
APPLIED SYSTEMS ANALYSIS

Engineering Planning and Technology Management

Richard de Neufville

*Massachusetts Institute
of Technology*

McGraw-Hill Publishing Company

New York St. Louis San Francisco Auckland Bogotá Caracas Hamburg
Lisbon London Madrid Mexico Milan Montreal New Delhi Oklahoma City
Paris San Juan São Paulo Singapore Sydney Tokyo Toronto

CHAPTER 3

INTRODUCTION TO CONSTRAINED OPTIMIZATION

3.1 THE PROBLEM

In practice, systems engineers always have to make their designs conform to a variety of constraints. They generally have to work within a budget. They will have to satisfy engineering standards designed to ensure the quality, the performance, and the safety of the system. They will also have to meet legal restrictions the community has imposed to achieve its health and other objectives. The design of the primary water supply system for New York City, for example, had to provide water at a minimum pressure of 40 psi (pounds per square inch) as measured at the curb, with less than a legal maximum for parts per million of various chemicals.

Constraints generally express some objectives for the system which could not be dealt with analytically in an obvious way. Setting constraints on a system is an alternative to the difficult problem of dealing explicitly with the multiple outputs of a system, which has been very difficult to do (see Section 2.5). Thus New York City's standard of 40 psi at the curb represents an effort to ensure that the water pressure will be adequate for all residents of the city.

The analysis will often demonstrate that many of these restrictions do not make much sense in detail. They may easily cost far more to meet than is worth spending on the objective they represent. New York's standard of 40 psi for water pressure thus entailed enormous extra costs for the construction of its huge underground aqueducts. Tens of millions of dollars could have been saved in the design if even slightly lower pressures had been permitted in a few areas. While agreeing with the principle behind the restrictions, in this case the idea that water

pressure should be "adequate," there is generally much scope for improvements at the level of detail (perhaps 38 psi would have been a better standard than 40 psi).

The reason that many restrictions on design do not make sense in detail is that they were never established on the basis of an analysis of the system. New York's standard of 40 psi was certainly not based on a demonstration that it was the optimum level, or that it was functionally preferable to 39 or 41 psi; it was almost certainly agreed to because it was somewhere in the right region and because it was a nice round number.

Part of the systems engineer's job in the future should be to help define the appropriate levels of restrictions or, even better, to optimize the multiple output production process directly (see Chapter 8). Meanwhile, however, the designer often will simply have to meet the restrictions.

The difficulty in meeting restrictions on the design is that they complicate the optimization process. To see this consider some function of many variables:

$$g(\mathbf{X}) = g(X_1, \dots, X_i, \dots, X_n)$$

which we assume to be continuous and analytic in that it has first and second derivatives. If this function is unconstrained, the necessary conditions for its optimum (a maximum or a minimum) is that all the first derivatives with respect to its arguments equal zero:

$$\frac{\partial g}{\partial X_i} = 0 \quad \text{all } i$$

On the other hand, if the function is constrained, these conditions for optimality may or may not apply. The optimum may then be at the intersections of the function with the constraint. Figure 3.1 makes the point. In a large scale system with hundreds of variables and constraints, the optimum could be located at an infinity of points.




FIGURE 3.1

In unconstrained situation (left) the maximum is at A; with constraints (right) it is either at B, if $X \leq b'$, or at A again if $X \leq b''$.

The problem is then to find ways to define optimality conditions which can pinpoint the optimum when there are constraints. For analytic functions there are two principal means of doing this, Lagrangeans and the Kuhn-Tucker conditions. These are explained in detail in Sections 3.3 and 3.4.

3.2 OBJECTIVE

The purpose of this chapter is to provide insight into the nature of the optimal design of constrained systems. It develops two key concepts that are essential in the practical application of optimization, particularly in sensitivity analysis (Chapter 6). These concepts are those of shadow prices and complementary slackness, defined in Sections 3.3 and 3.4 respectively.

These key ideas are developed by an incremental process. The presentation builds on the common understanding of unconstrained optimization in two steps of increasing complexity. The first, that of equality constraints solved by Lagrangeans, brings out the concept of shadow prices. The second, involving inequality constraints and the Kuhn-Tucker conditions, introduces complementary slackness.

This process focuses on the necessary conditions for optimality, as these are the ones that bring out the key issues. It avoids the second-order sufficiency conditions for optimality. These are both complex and beside the point since systems analysis essentially never uses these conditions in practical situations.

3.3 LAGRANGEANS

Lagrangeans constitute the easiest solution to optimization subject to constraints. Historically they are also the first. They may be considered an intermediate step to the general optimality conditions published by Kuhn and Tucker in 1951.

Lagrangeans solve the problem of optimizing a function subject to *equality constraints*, ones that must be met exactly. These routinely occur in practice. A standard example is the requirement to use one's budget entirely—otherwise it might be cut in the following period. Formally, the Lagrangean addresses the problem of optimizing $g(\mathbf{X})$ subject to a series of constraints $h_j(\mathbf{X}) = b_j$. Putting the matter into the standard vocabulary used with optimization, we call the function to be optimized the *objective function*. We then write the problem as

$$\begin{array}{ll} \text{Optimize:} & g(\mathbf{X}) \\ \text{Subject to:} & h_j(\mathbf{X}) = b_j \quad \text{and} \quad \mathbf{X} \geq 0 \end{array}$$

The condition that all \mathbf{X} be positive is a standard assumption that in no way limits practical applications: variables can always be redefined to meet this requirement. The purpose of this condition is to simplify the second order, the sufficiency conditions of optimality. In linear programming, discussed in Chapter 5, this sign convention is also exploited to facilitate the optimization procedure.

The problem is solved by constructing a new function to be optimized. This is called the *Lagrangean* and is

$$L = g(\mathbf{X}) - \sum \lambda_j [h_j(\mathbf{X}) - b_j]$$

Each parameter λ_j is a constant associated with a particular constraint equation, and is to be solved for. These parameters are known as the *Lagrangean multipliers*.

The marvelous aspect of the Lagrangean is that it simultaneously permits us to optimize $g(\mathbf{X})$ and to satisfy the constraints $h_j(\mathbf{X}) = b_j$. All we do is to apply the conventional process for unconstrained functions to the Lagrangean. The difference is that whereas in considering $g(\mathbf{X})$ we had i variables, one for each X_i , we now have $(i + j)$ variables, including one for each of the λ_j multipliers associated with the j constraints.

Applying the standard method of optimization to the Lagrangean we set all first derivatives equal to zero and obtain:

$$\begin{aligned} \frac{\partial L}{\partial X_i} &= \frac{\partial g}{\partial X_i} - \sum_j \lambda_j \frac{\partial h_j}{\partial X_i} = 0 \\ \frac{\partial L}{\partial \lambda_j} &= h_j(\mathbf{X}) - b_j = 0 \end{aligned}$$

The $(i + j)$ set of X_i and λ_j that satisfy these $(i + j)$ equations will determine the optimum. This solution will be designated (X^*, λ^*) where the *superscript* * indicates optimality.

Notice that the optimal solution to the Lagrangean clearly satisfies the constraints placed on the original problem; they are identical to the $\partial L / \partial \lambda_j$ equations that help define the optimal solution. Likewise, this solution also maximizes $g(\mathbf{X})$ since, at optimality where the constraints are met, the second term in the Lagrangean equals zero and optimizing L is then equivalent to optimizing $g(\mathbf{X})$ (see box on following page).

The Lagrangean multiplier has a most important interpretation. It gives the systems designer valuable information about the sensitivity of the results to changes in the problems. At the optimum, the Lagrangean multiplier λ_j equals the rate of change of the optimum objective function $g^*(\mathbf{X})$ as the constraint $b_j = h_j(\mathbf{X})$ varies. To see this, take the partial derivative of the Lagrangean with respect to this constraint and obtain

$$\frac{\partial L}{\partial b_j} = \lambda_j \quad \text{all } j$$

Now since the constraints are identically equal to zero at the optimum,

$$g^*(\mathbf{X}) = L \quad \text{and} \quad \frac{\partial g^*}{\partial b_j} = \frac{\partial L}{\partial b_j} = \lambda_j \quad \text{all } j$$

In the vocabulary of economics and systems analysis, the Lagrangean multiplier is known as the *shadow price* of the constraint. This neatly describes the parameter: it is the implicit price to be paid, in terms of changes of $g(\mathbf{X})$, per unit change of the constraint b_j . Being a derivative, the value of the shadow price

Example Use of Lagrangeans

Consider the problem:

$$\text{Maximize: } g(\mathbf{X}) = 3X_1X_2^2$$

$$\text{Subject to: } h(\mathbf{X}) = X_1 + X_2 = 3 \quad \text{and} \quad X_1, X_2 > 0$$

The Lagrangean is

$$L = 3X_1X_2^2 - \lambda(X_1 + X_2 - 3)$$

There is only one λ because there is only one constraint.

There are $(i + j) = 3$ equations defining the optimal solution. These are

$$\frac{\partial L}{\partial X_1} = 3X_2^2 - \lambda = 0$$

$$\frac{\partial L}{\partial X_2} = 6X_1X_2 - \lambda = 0$$

$$\frac{\partial L}{\partial \lambda} = X_1 + X_2 - 3 = 0$$

The first two equations imply that $X_2^* = 2X_1^*$. The third then leads to

$$\mathbf{X}^* = (1, 2) \quad \text{and} \quad \lambda^* = 12$$

The optimum is

$$g^*(\mathbf{X}) = 12$$

in general varies with the value of the constraint; the values calculated are thus instantaneous, at the margin, for any given value of the constraint.

Semantic caution: Although the Lagrangean multiplier is called the shadow price, it has no necessary connection with money. Its units are always in terms of those of the function to be maximized, $g(\mathbf{X})$, divided by those of the relevant constraint. Thus for the design of New York City's water supply system: when the objective was to maximize the capacity to deliver water subject to constraints on the pressure, the shadow price of the pressure standard was in terms of gallons per psi. The shadow price is expressed in terms of money only when $g(\mathbf{X})$ is some monetary amount such as dollars of benefit from a system.

The shadow price on a constraint is the quantitative indication of whether it is worthwhile to alter this requirement. It is the information that allows the designer to judge to what extent each of the constraints on a system make sense in detail. For New York City's water supply system, for example, the shadow price on the requirement that water be delivered throughout the city at a minimum pressure

of 40 psi at the curb was extremely high; unless the community for some reason thought this was an absolutely vital threshold, it would be reasonable to lower this standard for the areas of highest elevation so that the resulting improvements to the overall system could be shared by the whole community.

This issue of sensitivity of the optimum to changes in the definition of the problem is one of the central topics of all practical applications of systems analysis (see Section 3.1). Chapter 6 is devoted to this topic, and Sections 6.2 and 6.3 particularly focus on shadow prices.

3.4 KUHN-TUCKER CONDITIONS

The most usual situation in practice is that we have to deal with *inequality constraints*. Our plans must perform at least as well as some standards, but can also perform better if that is convenient. Thus New York City's standard for water pressure of 40 psi at the curb does not mean that pressure everywhere must equal 40 psi at the curb; it simply means that it can not be less and therefore that it may in many places actually be more than 40 psi.

The existence of inequality constraints significantly increases the number of solutions that we have to consider. They are not restricted to lie on a line or function, but may be anywhere inside the space defined by the restrictions. Figure 3.2 illustrates this point.

A formal statement of the optimization problem with inequality constraints is

$$\text{Optimize: } g(\mathbf{X})$$

$$\text{Subject to: } h_j(\mathbf{X}) \leq b_j \quad \mathbf{X} \geq 0$$

The inequality constraints are written as "less than or equal" because this form




FIGURE 3.2

The feasible region defined by an inequality constraint is larger than that defined by an equality constraint.

leads to a simple statement of the Kuhn-Tucker conditions, adequate to the purpose of this chapter.

The real inequalities can always be generalized to this form. In practice, constraints are often "more than or equal" to some quantity. New York City's water pressure standard is that way:

$$\text{pressure} \geq 40 \text{ psi}$$

This can also be written as

$$-\text{pressure} \leq -40 \text{ psi}$$

Formally, a constraint that requires an expression to be "less than or equal" to some amount is an *upper bound* in that it prescribes a maximum. Conversely, a "greater than or equal" constraint is a *lower bound*.

The strategy for solving the problem with inequality constraints is simple: we transform it into a problem we know how to solve. This is exactly the same path that Lagrange took in developing the Lagrangean; by means of some additional variables he changed the constrained problem to one that looked like an unconstrained problem.

The key to the solution for the problem with inequality constraints lies in changing these inequalities to equalities. This is done by introducing new variables, S_j , defined as the amount required to make up the difference between the constraint and the value of $h_j(\mathbf{X})$. They are thus called *slack variables*. The inequalities can then be written as equalities:

$$h_j(\mathbf{X}) + S_j^2 = b_j$$

where the slack variable is introduced as a squared term to ensure that it is a positive quantity.

The solution then consists of solving a Lagrangean to the revised problem:

$$\text{Optimize: } g(\mathbf{X})$$

$$\text{Subject to: } h_j(\mathbf{X}) + S_j^2 = b_j \quad \mathbf{X} \geq 0$$

This function is called the *generalized Lagrangean* and is

$$L = g(\mathbf{X}) - \sum \lambda_j (h_j(\mathbf{X}) + S_j^2 - b_j)$$

This equation involves $(i + 2j)$ unknowns with the addition of the j slack variables. The optimization proceeds as with standard Lagrangeans except that we now have j additional equations relating the Lagrangean and the slack variables. The new equations introduced in the generalized Lagrangean are

$$\frac{\partial L}{\partial S_j} = 2\lambda_j S_j = 0 \quad \text{all } j$$

They imply in general that either $\lambda_j = 0$ or $S_j = 0$. For this reason they are known as the *complementary slackness* conditions.

These complementary slackness conditions have important consequences for analysis and design. Their interpretation is as follows:

- If the optimal solution \mathbf{X}^* lies right on the constraint, $h_j(\mathbf{X}^*) = b_j$, then $S_j^* = 0$ and the shadow price $\lambda_j^* \neq 0$. In this case there is some value, in terms of better $g(\mathbf{X})$, to changing the constraint. Technically, we say this constraint is *binding*, and that the objective can be increased by changing this constraint.
- If the constraint is not binding, $h_j(\mathbf{X}^*) \leq b_j$, then $S_j^* \neq 0$ and $\lambda_j = 0$. That is, the nonbinding constraints do not affect the value of the objective $g(\mathbf{X})$.

These ideas are pursued in detail in Chapter 6, which discusses sensitivity analysis, and particularly in Section 6.2.

The *Kuhn-Tucker conditions* themselves are the entire set of equations associated with the generalized Lagrangean:

$$\frac{\partial L}{\partial X_i} = 0 \quad \text{all } i$$

$$\frac{\partial L}{\partial \lambda_j} = 0 \quad \text{all } j$$

$$\lambda_j S_j = 0 \quad \text{all } j$$

Given the sign convention adopted with the original statement of the problem, we can also include nonnegativity conditions on the shadow prices:

$$\lambda_j \geq 0 \quad \text{all } j$$

Since $\lambda_j = \partial g^*(\mathbf{X})/\partial b_j$, and since increasing b_j increases the feasible region and can thus only make $g(\mathbf{X})$ bigger, then λ_j must be positive. If we had a minimization problem or lower bound constraints, different restrictions on the shadow prices would exist. These are, in effect, part of the second-order conditions of optimality.

A more compact form of the Kuhn-Tucker conditions can be obtained by combining the complementary slackness and the constraint equations, $\lambda_j S_j = 0$ and $\partial L/\partial \lambda_j = 0$, into

$$\lambda_j [h_j(\mathbf{X}) - b_j] = 0$$

This is obviously true if $\lambda_j = 0$. It is also true if $\lambda_j \neq 0$ since then $S_j = 0$, and so $h_j(\mathbf{X}) = b_j$. This form obscures the meaning of the conditions; it is, however, the standard form the reader is likely to encounter elsewhere.

3.5 APPLICATIONS

Textbook. The following box provides a straightforward application of the Kuhn-Tucker conditions to a simple mathematical problem. Two points deserve particular attention. Notice first that the different combinations of ways of meeting the complementary slackness conditions must each be examined for feasibility, and then for optimality if feasible. Secondly, note that the restrictions on the shadow prices must be taken into account to determine which feasible solutions are optimal.

Example Use of Kuhn-Tucker Conditions

Consider the problem:

$$\begin{aligned} \text{Maximize: } & g(\mathbf{X}) = X_1 + 2X_2 \\ \text{Subject to: } & h_1(\mathbf{X}) = (X_1 - 1)^2 + (X_2 - 2)^2 \leq 5 \\ & h_2(\mathbf{X}) = X_1 \leq 4 \end{aligned}$$

The inequalities become equalities by means of the slack variables:

$$\begin{aligned} h_1(\mathbf{X}) &= (X_1 - 1)^2 + (X_2 - 2)^2 + S_1^2 = 5 \\ h_2(\mathbf{X}) &= X_1 + S_2^2 = 4 \end{aligned}$$

The corresponding Lagrangean is

$$L = (X_1 + 2X_2) - \lambda_1[(X_1 - 1)^2 + (X_2 - 2)^2 + S_1^2 - 5] - \lambda_2[X_1 + S_2^2 - 4]$$

The Kuhn-Tucker conditions are

$$\begin{aligned} \frac{\partial L}{\partial X_1} &= 1 - \lambda_1(2)(X_1 - 1) - \lambda_2 = 0 \\ \frac{\partial L}{\partial X_2} &= 2 - \lambda_1(2)(X_2 - 2) = 0 \\ h_1(\mathbf{X}) &= 5 \\ h_2(\mathbf{X}) &= 4 \\ \lambda_1 S_1 &= \lambda_2 S_2 = 0 \\ \lambda_1, \lambda_2 &\geq 0 \end{aligned}$$

The complementary slackness conditions offer the possibility of various combinations of λ_j , S_j being equal to zero. These must be investigated. Thus:

- If $S_2 = 0$, then $X_1 = 4$ from $h_2(\mathbf{X})$. However, this is infeasible from $h_1(\mathbf{X})$. Therefore, $S_2 \neq 0$ and $\lambda_2 = 0$.
- Since $\partial L / \partial X_2 = 0$ and $\lambda_1 \neq 0$, therefore $S_1 = 0$.

The first two conditions can thus be solved as

$$\lambda_1 = \frac{1}{2(X_1 - 1)} = \frac{2}{2(X_2 - 2)}$$

Thus, $X_2 = 2X_1$.

Substituting the previous in $h_1(\mathbf{X})$, knowing that $S_1 = 0$, we obtain

$$(X_1 - 1)^2 + (2X_1 - 2)^2 = 5(X_1 - 1)^2 = 5$$

Thus, $X_1 = 0$ or 2 .

These solutions lead to

$$\lambda_1 = -\frac{1}{2} \quad \text{or} \quad \frac{1}{4}$$

However, since $\lambda_j \geq 0$, the only feasible solution is

$$\lambda_1^* = \frac{1}{4} \quad X_1^* = 2 \quad X_2^* = 4 \quad g(\mathbf{X})^* = 10$$

Note that if we used $\lambda_1 = -\frac{1}{2}$, we would obtain

$$X_1 = 0 \quad X_2 = 0 \quad g(\mathbf{X}) = 0$$

which is not optimal.

Real-world. The problem of determining what price to charge for the products of a system illustrates the use of the generalized Lagrangean and the Kuhn-Tucker conditions in a practical setting. This is a most important issue in systems design. Although pricing may appear to be an economic issue quite distinct from engineering, the fact is otherwise. The demand for the products of a system, and thus the load on a system, depends on their prices. As it is not possible to design a system without considering its load, a full systems design, correctly done, will consider the prices that will be charged for the products of the system being designed.

The specific problem to be considered concerns the pricing of air transportation in an uncongested, developing area. It is drawn from a large study done by de Neufville and Mira (1975) for the United States and Colombian governments.

The problem assumes that the system uses a number of aircraft flights F , to transport a number of people, P . The value $V(\bullet)$ of the system, to the society as a whole, is a positive function of both F and P : it is better to serve more people and to have the greater convenience of more flights and higher frequency. On the other hand the system has a cost, taken to be the number of flights times their average cost, c . The net value of the system, the objective function $g(\mathbf{X})$, is then:

$$g(\mathbf{X}) = V(F, P) - Fc$$

The constraint on the system is that the capacity of the aircraft, C , cannot be exceeded by the number of persons per flight:

$$\frac{P}{F} \leq C$$

The preceding two equations constitute the statement of the problem, where net value is to be maximized.

The generalized Lagrangean is

$$L = V(F, P) - Fc - \lambda \left(\frac{P}{F} + S^2 - C \right)$$

so that the Kuhn-Tucker conditions are

$$\frac{\partial L}{\partial F} = V_F - c + \frac{\lambda P}{F^2} = 0$$

$$\frac{\partial L}{\partial P} = V_P - \frac{\lambda}{F} = 0$$

$$\frac{\partial L}{\partial \lambda} = \frac{P}{F} + S^2 - C = 0$$

$$\frac{\partial L}{\partial S} = 2S\lambda = 0$$

where the subscripted variable indicates the partial derivative with respect to that variable.

The solution is obtained by substitution. First we can recognize that

$$\lambda = FV_P \neq 0$$

since neither the number of flights nor the incremental value per passenger equals zero. This then implies that $S = 0$ and $P/F = C$: it would be better for society overall if planes flew full and empty seats were not wasted. Substituting the value of the shadow price in the first equation leads to

$$V_F + CV_P = c$$

Economic theory indicates that the price, p , charged for a service should be its marginal value to people, $p = V_P$. Therefore, the analysis defines the optimum price to charge a person for a flight is

$$p^* = V_P = \frac{1}{C} \{c - V_F\}$$

The implications of the formal analysis are quite interesting. Firstly, we may see that when there is a positive value to greater frequency, $V_F > 0$, as there is in a developing area, the optimum price for a ticket is less than the average cost:

$$p^* = \frac{1}{C} \{c - V_F\} \leq \frac{c}{C}$$

A subsidy to the system is thus in order to ensure that, from the point of view of society, adequate transportation services tie the region together. The optimal amount of this subsidy per flight is the difference between the average cost and the amount paid per flight:

$$\text{optimal subsidy} = c - Cp^* = V_F$$

the value to society of the extra flight.

When there is sufficient frequency so that $V_F = 0$, then the optimal price for a ticket is its average cost. Finally, if the air traffic system is saturated so

that extra flights cause delays, $V_F \leq 0$, then the optimal price is greater than its average cost: the optimal policy is to tax the system to dampen demand, lessen congestion, and improve the overall performance of the system.

REFERENCES

- Baumol, W. J., (1977). *Economic Theory and Operations Analysis*, 4th ed., Prentice-Hall, Englewood Cliffs, NJ.
- de Neufville, R., and Mira, L.-J., (1974). "Optimal Pricing for Air Transport Networks," *Transportation Research*, September, pp.1-11.
- Winston, W.L., (1987). "Non-Linear Programming," Ch. 11, *Operations Research: Applications and Algorithms*, PWS Publishers, New York.

PROBLEMS

3.1. Constrained Optimization I

Given the problem:

Maximize: profit = XY \$

Subject to: $X + 4Y \leq 12$ (units of resource) $X, Y \geq 0$

- Formulate the problem, state the optimality conditions, and solve for the optimal values of X, Y and profit.
- If you were offered another unit of resource for \$8, would you buy it? Why?
- What is the maximum you would pay for a unit of resource?

3.2. Constrained Optimization II

Given the problem:

Maximize: $Z = 2X + 3XY$

Subject to: $XY + X^2/9 = 13$ $X, Y \geq 0$

- Formulate the problem, state the optimality criteria and solve.
- What is the practical significance of the Lagrangean multiplier λ ?
- How would the formulation change if the constraint were $XY + X^2/9 \leq 13$?

3.3. Constrained Optimization III

Formulate the problem, state the optimality criteria and solve:

(a) Maximize: $Z = XY$

Subject to: $X + Y^2 \leq 48$ $X, Y \geq 0$

(b) Maximize: $Z = XY$

Subject to: $X + Y^2 = 12$ $X, Y \geq 0$

(c) Maximize: $Z = X^2 + 3Y^2$

Subject to: $X + Y^2 \leq 10.5$ $X, Y \geq 0$

3.4. Constrained Optimization IV

Solve the following optimization problem by Lagrangean analysis, not graphically. Clearly indicate the value of all variables and the objective function at the optimum.

Maximize: $Z = 4ab^2$

Subject to: $3a + b \leq 10$ $a, b \geq 0$

3.5. Road Work

Assume that the number of highway miles that can be graded, H , is a function of both the hours of labor, L , and of machines, M :

$$H = 0.5L^{0.2}M^{0.8}$$

- (a) Minimize the cost of grading 20 miles, given that the hourly rates for labor and machines are: $C_L = \$20$; $C_M = \$160$.
 (b) Interpret the significance of the Lagrangean multiplier.

3.6. Power Plant

A cooperative operates an oil-fired power plant, selling electricity and steam to its members. The value of a unit of electricity, e , is a constant P_e per unit, which is its price from an alternative supplier. The value of steam, s , is a monotonically decreasing function $P_s(s)$. The cooperative wishes to maximize the net benefits of its members:

$$Z = P_e(e) + \int_0^s P_s(s) ds - P_q Q$$

where P_q and Q are the unit price and quantity of oil used in the plant.

The technical constraints on the operation are that the electricity and steam produced are less than their thermal efficiency times the quantity of oil used:

$$e \leq E_e Q \quad s \leq E_s Q$$

and that maximum production of electricity and steam are incompatible:

$$\frac{e}{E_e} + \frac{s}{E_s} \leq 1.5Q$$

- (a) Formulate the problem and state the Kuhn-Tucker conditions.
 (b) By setting various combinations of the $\lambda_i = 0$, eight potential solutions are logically possible. Examine each to determine which are feasible, and explain their significance.

3.7. Tim Burr, Jr.

See Problem 2.10. Junior, happy with his father's identification of significant returns to scale, is all set to maximize production to take advantage of their benefits. Do so for him, with a budget B and input costs for pine and balsam proportional to the $\frac{5}{4}$ power of their quantity—since Tim can only get more lumber by acquiring it further away. Do his economics improve as he builds a larger plant?

3.8. Heat Exchanger

160 m of tubes must be installed in a heat exchanger to provide the necessary surface area. The cost of the installation is the sum of the cost of the tubes, the shell of the exchanger, and of the floor space it occupies:

$$\text{Cost} = 700 + 25 D^{2.5} L + 12.5 DL$$

where D and L are the diameter and the length of the cylindrical exchanger. No more than 20 tubes can fit into 1 m² of cross-section. Find the dimensions that minimize cost.

CHAPTER 4

MARGINAL ANALYSIS

4.1 CONCEPT

Marginal Analysis is a basic form of optimization of design. It is a means of selecting the best choice from among many technically efficient ways to achieve a stated objective or product.

Marginal analysis combines two sorts of models, as all procedures for optimizing design must. One model represents the technical possibilities; the other represents the relative values of the several inputs of the production process. Specifically, marginal analysis combines the production function, which represents only the technically efficient production possibilities; and the *input cost function*, which describes the cost of the inputs used (see Section 4.2).

Marginal analysis is based on three specific assumptions. The first is that the feasible regions are convex for the portions of both models being considered. A key word in this statement is "portion." As Sections 4.2 and 4.3 explain, this qualification permits the application of marginal analysis even for processes that have increasing returns to scale, and for which the entire feasible region is not convex.

The second assumption of marginal analysis is that the only constraint on the system is the amount of money available, the budget. The resources themselves are presumed to be available indefinitely, provided there is enough money to buy them. This situation is quite common in practice, since the materials required for any one system are generally far less than the supply available. A company setting out to build a factory can thus presume that there will be as much steel and concrete available for purchase as it might require, and that it can thus optimize its design for the factory as if the supply were unconstrained.