Chapter 1

Introducing the Simplex Algorithm

1.1 Constrained Optimisation

Almost the whole of this course is concerned with the following general question:

> Find the maximum, or minimum value of a function $f = f(x_1, \ldots, x_n)$ of n real variables subject to the *constraints* that $g_i(x_1, \ldots, x_n) \geq 0$ for $i = 1, \ldots, m$.

Such a problem will be referred to as a **constrained optimisation** problem, and the corresponding value as the **optimum value**. The function we are to optimise is called the **objective function**. The set of vectors $\mathbf{x} = (x_1, \ldots, x_n)$ which satisfy the constraints is known as the **feasible region**. A vector at which the objective function attains its optimal value is known as an **optimal feasible vector**. Of course there is no guarantee that an optimal feasible vector exists for a given problem.

Note first that there is no need to deal with maximising and minimising separately, since finding a minimum of $f(x_1, \ldots, x_n)$ is the same as finding a maximum of $-f(x_1, \ldots, x_n)$. Note also that if we have two constraints $g_1 \geq 0$ and $g_2 \geq 0$ and in addition we know that $g_2 = -g_1$, then $g_1 = 0$; in other words an equality constraint can be reformulated as a pair of "$\geq$" constraints.

As a simple example, consider the problem of maximising $x^2 - 5x + 6$ for $1 \leq x \leq 2$. Such a problem is familiar from first year calculus; it can be written as a constrained optimisation problem by writing $f(x) = x^2 - 5x + 6$ and defining constraints $g_1(x) = x - 1 \geq 0$ and $g_2(x) = 2 - x \geq 0$.

In the first part of the course, we concentrate on a special case of the problem in which both the objective function, and the constraints are *linear* functions. Such a problem is known as a **Linear Programming Problem**, or linear optimisation problem,

1.2 Some Sample Problems

We give a number of different examples of problems with linear constraints, aiming to show that a large class of "interesting" problems are of this type.

1.2.1 Maximising a Function of Two Variables with Constraints

Our first example can be solved quite simply by geometric methods.

The problem is to maximise $x+y$ subject to the constraints that $2x+y \leq 8$, $x+2y \leq 7$ and $x-y \geq -2$. This is illustrated in Fig. 1.1. The feasible region is shaded, and consists of the portion in the first quadrant (so $x \geq 0$ and $y \geq 0$), which lies below each of the three thick lines. These lines, $y = x+2$, $y+2x = 8$ and $2y+x = 7$ correspond to when the constraint becomes **tight**. The objective function - the thing we are trying to maximise is $x+y$, and the three parallel lines represent lines $x+y = k$ for different values of k. The largest value of k is attained at a point within the feasible region when the line just touches the vertex M; in the diagram, this corresponds to the largest value of k, with increasing values of k corresponding to lines which are further up and to the right.

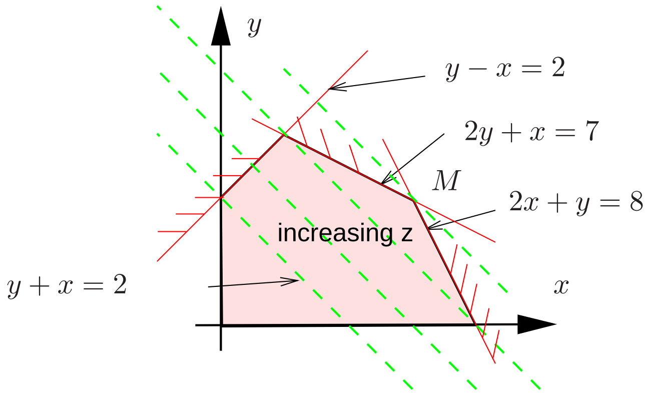

Figure 1.1: Two dimensional optimisation problem.

1.2.2 Machine Shop Scheduling

Our next task is to show that problems of this sort can occur in "practical" situations; that problems of interest in the "real world" lead to this type of *constrained optimisation* problem.

A machine shop makes two products called (rather unimaginatively) A and B. Product A can be made with two options — as A_1 and A_2, while product B is available in options B_1, B_2 and B_3. The machine shop makes the two products using an appropriate combination of three machines, which can be used in any order. The production contract requires that 60 units of item A and 85 units of item B be produced per week, although they can be produced in any of the various options. The objective of the exercise is to determine the product mix that is most profitable. The situation is summed up in Table 1.1.

In order to write down in detail what is required, we need to introduce suitable variables.

> *Choosing variables is often the hardest part of the whole process. One way is to think what you need to know in order to solve the problem — give the orders or instruct the foreman. Such variables are often known as* ***decision variables*** *because knowing their values enables a decision to be made. In this case, the decision is "how many of each option of each product do we make each week?"*

It is thus natural to introduce the following variables. Let x_1 be the number of units of product A_1 to be produced per week, x_2 be the number of units of product A_2 to be produced per week, x_3 be the number of units of product B_1 to be produced per week, x_4

Product	Option	Unit production time on machine number			Unit Profit
		1	2	3	
A	1	0.5	-	0.2	2
	2	-	0.4	0.2	2.5
B	1	0.4	0.3	-	5
	2	0.4	-	0.3	4
	3	-	0.6	0.3	4
Hours per week that machines are available		38	31	34	

Table 1.1: Machine shop costs.

be the number of units of product B_2 to be produced per week and x_5 be the number of units of product B_3 to be produced per week.

The profit from such a product mix is given by

$$P = 2x_1 + 2.5x_2 + 5x_3 + 4x_4 + 4x_5,$$

and this is the function (of $x_1, x_2, \dots, x_5$) that we wish to maximise. The constraints are of three sorts:

$$\begin{aligned} x_1 + x_2 &= 60, \\ x_3 + x_4 + x_5 &= 85, \end{aligned} \qquad \text{(Required production)}$$

$$\begin{aligned} 0.5x_1 + 0.4x_3 + 0.4x_4 &\leq 38, \\ 0.4x_2 + 0.3x_3 + 0.6x_5 &\leq 31, \\ 0.2x_1 + 0.2x_2 + 0.3x_4 + 0.3x_5 &\leq 34, \end{aligned} \qquad \text{(Machine time)}$$

$$x_i \geq 0 \quad \text{for each } i. \qquad \text{(Reality)}$$

Solving this constrained optimisation problem then gives the values of $x_1, x_2, \dots, x_5$ which give the most profit for this particular contract.

1.1. Remark. Much of the remainder of the course is devoted to solving such problems. When you can, and have enough facility with MAPLE, come back to this problem. You should find that the problem is feasible and that the maximum profit is 520 units.

1.2.3 A Transport Problem

In the next introductory example, we give the data in purely symbolic form. As such this describes a *class* of problems, known as **transport problems**.

A firm has warehouses $W_1, W_2, \dots, W_m$ to supply retail outlets $R_1, R_2, \dots, R_n$ with a certain product. The warehouse W_i has a supply s_i of the product ($i = 1, \dots, m$), measured in some convenient units, and we assume that **all** the supply is to be shipped to the retail outlets. In doing this, a demand d_j at outlet R_j **must** be satisfied for each $j = 1, \dots n$. It is given that the cost of shipping the product from warehouse W_i to retail outlet R_j is proportional to the amount shipped, and that shipping a unit amount costs c_{ij}.

We formulate the problem of determining how much of the product should be sent from each warehouse to each retail outlet, so that all demands are satisfied, all the product is shipped, and the transportation cost is minimised.

> *Again we have the problem of choosing variables which enable us to describe a solution. And again they are suggested by the information that you would need to pass to the manager of each warehouse*

Let x_{ij} be the amount of product shipped from warehouse W_i to retail outlet R_j. The the total shipping cost is

$$C = \sum_{i=1}^{m} \sum_{j=1}^{n} c_{ij} x_{ij}.$$

Our supply constraint, that we ship all the product from warehouse W_i becomes

$$\sum_{j=1}^{n} x_{ij} = s_i,$$

while the demand constraint, to meet the demand specified at retail outlet R_j is

$$\sum_{i=1}^{m} x_{ij} \geq d_j.$$

Note we have the feasibility constraints that for each pair (i, j), $x_{ij} \geq 0$, since we must ship a non-negative amount of the product.

The problem then becomes one of minimising the cost C subject to these constraints.

1.2.4 A Blending Problem

Wine from three European countries is to be blended. We express all costs in £ (perhaps it should be Euros?), so the three wines cost respectively C_1, C_2 and C_3 per litre. The wines are to be blended and sold for d per litre. The wines have acidities α_1, α_2, and α_3, and the blended wine must have an acidity $\leq \alpha$. *Assuming* that acidity blends by volume, so that if the three wines are mixed in the proportions $x_1 : x_2 : x_3$, then, by volume, the mixture has acidity $\alpha_1 x_1 + \alpha_2 x_2 + \alpha_3 x_3$.

1.2. Remark. A separate question is whether the assumption is realistic. You could treat this as an excuse to do some practical work if you wish. Even if the assumption is false, there is a great temptation to behave as if it is true; because anything else is *very* much harder to manage. It should come as no surprise to you that most "scientific" assessment procedures behave as though their problem is linear even when it clearly isn't: do *you* think that *all* the marks on a given exam question are equally easy to get? Note that our system invariably assumes that this is the case.

To continue with the problem, the cost per litre of wine which is sold at d is $C_1 x_1 + C_2 x_2 + C_3 x_3$. Our problem is to maximise the profit, which is thus

$$d - \sum_{i=1}^{3} C_i x_i,$$

subject to the constraints that $x_i \geq 0$ and $\alpha_1 x_1 + \alpha_2 x_2 + \alpha_3 x_3 \leq \alpha$.

Another version This is the same problem, but there are only Q_i litres of each wine available. In this case, we use as variables q_i, the number of litres of each wine to be blended. The problem then becomes that of maximising

$$d\sum_{i=1}^{3} q_i - \sum_{i=1}^{3} C_i q_i$$

The limited volume gives the constraint that $q_i \leq Q_i$ for $i = 1, 2, 3$ while the acidity constraint is

$$\alpha_1 \frac{q_1}{\sum_{i=1}^{3} q_i} + \alpha_2 \frac{q_2}{\sum_{i=1}^{3} q_i} + \alpha_3 \frac{q_3}{\sum_{i=1}^{3} q_i} \leq \alpha.$$

As it stands, this last constraint is *not* linear, but can be made so by multiplying through by $\sum_{i=1}^{3} q_i$.

1.3 A More Elaborate Example

In the past, the following example has been set as continuous assessment for this class. It is presented here as a relatively realistic example of how the simplex algorithm might be used in practice. You are invited to work the various parts of the example at the appropriate time during the course. At present you can do the "formulation" part. Here then is the "story".

You have been engaged by a manufacturing company because they value your expertise in Linear Programming. This may be a little premature, but they don't need your report until after you have finished Section 3, when you will have the necessary expertise. The company is the Mendip Metals Manufacturing PLC. This example is borrowed, so I have deliberately left in the original location! And the Muchals Metals Manufacturing Company didn't have quite the same foreign ring to it! The details provided by MMM are given in question 1.3 on tutorial sheet 1.

Your eventual aim is to produce a report addressed to the MMM management, advising how much of each of the raw materials to buy, in order to maximise profits. You should also discuss the effects of changes in market conditions. You may in addition offer other advice based on your calculations. At this stage, you should show that the problem of deciding which, if any, of the alloys to manufacture, and from which raw materials, can be expressed as a Linear Programming Problem. Overall, the production of the report is divided into three parts:

- formulating the problem;
- obtaining a solution of the problem in MAPLE; and
- obtaining a useful collection of solutions, and on the basis of your results, writing the report.

There should be no mathematics in the body of the report. The mathematical formulation, and a brief statement of the solution, should go into Appendices, together with any sensitivity results you have calculated. Do *not* go into details of the simplex calculation, which should be done using MAPLE, and not by hand. The entire report will be assessed