IEEM 101: Linear Programming and Its Applications

Outline of this series of lectures:

1. How can we model a problem so that it can be "solved" to give the required solution
2. Motivation: examples of typical ***linear*** problems in IE
3. Story by pictures: a graphical interpretation of ***Linear Programs*** (LP)
4. Solving LP using the Simplex Method
    4.1. How to put the LP in "standard form" for Simplex method
    4.2. How to use a computer program to solve an LP
5. Some theoretical basics of Simplex: the beauty of convexity

---

These notes are in two parts: this part has topics 1-3 above.

Modeling of Problems

As we mentioned in our introduction, our interest is always to either find methods to *improve* a product, or a process. It is natural to ask: how much should we improve it ? Obviously, we would like to improve it as much as possible, under the give constraints of time, cost, etc.

In other words, we are always seeking for the ***maximum*** improvement. Thus we will often be faced with a situation that we are given some resources to do a certain task, and some constraints (e.g. time constraint such as deadlines). It is possible that there are several ways to do the task, but some ways are better than others. So how do we guarantee that the way we select is the "best"?

It all depends, of course, on the model that we construct. This is tricky – for one thing, we are often not very clear about how to define "best solution"; and even when we do, the same problem may be modelled in different ways, and it could be easier to "solve" one model than another. An example of the first kind of difficulty is seen in the planning of many governments: they want to minimize their expenses, but they also want to maximise employment, even if that means hiring many people to work on public sector projects; yet these two "goals" often conflict each other over the period of the planning horizon. In any case, we shall ignore such "multi-objective" problems in this course: even when there is a single, well-defined objective, there are many interesting problems.

We look at some examples to learn how problems can be modelled. In particular, we shall only look at problems whose models will turn out to be ***linear*** (I will define this soon). There are many other IE problems that are non-linear; we shall meet them later on the course.

Example (The Product Mix Problem):

A fertilizer factory manufactures two types of fertilizers, Type A is high in phosphorus and Type B is low in phosphorus. Each has three key ingredients: urea, rock phosphate and potash. These ingredients are also produced or mined by the company at other sites. Table 1 gives the information about the fertilizers.

Item	Tons Required/ ton of Fertilizer		Maximum amount available per day
	Type A	Type B	
Urea	2	1	1500
Potash	1	1	1200
Rock Phosphate	1	0	500
Net profit per ton	15	10	

Assuming that the company can sell all the fertilizer it can produce, how much of each type should it make, and how much of each raw material will it need per day?

Model:

If we know how much of each type we make per day, then it is easy to compute the raw material requirements.

Step 1. The decision variables

Let us denote the daily production of Type A by x tons, and of Type B by y tons. These variables (whose values we must determine) are called the decision variables. Identifying the decision variables, what type of values they can have, is the first key step in modelling a problem.

Here, x and y are real numbers, which must be positive.

Step 2. The objective function

We would like to maximize our profit. Based on table 1, our profit per day = $15x + 10y$

Thus the objective function is to maximize the value of the function, $z = 15x + 10y$

Step 3. The constraints

Looking at the objective function, it is clear that if we keep increasing the value of x, y, or both, then our profit also keeps increasing. However, we cannot produce arbitrarily large amount of fertilizer, because our raw material supplies are limited. These, therefore, are constraints on our solution.

Since the supply of each raw material is limited, we need one constraint equation for each, as follows:

Urea: $2x + y \leq 1500$

Potash: $x + y \leq 1200$

Rock Phosphate: $x \leq 500$

Step 4. Complete the formulation

The problem is therefore completely specified as follows:

maximize	$z(x, y) = 15x + 10y$
subject to	$2x + y \leq 1500$
	$x + y \leq 1200$
	$x \leq 500$
	$x \geq 0$,
	$y \geq 0$

So we have managed to convert words into algebra; but how do we solve this problem? Later, we shall see how such formulations can be solved.

For now, let us just concentrate on ***two interesting aspects***. First, we call this formulation a linear system. This is because the objective function, as well as each constraint is a ***linear function*** in the variables x and y (that is to say, there is no term containing, say, x^2, y^2, xy, xy^2, etc.) Secondly, the constraints are not equations, but ***inequalities***. What does that mean? It means that it may not be necessary for us to use all the raw material we can obtain, in order to maximize our profit (why is that?).

We define some useful terms before continuing. Note that the way that we write our constraints, the company may produce nothing, i.e. x = 0 and y = 0 (verify that substituting x = 0 and y = 0 in the constraint inequalities does not violate any of them). The profit in this case = 0. Likewise, x = 1, y = 1 also does not violate any constraint, and gives a profit, z = 25. In fact, we can find many pairs of values of (x, y) that don't violate any constraint, each with a different value of profit. Each such pair of values of the decision variables is called a ***feasible solution***. On the other hand, x = 600, y = 500 is not a feasible solution, since it violates the constraint on the consumption of urea [2x 600 + 500 = 1700, which is larger than 1500]. Thus x = 600, y = 500 is ***infeasible***. Among all the feasible solutions, there is one (or many) that will correspond to the maximum value of the profit. We don't know the best combination yet, but if we denote one such "best" solution as (x*, y*), then this value will be called an ***optimum value***.

Later, we shall see how to get the optimum value for this and other similar problem formulations. For now, let's look at some more examples.

Example (The Blending Problem):

Here is an extended form of the product mix problem. It is also very common in many industries, e.g. food processing companies making canned soups, animal food, chemical plants, paint manufacture, pharmaceutical companies making drugs, perfume manufacturing and so on. Our example is from the petroleum industry.

In a petroleum refinery, crude oil is separated into many different oils (based on how volatile they are). Several of these oils can be mixed in different proportions to make petrol that is sold from petrol stations. Each oil has a property called octane rating (a number like 60, 80, 90, 99, etc.) Higher octane rating is more desirable for cars – if you visit the petrol station, you will see 2 or 3 grades of petrol sold at different prices; the main difference is in the octane rating.

An interesting property of petrol is that the octane rating (OcR) of a mixture of oils is just the proportional rating of the constituents of the mixture. In other words, if we mix 2 L of Oil 1 (OcR = 60) with 3 L of Oil 2 (OcR = 80), then the OcR of the resulting mixture is given by (2x60 + 3x80)/(2+3) = 72.

The refinery has four types of oils that can be used to produce petrol; they sell three different types of petrol at gas stations, whose prices are based on the minimum OcR of each type. The data for the manufacture cost of each type of oil and the amounts available per day are given in the tables below.

Raw oil	OcR	Available amount (barrels/day)	Price/barrel	Sale price
1	68	4000	31.02	36.85
2	86	5050	33.15	36.85
3	91	7100	36.35	38.95
4	99	4300	38.75	38.95

Petrol Type	Min OcR	Selling Price	Demand (barrels/day)
1 (Premium)	95	45.15	≤ 10,000
2 (Super)	90	42.95	No limit
3 (Regular)	85	40.99	≥ 15,000

Note that the oils can be sold in raw from, or they can be blended to make petrol, and the petrol can then be sold. The objective is to find the compositions of the three types of petrol to maximize the profits.

Step 1. The decision variables:

At first glance, it seems we need to decide how many barrels of each type of petrol we need to produce, so we would need four variables, say, w, x, y, z. However, notice that the petrol is only rated in terms of its minimum OcR, in other words, it could be possible to blend some extra high grade oil into a low grade petrol (e.g. making regular petrol with OcR > 85), yielding higher profits than selling the oil directly. Thus, what we really need is to figure out how much of each type of oil is blended into each type of petrol. Thus we need 4x3 = 12 variables, for which it is easier to adopt a subscripted notation:

x_{ij} = barrels/day of oil of type i used in making petrol of type j (i = 1 to 4, j = 1 to 3)

Thus the total amount of Premium petrol made per day = $x_{11} + x_{21} + x_{31} + x_{41}$.

It's OcR = $(68x_{11} + 86x_{21} + 91x_{31} + 99x_{41})/ (x_{11} + x_{21} + x_{31} + x_{41})$, which must exceed 95.

Thus:

$(68x_{11} + 86x_{21} + 91x_{31} + 99x_{41})/ (x_{11} + x_{21} + x_{31} + x_{41}) \geq 95$, in other words,

$68x_{11} + 86x_{21} + 91x_{31} + 99x_{41} - 95(x_{11} + x_{21} + x_{31} + x_{41}) \geq 0$.

Note that: (a) this is a linear constraint; (b) ($x_{11} + x_{21} + x_{31} + x_{41}$) is always positive (or zero), since each x_{ij} is positive or zero; thus, if ($x_{11} + x_{21} + x_{31} + x_{41}$) > 0, we can multiply either side of the inequality by it without changing the sign of the inequality; on the other hand, if we make 0 barrels of Premium petrol, then the linear form of the constraint is still true.

Step 2. The objective function

We want to maximize our profits. From the data, it is clear that any leftover oil should be sold in raw form, since that will provide some profit. Therefore, our total cost is fixed, and equals 4000x31.02 + 5050x331.5 + 7100x36.35 + 4300x38.75. Thus, to maximize profits, it is sufficient to maximise the revenue from sales. There are two types of revenues; petrol and raw oil. The revenue for petrol e.g. Premium type, is given by: 45.15($x_{11} + x_{21} + x_{31} + x_{41}$); similarly, the revenue from sale of, e.g. raw oil type 1 is given by: 36.86 (4000 - ($x_{11} + x_{12} + x_{13}$)). Following this, the objective function can be written as:

Maximize:

$$45.15(x_{11} + x_{21} + x_{31} + x_{41}) + 42.95(x_{12} + x_{22} + x_{32} + x_{42}) + 40.99(x_{13} + x_{23} + x_{33} + x_{43}) +$$
$$36.85\ (4000 - (x_{11} + x_{12} + x_{13})) + 36.85\ (5050 - (x_{21} + x_{22} + x_{23})) +$$
$$38.95\ (7100 - (x_{31} + x_{32} + x_{33})) + 38.95\ (4300 - (x_{41} + x_{42} + x_{43}))$$

The first line gives the money we make by selling each type of petrol, the second and third lines give the money we make from selling the raw oils of the four types.

Step 3. The constraints

There are several types of constraints:

(a) The OcR constraints:

$$68x_{11} + 86x_{21} + 91x_{31} + 99x_{41} - 95(x_{11} + x_{21} + x_{31} + x_{41}) \geq 0$$
$$68x_{12} + 86x_{22} + 91x_{32} + 99x_{42} - 90(x_{12} + x_{22} + x_{32} + x_{42}) \geq 0$$
$$68x_{13} + 86x_{23} + 91x_{33} + 99x_{43} - 85(x_{13} + x_{23} + x_{33} + x_{43}) \geq 0$$

(b) Can’t use more oil than we have:

$$x_{11} + x_{12} + x_{13} \leq 4000$$
$$x_{21} + x_{22} + x_{23} \leq 5050$$
$$x_{31} + x_{32} + x_{33} \leq 7100$$
$$x_{41} + x_{42} + x_{43} \leq 4300$$

(c) The demand constraints:

$x_{11} + x_{21} + x_{31} + x_{41} \leq 10{,}000$

$x_{13} + x_{23} + x_{33} + x_{43} \geq 15{,}000$

(d) Allowed values of variables

$x_{ij} \geq 0$ for I = 1, 2, 3, 4, and j = 1, 2, 3.

Steps 2 and 3 above show the complete formulation of the blending problem. As you may guess, even if we had some hope of solving our product mix problem by hand, our blending problem is already too large to try and solve using paper and pencil. Yet the two problem formulations look very similar in nature (they both have a linear objective function, and a set of linear constraints); the difference is that some constraints in the blending problem are '≥', while others are '≤'; however, this is not a problem, for we can always multiple both sides of a '≥' constraint by -1, to turn it into a '≤' constraint [why?].

You can see that a common method to solve problems of this form would be nice. Let us look at one last example of a similar nature, before we begin to explore how to solve these problems.

Example: Transportation problem

Steel plant uses coal and iron ore to produce steel. There are three steel plants in a country, which receive iron ore from two different mines. The cost of transporting the ore between any mine and plant is as shown in the table below. The table also shows the maximum production at each mine, and the requirement of each plant.

		transportation cost per ton		
	mine capacity/day	plant 1	plant 2	plant 3
Mine 1	800	11	8	2
Mine 2	300	7	5	4
daily ore requirement at each plant:				
		400	500	200

The problem is to find how much ore must be shipped from each mine to each plant daily, so that the costs are minimized.

Step 1. The decision variables

We need to figure out how much ore is sent from mine i to plant j, so let us denote:

x_{ij} = amount of ore (tons) shipped from mine i to plant j per day.

Step 2: The objective function

Again, this is fairly straightforward. We need to minimize the transportation costs:

Minimize: $11x_{11} + 8x_{12} + 2x_{13} + 7x_{21} + 5x_{22} + 4x_{23}$

Step 3: The constraints

There are three types of constraints: (a) ore shipment from each mine should not exceed its daily production capacity, (b) demand of each plant must be met, and (c) no decision variables can be negative. These are easily written as:

$$
\begin{array}{llllllll}
x_{11} + x_{12} + x_{13} & & & \leq 800 & \text{[capacity of mine 1]} \\
 & x_{21} + x_{22} + x_{23} & & \leq 300 & \text{[capacity of mine 2]} \\
x_{11} + & x_{21} & & \geq 400 & \text{[demand at plant 1]} \\
\quad x_{12} & \quad + x_{22} & & \geq 500 & \text{[demand at plant 2]} \\
\qquad x_{13} & \qquad + x_{23} & & \geq 200 & \text{[demand at plant 3]}
\end{array}
$$

$x_{ij} \geq 0$ for all i= 1, 2, and j = 1, 2, 3.

Notes:

1. In this case, we can easily see that the total mine output = 800 + 300 = 1100, which is exactly equal to the total demand = 400+500+200 = 1100. Thus all the ore that is produced must be shipped, and so we can just replace the inequalities in the capacity and demand constraints by ‘=’ signs.

2. The problem is a simple one to formulate; it has some historical importance in the field of linear programming, though. The first time these problems were solved in a systematic method (these methods were initial forms of the Simplex method to be later discovered by Dantzig) by Kantorovitch in USSR in the 1930’s, and later by Koopmans in 1940’s: contributions for which they were awarded the Noble prize in Economics.

Graphical illustration of the LP formulations

Let us first refresh some high-school geometry:

Point in a 1D space: x = c, where c is a real number on the line representing the 1D space.

Line in 2D: $ax + by = c$

The equation above represents a line using algebra; the figure below represents a line geometrically. The figure also shows what happens to the line when we change the value of c.

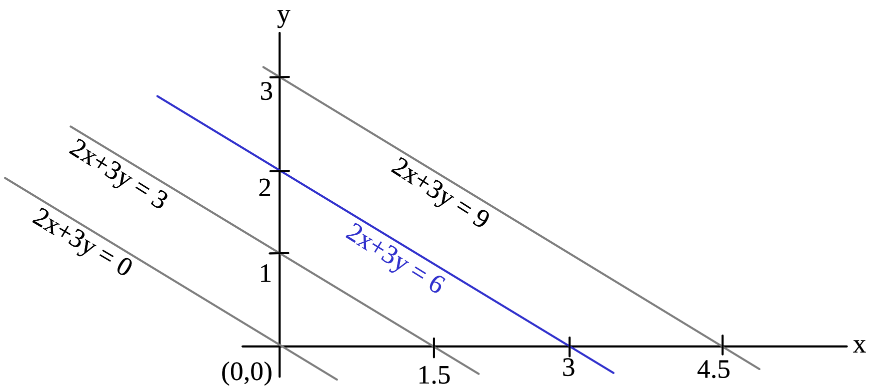

Plane in 3D: $ax + by + cz = d$

Once again, the algebra equation can be represented by geometric figures; the images show how the plane moves as the value of the constant, d, changes (keeping the values of a, b, and c constant). Note also that while the plane is embedded in 3D space, we have managed to represent it on a 2D paper (or computer screen) by drawing its "projected view"; since our childhood, humans adapt to the task of imagining 3D objects by just looking at their 2D projected views.) Also, for different values of a, b, and c, we get planes with different orientation.

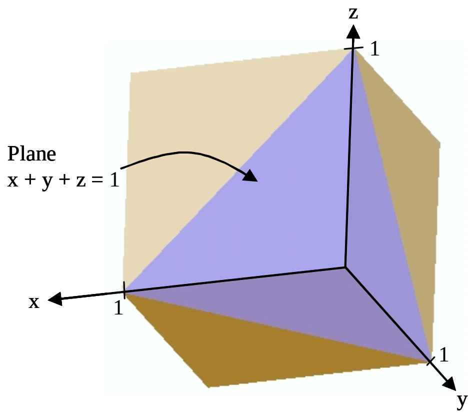

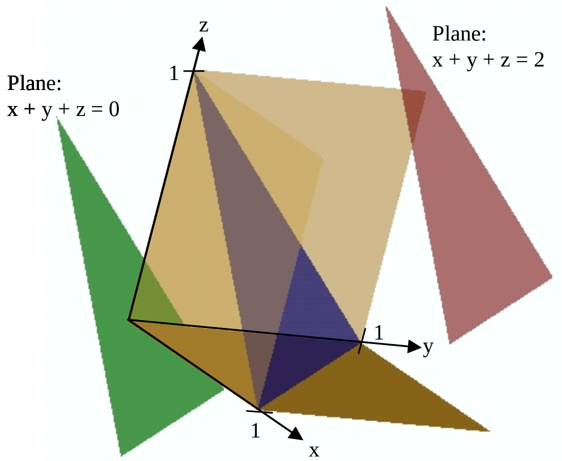

Hyper-plane in nD: $a_1x_1 + a_2x_2 + \ldots + a_nx_n = c$

Now it becomes difficult to draw even projected images of such "hyper-planes" – planes in 4D or higher dimension. We will not try to imagine such figures, and rather rely upon algebra to handle them.

[Note: you may have learnt in high school that it was the brilliant French mathematician Descartes who first made the connection between algebra and geometry, by immersing the geometry into what we now call, to honour him, the Cartesian coordinate space. In the study of LP, we shall see the beauty of this correspondence – the low dimension examples can be solved geometrically to understand the "meaning" of things in the algebra model; this helps us to formulate the method; then, for problems with large number of variables, i.e. problems in high dimensions, we use the same method in its algebraic form].

Half spaces:

Notice that a line, $ax + by = c$, divides the plane into two halves; in one half are all points with coordinates (x_u, y_u) such that $ax_u + by_u > c$; similarly, for each point (x_d, y_d) in the lower half, the coordinates satisfy $ax_d + by_d < c$. Likewise, every plane divides the 3D space into two halves, with the coordinates of points in each half satisfying the same type of relationship. The 2D case is shown in the figure below.

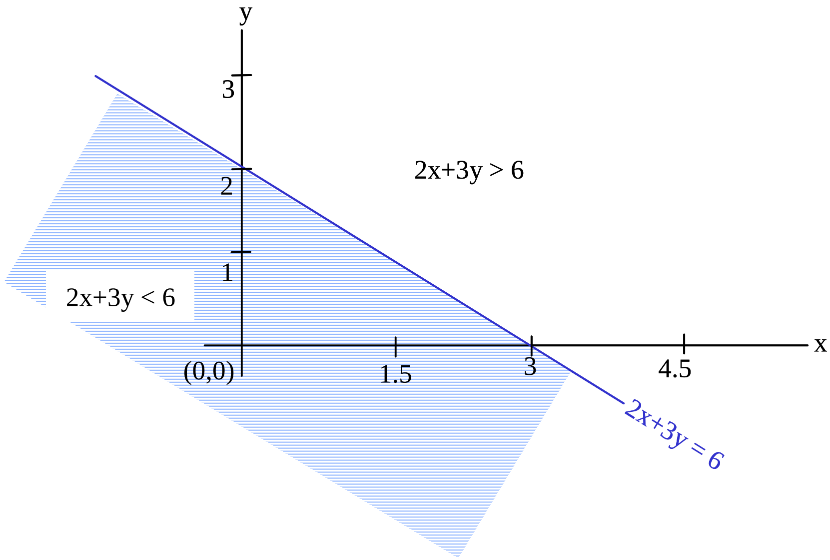

Let us therefore return to the 2D case, and re-visit our product mix problem. The formulation is repeated below.

maximize $z(x, y) = 15x + 10y$

subject to $2x + y \leq 1500$

$x + y \leq 1200$

$x \leq 500$

$x \geq 0,$

$y \geq 0$

What is a feasible solution to this problem?

Any point on the XY plane whose coordinates satisfy each of the 5 constraints.

Notice that since each constraint is a linear equation, it can be represented by a half-plane. Thus if we want to satisfy two constraints simultaneously, then all we need to do is to find the set of points in the XY plane that are in the half plane of each of the constraints. If we look at the half-plane as a set of points, then we look at it as the intersection of the two sets. Extending this idea, satisfying n-constraints requires us to take the intersection of n half-planes. In the figure below, we draw the set of points related to each constraint. The set of points satisfying all the constraints is marked by hatched lines.

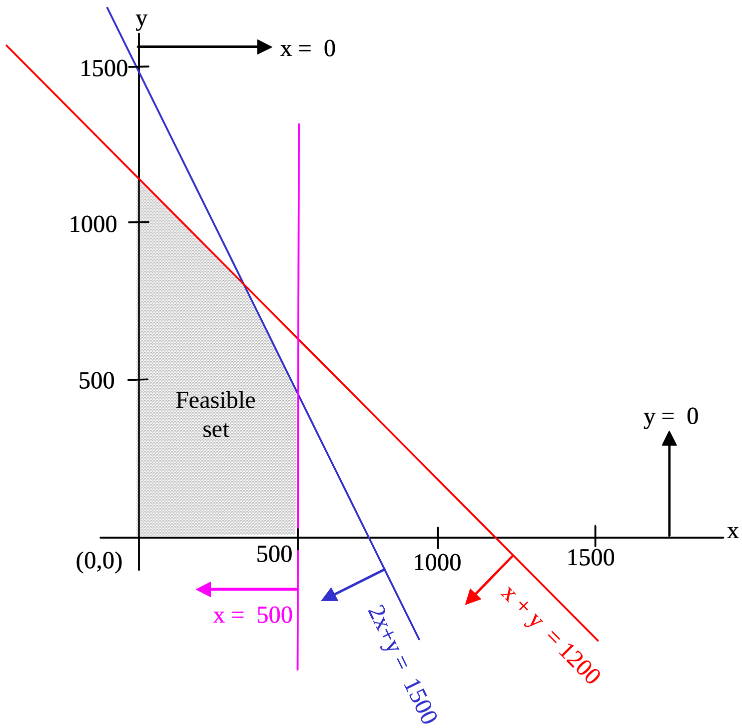

Thus any point inside the hatched region is a feasible solution. There are infinite such points. Which one will yield the best value for our objective function?

Let us arbitrarily try the point (0, 0). At this point, the value of the objective function z = 15x + 10y = 15x0 + 10x0 = 0; graphically, this means that the line 15x + 10y = 0 will pass through the point (0, 0), as shown in the figure below.

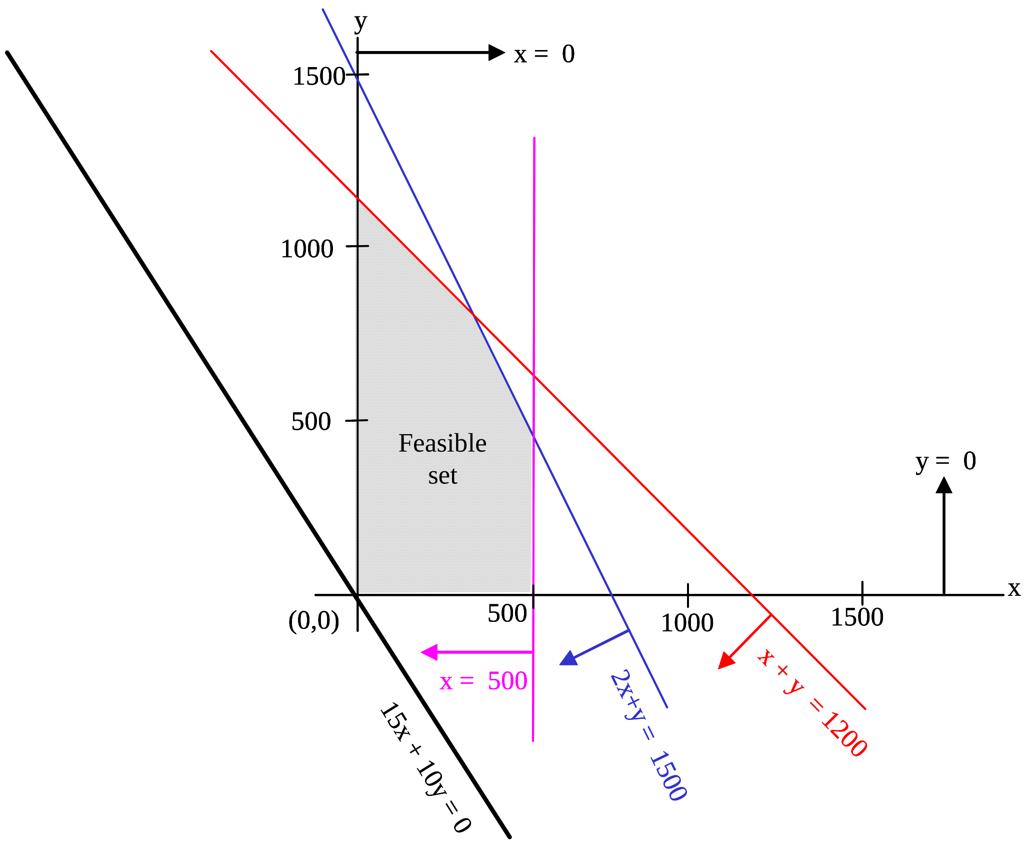

Now note that our objective function is of the form ax + by = c; we learnt earlier that if we keep sliding this line to the right-hand side, we progressively increase the value of c. Since our objective is to maximize 15x + 10y, let us try the line 15x + 10y = 5000, as shown in the figure below.

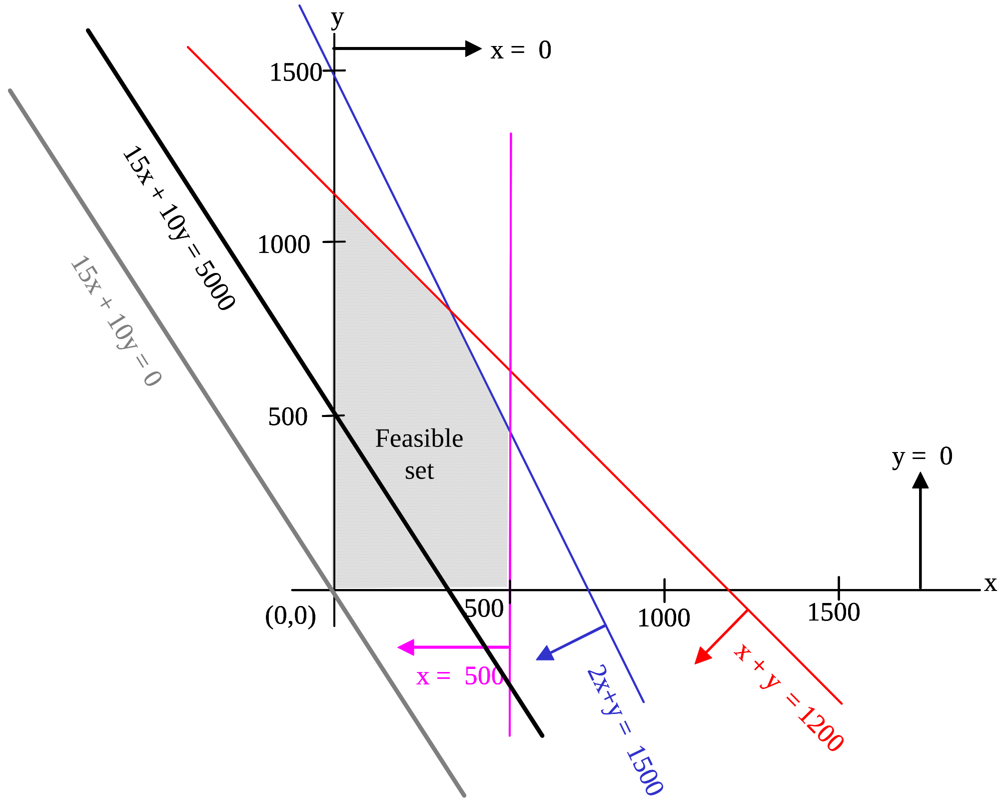

We can see that there are still some grey hatched points to the right-side of this new line. In other words, if we pick one of the points to the right of the new line, but still in the grey region, it must lie on some line with the equation 15x + 10y = d, where d > 5000. In other words, we can find a better solution still!

So we keep sliding the line to the right (i.e. increasing the value of the objective function), until it only just touches some point in the feasible set. Moving it any further right will cause it to go entirely through the infeasible region. The value of c corresponding to this line (shown in the figure below) is the maximum value we can get for our objective function. The point where this line touches our feasible set is called the optimum point. You can verify that this is the point (300, 900) in our case. It is at the intersection of two lines that denote two of our constraints, namely: $2x + y \leq 1500$, and $x + y \leq 1200$.

The conclusion is that our fertilizer company should produce 300 tons of Type A, and 900 tons of Type B fertilizer.

Now let us verify the constraint equations for this mix:

(i) $2x300 + 900 = 1500 \leq 1500$ [OK]

(ii) $300 + 900 = 1200 \leq 1200$ [OK]

(iii) $300 \leq 500$ [OK]

(iv, v) $300 \geq 0$, and $900 \geq 0$ [OK].

Notice in particular constraint (iii), which limits the maximum amount of Rock Phosphate we can use per day to 500 units. In the optimum mix, we only require 300 units. We say that there is 200 units of ***slack*** in this constraint. We shall meet this concept of slack again, later on, since it plays an important role in the algebraic method of solving the linear program.

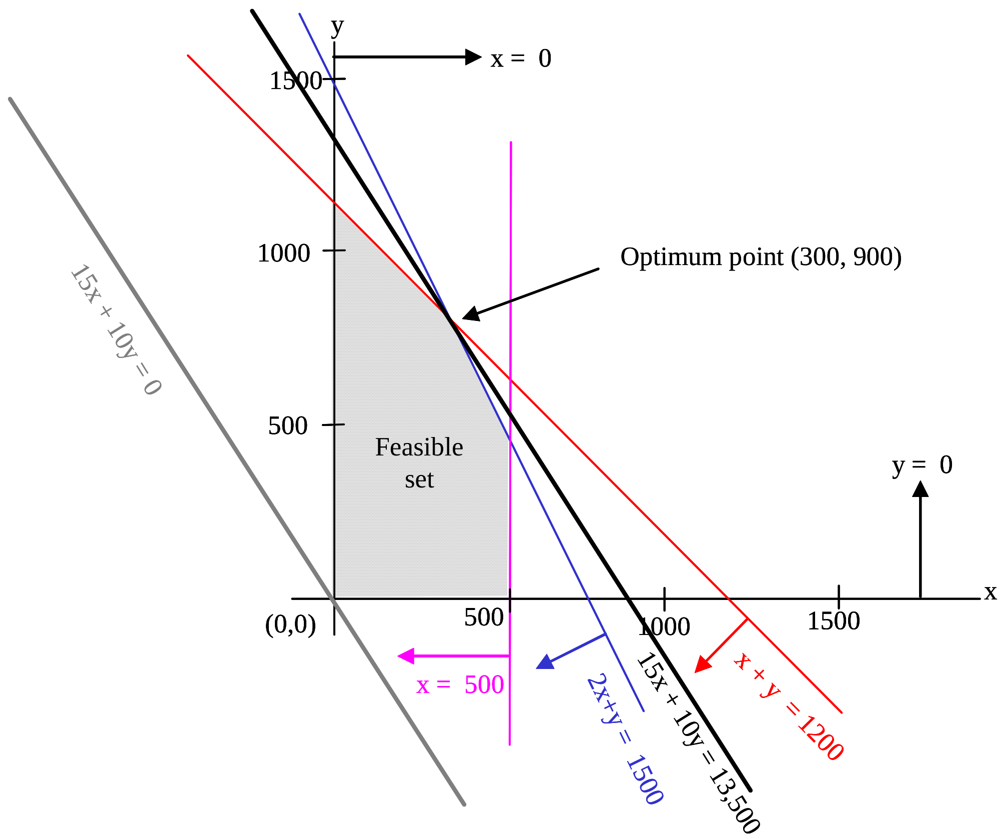

Main reference:

Operations Research, K. G. Murty, Prentice Hall.