a furnace charge. This information can be fed into the inventory system, either manually or mechanically, to update the on-hand figures. The updated inventory then serves as the basis for replenishing inventory through purchase or internal production and for updating the LP model master records to reflect the changes in raw materials availability.

The same information can be used for inventory leveling studies. Normal demand rates can be established for all materials carried in inventory. For popular materials, most economical order quantities can be determined to reduce the frequency and related cost of purchase; the level of rarely used materials can be reduced to reasonable amounts through sale or forced allocations. This not only converts unneeded inventory to cash but also reduces the costs of inventory maintenance.

PURCHASING AND COST ANALYSIS

The cost ranges for each material used in an optimal furnace charge can easily become the most widely used data produced by LP. Either directly or indirectly, they provide answers to these typical questions:

1. How much can the cost of zinc fluctuate without affecting the cost of the alloy?
2. What is the price at which any individual material must be purchased in order to possibly reduce alloy production cost?
3. What metals are most sensitive to cost change (smallest cost ranges)? Which are least sensitive (widest cost ranges)?
4. How much will it cost to substitute a non-optimal metal for an optimal one?

These and other questions arise many times in actual day-to-day operations. Timely and accurate LP solutions provide many answers which aid in deciding the best action. The same questions can be asked and answered experimentally to study the affect of a variety of possible conditions on the costs of producing an alloy or alloys.

PRODUCT RESEARCH AND PRICING

As noted before in the discussion of quality controls, the weight restrictions (RHS) for an alloy blend can be relaxed or tightened to study the cost of quality. Similarly, substitutions of a variety of equivalent metals for a standard metal can be studied for cost, feasibility, and pricing.

The LP technique also provides an accurate means of pricing the standard alloys and the infrequently blended special alloys. Costing and pricing an alloy that has never before been blended can be an elusive task. However, by building a model for the unusual blend and solving it, the pricing problem is made considerably more manageable.

MANAGEMENT STUDIES

Linear programming is as applicable to hypothetical situations as it is to real ones. Indeed, the use of the technique to test the affect of proposed changes on current operations is one of its most powerful advantages.

Use of LP in this manner amounts to playing the "what if" game: What if one material becomes unavailable? What if we purchase certain raw materials instead of making them? What if we make instead of purchase? What effect on inventory levels and costs would be felt if furnace capacity were increased or decreased by 50%? What if a 10% reduction were made in the selling price of alloy 7000? How could it be most readily recouped through lower production costs?

By building experimental models or modifying existing ones, many different courses of action for many different situations can be studied before a change is actually made.

SUMMARY

It is clear that the computation and recomputation of optimal furnace charges based on available raw materials is the heart of linear programming application. But the savings through better furnace scheduling, increased capacity, higher product quality, and more accurate inventory and purchasing information can amount to many thousands of dollars.

POSTOPTIMAL OPERATIONS AND GENERAL PROCEDURES

A simple diagram of the LP processing cycle is shown in Figure 16. Operations connected by solid lines are basic to any run; those indicated by dotted lines may or may not be required, depending upon existing conditions.

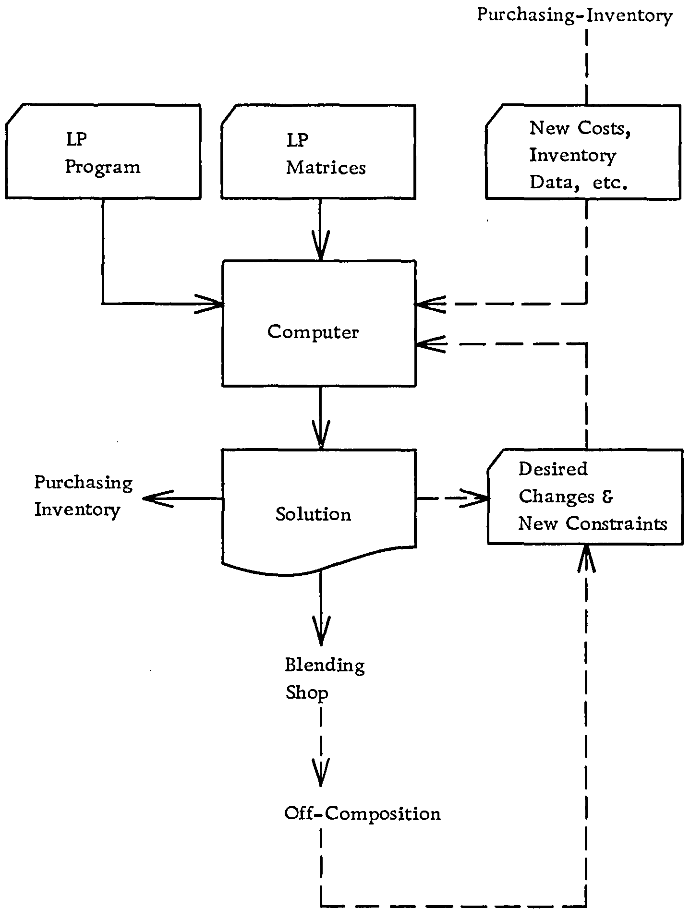

Figure 16. Processing cycle

OPERATING MODES

Actual production runs may be made under either of two modes. First, constraints on inventory availability and inventory state (ingot or loose form) can be omitted from the matrices. Secondly, these constraints can be imposed upon the problem before an initial solution is obtained.

The advantages of one mode over the other depend upon the existing situation. Generally, however, an initial assumption of unlimited available resources and no ingot restrictions would seem more beneficial for the following reasons:

1. The LP code is given considerably more freedom in determining the least expensive raw materials mix.
2. There is an opportunity, through study of the initial solution, to determine whether such constraints are actually required.
3. If constraints are desirable, the initial solution pinpoints the applicable raw materials and the magnitudes of the over-inventory allocations.
4. There are indications where "tradeoffs" among raw materials can best be attempted to achieve the desired ingot allocations or to balance allocations to known inventory levels.
5. The cost of such constraints becomes immediately apparent by comparing initial and secondary solutions.
6. Re-solution of a problem once it has been solved initially can be done very economically.

Under the unlimited mode, then, the approach is to begin with an initial optimal solution and work backwards to one which is best for the additional constraints.

The following describes how inventory and ingot constraints might be obtained on the basis of an initial optimal solution.

Ingot Constraints

By referring to the initial optimal solution, activities whose optimal levels are to be rounded to ingot weights can be quickly determined. Assume that the level of SC4 in the solution illustrated below is to be increased to 500 pounds, or five ingots.

NAME	ACTIVITY LEVEL
SC4	450.0000
SC6	150.0000

A first step would be to add the constraint equation below to the problem:

	SC4		RHS
Ingot Bal.	1	=	500

If it is desired to reduce or limit the allocation of SC6 to 100 pounds, another constraint could also be added. In fact, this might contribute heavily to achieving the 500-pound allocation of SC4:

	SC4	SC6		
Ingot Bal	1		=	500
Level		1	=	100

It should be noted that such constraints will not guarantee the desired results. In many cases, it may become necessary to obtain several solutions, changing the constraints according to the results of the previous solution. However, this approach to controlling ingot allocations has proved highly effective in actual application.

Inventory Constraints

The techniques for writing these constraints were explained earlier under LP Model Formulation.

FREQUENCY OF SOLUTION

The frequency with which charges are computed will be a function of the arrival of new information. For correcting off-compositions, data feedback may occur every few minutes. New inventory data will be entered as frequently as materials are received, or as a by-product of the daily charges.

CONTINUOUS POURING

During the course of developing the techniques for aluminum alloy blending, it was conjectured that it may someday be possible to further apply the system to achieve continuous pouring. Under such a system, several small furnaces are linked with one large furnace. Each of the smaller furnaces are continuously charged. These in turn continuously charge the larger furnace from which the finished alloy is poured.

It is beyond the scope of this paper to treat the subject of continuous pouring in detail. Suffice it to say that the possible advantages are many — reduced setup cost, a higher rate of production, small fluctuation in alloy quality, and tight, uniform furnace scheduling.

E20-0127-0

Printed in U.S.A. E20-0127-0

IBM

International Business Machines Corporation
Data Processing Division
112 East Post Road, White Plains, N. Y. 10601