

Chapter 1, Operations Research (OR)

Kent Andersen

February 7, 2007

The term “Operations Research” refers to “research on operations”. In other words, the study of how to “operate something” in the best possible way.

A key idea in OR is to extract “the essentials” of a practical problem, and to formulate a mathematical model that represents a practical problem.

This, of course, has limitations, but the approach has nevertheless proved effective in a wide range of applications.

OR has been applied in areas such as manufacturing, transportation, railway scheduling, airline crew scheduling, construction, telecommunication, financial planning and in the military - to name just a few.

Historically, OR started during World War II, where the US military used mathematical models to plan its operations.

OR is now routinely applied in both the private sector and the public sector. Two important factors have played an important role in the development of OR:

- (a) The invention of the computer: this has made it possible to address very large and complex practical problems.
- (b) The invention of a method for solving linear programs (this method will be presented later): This invention led to huge amounts of research on how to solve mathematical models in general. The methods for solving linear programs are still used for solving practical problems today.

1 Main steps in an OR approach

We now outline the main steps that are used in “an OR approach” to address a practical problem:

- (1) Define the problem of interest, and the data that is needed to describe the relevant part of the problem. This step is usually performed through communication with the management of the organization, or with the relevant people of a logistics department of a larger organization.
- (2) Design a mathematical model that captures the main aspects of the problem. It is important that the model does not ignore important aspects of the problem, so the management of the organization must, at least, accept the simplifications involved in the model.

- (3) Develop appropriate methods for deriving solutions to the problem. This step is often performed by OR experts.
- (4) Test the model: The methods developed produce some initial, suggestive, solutions to the problem, and this often leads to a realization of the limitations of the model. Hence, this step often leads to modifications of the model, and a repetition of steps (1)-(4).
- (5) Once an appropriate model is agreed upon, a number of methods are chosen. This choice is based on a trade-off between the quality of the solutions obtained, and the time needed to obtain them.
- (6) Implement the chosen methods (on a computer), and construct software for practical usage. In other words, the software must be easy to use for the end user, which is typically not an expert in OR.

Some of the key questions that need to be answered are as follows:

- (a) Variables: Which numbers need to be known in order to have a solution to the problem? For example, in resource planning, the amount of resource i which is devoted to activity j could be the variable of choice. Knowing the values of all these variables would then give an allocation of resources to activities. Another example could be in a time-table construction problem, where the variables could answer the question “who works at what time?”. The result is a number of variables x_1, x_2, \dots, n whose values are supposed to describe “a solution”.
- (b) Constraints: It is likely not all values of the variables that gives a useful solution. For instance, there is likely a limit on the amount of resources that every solution has to respect. For instance, in the construction of a time table, there is likely a limit on how long time an individual can work, and also, each individual probably needs a break once in a while.
- (c) Objective: In the case when there are several possible solutions, how is the quality of a solution measured? Typically this is either the “cost” of a solution, which one wants to minimize, or the “profit” of a solution, which is then maximized. However, things can be more complex than that. For instance, in train scheduling, a schedule where the trains are as close to being “on time” as possible might involve two trains that cross the same physical point almost at the same time, and this might be considered very risky for safety reasons.
- (d) Robustness and uncertainty: How likely is it that the input data of the model changes, and what are the consequences? A solution that is “cheap” might become very expensive if the data changes slightly, and this would, needless to say, be very unfortunate. In the construction of a time table, what happens if a person becomes sick? Does this ruin a given schedule? Or what happens to a schedule in airline crew scheduling if all the pilots in a given city decide to go on a strike?

The outcome of the above steps is often an optimization problem:

$$\begin{aligned} \min \text{ or } \max f(x) \\ \text{such that } x \in F \end{aligned}$$

where:

x is the variable (typically a vector) that defines a solution.

$f(x)$ is a function that describes the quality of a solution.

F denotes the set of vectors x that satisfy all the constraints.

The above formalization (model) then says: find the variable vector x that satisfies all constraints of the problem ($x \in F$, where “F” stands for “Feasibility”), and has the best possible objective function value $f(x)$.

The point of steps (1)-(6) in the “OR procedure” is then to formalize the variables x , the function $f(x)$ and the set F in such a way that a “best possible x ” for the mathematical problem also makes sense in a practical setting. Also, it is important that the data needed to define the optimization problem can be supplied easily, and that a solution x to the optimization problem is delivered in a form that can be easily understood by an “end user”.

Given a model, the next step is to develop methods for finding “good” solutions for the problem (where “good” is defined by the objective function $f(x)$). The main issue for choosing such methods is often time:

- (i) How much time can be allowed between the point in time when the data is supplied by the “end user” and the time when a solution is supplied by the software. In other words, how much time is allowed for finding a good solution to the optimization problem.
- (ii) How much time is available for developing methods for finding good solutions?

The following categorization is often used for applications:

On-line problems: In these applications of OR, solutions are required to be delivered very fast. Imagine a computer program, where a user clicks a button, and then waits for a solution.

Strategic problems: For such problems, more time is available. An example could be the construction of a yearly plan for a production company. Another example could be the construction of a monthly plan for an airline company (which planes and which pilots should service which flights).

Given the importance of being able to provide good solutions to optimization problems, a lot of research has been done on this issue. Optimization problems are categorized according to their characteristics:

Linear programs: The objective function ($f(x)$) and all constraints (the set F) can be described with linear functions. This was the pioneering optimization problem for OR, and we are going to study this problem intensively in this course.

Integer programs: Some variables (or all) have to be integers in a solution. An important special case is a *decision variable*, which must be either zero or one. In the construction of a time table for, say, a hospital, such a variable could decide “whether or not Paul works on Tuesday from 5 to 6”.

Non-linear programs: This simply means that the problem definition uses some functions that are not linear.

Although software is available in all of the above three categories, a direct application is rarely feasible for a given “practical” problem, and it is therefore necessary to design problem specific software.

Once some methods have been suggested for generating solutions, the first, suggestive, solutions become available. This, very often, leads a re-consideration of the underlying model, and often this leads to a change in the model. In other words, it is important to remember that a model is not a constant. A model should be re-considered regularly.

2 Linear Programming

We now introduce *Linear Programs* (LP's). As mentioned, these problems are characterized by the fact that they can be described completely by using linear functions only. We start by considering an example.

Example 1 (*The Wyndor Glass Company (WGC)*)

The WGC produces two types of glass products - windows and glass doors. These products are produced in three plants. Plant 1 produces aluminum frames and hardware, Plant 2 makes wood frames and Plant 3 produces glass and assembles the products.

To increase their profits, the WGC has decided to introduce two new products:

- (1) *An 8-foot glass door with aluminum framing.*
- (2) *A 4 × 8 foot wood-framed window.*

Product 1 requires production capacity from Plant 1 and Plant 3, and Product 2 requires production capacity from Plant 2 and Plant 3. The marketing department has concluded that the WGC can sell as much as is produced of the two products (they expect the products to be popular among the customers). However, Product 1 and 2 are competing for the same production capacity at Plant 3, so the WGC is not sure about how to mix the two products. We have been asked to analyze this problem for the WGC.

Our discussions with management has led to the following definition of the problem:

- (i) *Management would like to know the production rate of each of the two products. The production rate is defined to be the number of batches to produce of each product per week, where each batch has size 20.*
- (ii) *They would like to choose the production rate so as to maximize the profits for the WGC.*

Each plant has a limited number of production hours available per week to produce the two products. The WGC has other products, and a large amount of the production hours are tied up on the production of these products. Plant 1 has 4 hours available for the two new products, Plant 2 has 12 hours available, and Plant 3 has 18 hours available. These numbers were obtained by talking to the management of the three plants.

To solve the problem of the WGC, we realized that the following questions need to be answered:

- (i) *How many production hours does it require to produce one batch of each product?*
- (ii) *How much profit does one batch of each product generate?*

After discussing these questions with management, we were pointed to key staff members that could answer the questions. The marketing department informed us that they expected one batch of product 1 to generate a profit of \$3000, and each batch of product 2 was expected to generate a profit of \$5000.

To answer question (i), we were pointed to staff in the manufacturing division. They informed us that one batch of Product 1 required 1 production hour at Plant 1 and 3 production hours at Plant 3. To produce one batch of product 2, it would require 2 production hours at Plant 2 and 2 production hours at Plant 3.

With the above data, we realized that we now had enough information to formulate the problem of the WGC as a Linear Program (LP). We chose the following variables:

$x_1 =$ *number of batches of product 1 to produce each week*

$x_2 =$ *number of batches of product 2 to produce each week*

$Z =$ *total profit per week (measured in thousands of dollars) obtained from producing the two products*

The variables above are the decision variables of the WGC problem, and our task is to determine the best values for these variables so as to maximize profits. Given x_1 and x_2 , the value of Z can be calculated as $Z = 3x_1 + 5x_2$. Since the plants have a limited number of production hours, the assignment of values to x_1 and x_2 needs to respect these capacities. Given the values of x_1 and x_2 , the total consumption of production hours at Plant 1 is x_1 (one batch of Product 1 requires one production hour at Plant 1, and Product 2 does not pass through Plant 1). Hence, for x_1 and x_2 to respect the capacity of 4 hours at Plant 1, we must have $x_1 \leq 4$. Similarly, for x_1 and x_2 to respect the capacity at Plant 2, we must have $2x_2 \leq 12$, and finally, the capacity of 18 at Plant 3 means that x_1 and x_2 must satisfy $3x_1 + 2x_2 \leq 18$.

Naturally, since the variables x_1 and x_2 are supposed to represent production rates, they must be non-negative, or in other words, x_1 and x_2 must satisfy $x_1 \geq 0$ and $x_2 \geq 0$. Written as a Linear Program, the problem of the WGC is to

$$\begin{aligned} & \text{Maximize } Z = 3x_1 + 5x_2 \\ & \text{subject to the restrictions that} \\ & \qquad x_1 \leq 4 \qquad \qquad \text{(Capacity of Plant 1)} \\ & \qquad 2x_2 \leq 12 \qquad \text{(Capacity of Plant 2)} \\ & \qquad 3x_1 + 2x_2 \leq 18 \qquad \text{(Capacity of Plant 3)} \\ & \qquad x_1 \geq 0, x_2 \geq 0 \qquad \text{(Non-negativity)} \end{aligned}$$

2.1 Graphical solution

Since our problem only involves two variables, it is possible to use a graphical procedure to solve it. Consider the graph in Fig. 1.(a). In the graph, the lines $x_1 = 4$, $2x_2 = 12$ and $3x_1 + 2x_2 = 18$ have been drawn. We now identify the values of x_1 and x_2 (the solutions) that satisfies the restrictions of our linear program. Since we must have $x_1 \geq 0$ and $x_2 \geq 0$, all solutions must lie on the positive side of the axes (including actually *on* the axes). Furthermore, the restriction $x_1 \leq 4$ means that (x_1, x_2) must be on the left side of the line $x_1 = 4$, and similarly, (x_1, x_2) must be below the line $2x_2 = 12$. Finally, the restriction $3x_1 + 2x_2 \leq 18$ gives that (x_1, x_2) must be below the line $3x_1 + 2x_2 = 18$. We conclude that the pairs (x_1, x_2) that define a solution to our linear program must be in the shaded area of Fig. 1.(a).

To solve the linear program, we therefore have to answer the question: which of the solutions in the shaded area of Fig. 1.(a) maximize $Z = 3x_1 + 5x_2$. Consider Fig. 1.(b), where we, in addition, have drawn the lines $Z = 20 = 3x_1 + 5x_2$ and $Z = 36 = 3x_1 + 5x_2$. Observe that the points (x_1, x_2) on the line $Z = 20 = 3x_1 + 5x_2$ that lie *within* the shaded region are all feasible solutions to the linear program and have objective function value equal to 20. Also observe that, if we increase the value of Z from 20, the line $Z = 3x_1 + 5x_2$ moves upwards. To solve our linear program, we therefore need to move the line $Z = 3x_1 + 5x_2$ as high as possible such that the line still intersects the shaded area. This is achieved at the value $Z = 36$, where the only intersection point between the line $Z = 3x_1 + 5x_2$ and the shaded area is the point $(x_1^*, x_2^*) = (2, 6)$. The best solution for the WGC is therefore to produce two batches of Product 1, and six batches of Product 2. We therefore report this solution to the management of the WGC, which it can then implement, and maximum possible profit is $Z = 36$.

Observe that the lines $Z = 20 = 3x_1 + 5x_2$ and $Z = 36 = 3x_1 + 5x_2$ are parallel. This is not a coincidence. *Any* line of the form $Z = 3x_1 + 5x_2$ can be written as $x_2 = -\frac{3}{5}x_1 + \frac{1}{5}Z$, and therefore *all* lines of this form have a slope of $-\frac{3}{5}$.

The above procedure is often referred to as the *graphical method*. It can be used to solve *any* linear program in two variables. It can be extended to linear programs with three variables, but this is the limit. We will later present a more general method, called *the simplex method*, which can be used for any number of variables.

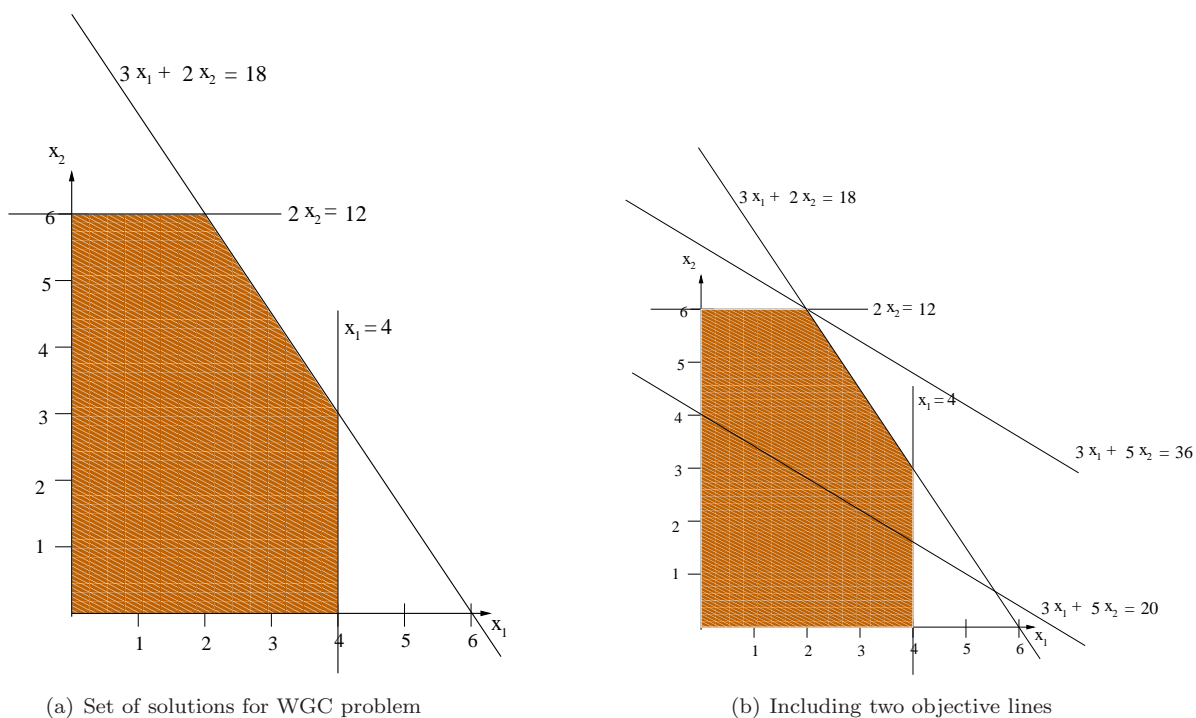


Figure 1: The WGC problem in two dimensions

We are not done with the example of the Wyndor Glass Company. As mentioned, all six phases of the OR procedure needs to be performed in a well-conducted OR study. In particular, we will do a post-optimality analysis, or in other words, we will examine the sensitivity of our *model* to the data (the numbers that are used in defining the linear program).

2.2 The general linear programming model

The Wyndor Glass Company (WGC) problem illustrates a typical linear program (in a miniature version). We now define and discuss how a general linear program looks like.

In the WGC problem, there were three resources (The number of production hours at each plant) and two activities (the two new products), and the activities were competing for the resources.

Now suppose we have n activities and m resources. Typical resources could be money, machines, equipment, vehicles and personnel. Typical activities could be amount of investment in a given project, amount of advertising to put in particular media, and the amount of good to ship from a given source to a given destination. As in the WGC example, there is also a profit of choosing one unit of a particular activity.

The linear program is then the problem of choosing levels for each of the n activities such that none of the m resources are over used. The data required to formulate this linear program (which should be supplied by the organization that is interested in an optimal solution to the problem) are given as follows. This notation is fairly standard, and will be used intensively in the next chapters. The activities are indexed $j = 1, 2, \dots, n$, and the resources are indexed $i = 1, 2, \dots, m$. Furthermore:

- c_j = profit from each unit of activity j
- b_i = amount of resource i that is available for allocation among the activities
- $a_{i,j}$ = amount of resource i consumed by each unit of activity j

Recall that, for a linear program, the assumption is that everything is *linear*. Therefore, if activity j is chosen to be at level $l_j > 0$, then the contribution to the profit of activity j is given by $l_j c_j$, and activity j uses $l_j a_{i,j}$ of resource i . This assumption is definitely *not* necessarily true for all OR applications.

The variables needed for the general version of the WGC example are as follows:

- x_j = level of activity j
- Z = The overall profit obtained from a given choice of levels for the activities

Since our assumption is linearity, the total profit can be obtained as the sum of the contributions to the profit from each activity. In other words, we have

$$Z = c_1 x_1 + c_2 x_2 + \dots + c_n x_n.$$

Similarly, the total consumption of resource i , where $i \in \{1, 2, \dots, m\}$, can be calculated as the sum of the amounts consumed by each activity. In other words, given the levels x_1, x_2, \dots, x_n of the activities, the consumption of resource i is given by

$$a_{i,1} x_1 + a_{i,2} x_2 + \dots + a_{i,n} x_n.$$

For a given choice x_1, x_2, \dots, x_n of levels for the activities, the availability of resource i for $i = 1, 2, \dots, m$ must be respected. In other words, we must have $a_{i,1} x_1 + a_{i,2} x_2 + \dots + a_{i,n} x_n \leq b_i$ for $i = 1, 2, \dots, m$. The mathematical model for allocating resources to activities can now be formulated as follows. Choose values for x_1, x_2, \dots, x_n so as to:

$$\begin{aligned} & \text{Maximize } Z = c_1 x_1 + c_2 x_2 + \dots + c_n x_n \\ & \text{subject to the restrictions that} \\ & a_{i,1} x_1 + a_{i,2} x_2 + \dots + a_{i,n} x_n \leq b_i \qquad \text{for } i = 1, 2, \dots, m \\ & x_1 \geq 0, x_2 \geq 0, \dots, x_n \geq 0 \qquad \text{(Non-negativity)} \end{aligned}$$