Addendum

Since this booklet was written, huge improvements have been made in the area of software speed.

All versions of **Three-D** and **Four-D** are now faster and new versions are available which take full advantage of the power of AT computers and extended memory.

The actual run time of our software still depends on your hardware and the complexity and size of your models, however it is now 'normal' for **Three-D** runs to take minutes instead of hours.

THE FACTS AND FALLACIES OF OPEN PIT OPTIMIZATION

by

J. WHITTLE

1-JAN-89

INTRODUCTION

Optimization makes pit design easier, and makes pits more profitable.

The design is easier because, after optimization, you know exactly what you are aiming for, and you know that small deviations from the target will have almost no effect on the value of the pit.

The pits are more profitable because profit is what is maximized, and there can be no other design which offers higher profitability. The difference in value between a design based on optimization and a traditional design can often be several per cent - millions of dollars.

This booklet shows how pit optimization can produce optimal results whilst taking into account realistic pit design constraints and using modest computer power.

THE FACTS OF OPEN PIT OPTIMIZATION

What is Pit Optimization?

The first thing to recognise is that any pit outline which conforms to safe slope limits has a dollar value.

Provided that you know the nature of the ore body and its surrounding waste, and provided that you know the costs and revenues associated with mining and processing, you can calculate all the costs and revenues for a particular outline and thus get a total value.

Of course, it isn't quite as simple as that. Because the dollars you spend or receive now are worth more to you than the dollars you spend or receive in a year's time, the total value can depend on the sequence in which you mine the pit. However, once such a sequence is established, you can calculate the value of every block within the pit in today's dollars. The value of the pit as a whole is the total value of the blocks which lie within it.

Pit optimization aims to find the pit outline with the highest total value.

Now, you can't find the pit outline with the highest total value until you know the block values; you don't know the block values until you have worked out a mining sequence; and you can't work out a mining sequence until you have a pit outline!

We discuss this apparently insoluble problem during the course of this booklet.

A Simplified Example

Consider the idealized situation shown in the east-west section in Figure 1, where there is a flat topography and a rectangular ore body of constant grade. We will assume that the ore body has a long north-south strike, so that the effects of the ends of the pit are small and we can design on a section basis.

The horizontal lines represent bench levels, and there are eight possible pit outlines.

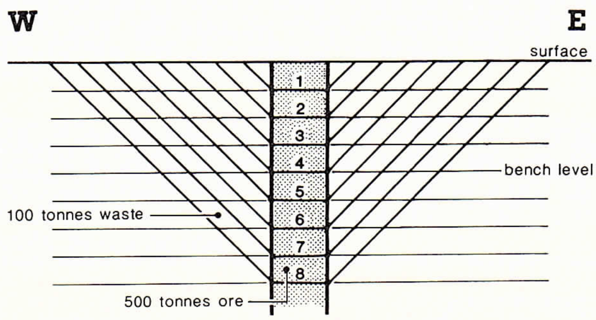

FIGURE 1

If each bench contains 500 tonnes of ore and each diamond-shaped block of waste represents 100 tonnes, the totals in Table 1 apply.

Pit	1	2	3	4	5	6	7	8
Ore	500	1,000	1,500	2,000	2,500	3,000	3,500	4,000
Waste	100	400	900	1,600	2,500	3,600	4,900	6,400
Total	600	1,400	2,400	3,600	5,000	6,600	8,400	10,400

Table 1: Tonnages for the possible pit outlines.

If ore is worth $2.00 per tonne after all ore mining and processing costs have been paid, and if waste costs $1.00 per tonne to remove, then we have the total values shown in Table 2 - in today's dollars.

If we plot pit value against pit tonnage, we get the graph shown in Figure 2.

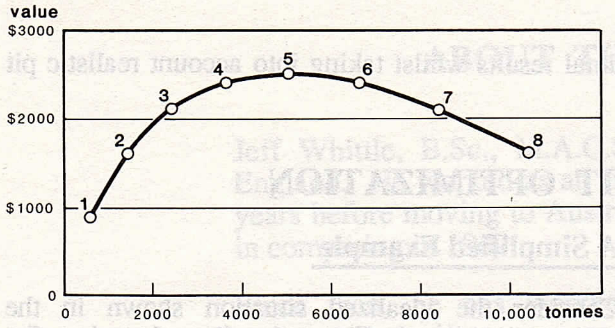

FIGURE 2

Clearly, we make the most money by mining down to Pit 5 - or do we? The fact is that this depends on the mining sequence and on the mining and processing capacities.

If the mining and processing capacities are such that we can mine out the whole pit within a few months, then the sequence in which we mine has little effect on the total cash flow, and Pit 5 is still the optimal pit.

On the other hand, if the mining and/or processing capacity are limited and the mine is going to take five years to complete, then the mining sequence can not only affect the total cash flow; it can affect the size of the optimal pit.

If we mine the pit in strict bench order, (i.e. we complete the mining of each bench before starting the next), then in the early years we will be mining waste at the edge of the pit and paying the costs in today's dollars. However, the ore that this eventually releases will not be reached until four or five years later, and because of discounting, its value will be much less than is required to pay for the earlier waste removal. In short, we would have been better off not mining the edge waste in the first place, so the optimal pit is no longer correct. On the other hand, if we first mine Pit 1 and then incrementally mine out Pits 2, 3, 4 etc., the effect is much smaller.

To illustrate this, we have worked out two idealized mining sequences with each of two different tonnage limits for each of the eight pits. We have used 'flat' mining, where each bench is completed before the next is started; and we have used 'incremental' mining, where we mine out Pit 1 then push back to Pit 2 etc.

For each mining sequence, we have considered the case where processing is limited to 500 tonnes per year, but total mining capacity is unlimited; and the case where processing capacity is unlimited, but the mining capacity is set at 1,000 tonnes per year.

We have worked out the cash flows and the discounted cash flows using a discount rate of 10% per year. Note that 10% is not unreasonably low because in this case the 'Real' rather than the 'Notional' discount rate should be used. The results are shown in Table 3.

Pit	1	2	3	4	5	6	7	8
Value	900	1,600	2,100	2,400	2,500	2,400	2,100	1,600

Table 2: Values of the pits if ore is worth $2.00 per tonne after all ore mining and processing costs have been paid, and if waste costs $1.00 per tonne to remove.

Pit	1	2	3	4	5	6	7	8
Flat Mill 500	900	1,530	1,917	2,085	2,057	1,851	1,486	978
Inc. Mill 500	900	1,530	1,935	2,154	2,220	2,161	2,002	1,763
Flat Mine 1,000	900	1,520	1,868	2,011	1,956	1,552	1,049	449
Incr. Mine 1,000	900	1,520	1,898	2,068	2,105	1,993	1,790	1,540

Table 3: Discounted cash flows for the pits using 'flat' or 'incremental' mining, and with milling or mining capacity limited.

The value of the optimal pit is underlined in each case, and two things are clear:

1. Unless incremental mining costs are very much higher than flat mining costs, incremental mining will make more money for us than flat mining. From the cash flow point of view, incremental mining is optimal.

2. If we mine a pit in a sub-optimal sequence (i.e. by flat mining) the optimal pit gets smaller.

Note that any practical mining sequence for a particular pit outline will produce a value which is between the two extremes. Unless you mine waste and leave exposed ore untouched, it is not possible to produce a worse value than that produced by flat mining, and it is not possible to produce a higher value than that produced by incremental mining.

An Actual Example

The first case where the Whittle Programming Lerchs-Grossmann Package - **Three-D** - was used was on a design, in 1985, for a gold mine south of Kalgoorlie in Western Australia.

A team of experienced engineers had spent some weeks producing and evaluating a range of alternative designs before finally completing a detailed design which they believed to be the best which could be done. They had used computer-assisted techniques, but not optimization.

In three days, using optimization, a new fully detailed design was produced. When compared with the original, it proved to have a value which was five per cent higher!

This five per cent extra profit - representing millions of dollars - could be obtained not by working harder, not by outlaying more money, not by taking more risk .... but by mining smarter!

Although the company involved has many open pits and is regularly designing and re-designing them, they have not done any more manual designs. This is because optimization is better, and also because it is easier.

Detailed Design

Pit optimization works from a block model and gives a pit outline in the form of a list of blocks to be mined. Clearly this jagged outline is not a practical pit design. The outline must be smoothed, and haul roads and safety berms must be put in.

Smoothing the outline is easy.

If you are drawing the final pit by hand, then you merely draw a smooth curve approximately through the mid-points of the block edges.

If you are using computer-aided design, you contour the surface defined by the centre of the base of the lowest block in each column.

Now you might object that, having optimized to precise block outlines, you are now negating the optimization. Not so!

In the first place, tonnage is added and subtracted alternately along the outline in almost identical quantities, as is clearly illustrated in Figure 3.

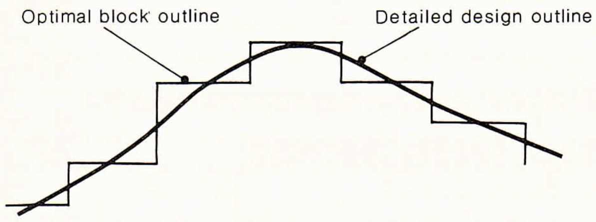

FIGURE 3

If a particular part of the outline is entirely in ore, or entirely in waste, then these balanced additions and subtractions have no effect whatever on the practical value of the pit. Even at the boundary between ore and waste, what is lost in one part of the pit will very likely be gained in another.

Provided that the optimal block outline has been generated with average required slopes, then similar balanced additions and subtractions occur when haul roads and safety berms are incorporated into the pit walls, as can be seen in Figure 4.

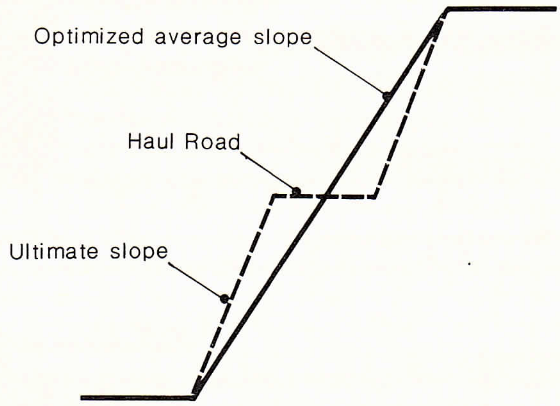

FIGURE 4

Points where the haul road is more than half way down and thus causes extra waste to be mined, are balanced by points where it is less than half way down and thus causes less waste to be mined.

In the second place, optimal pit outlines are very robust in the face of dimensional change.

Refer again to the curve linking pit value and tonnage, repeated in Figure 5. Note that this curve is smooth in all real cases.

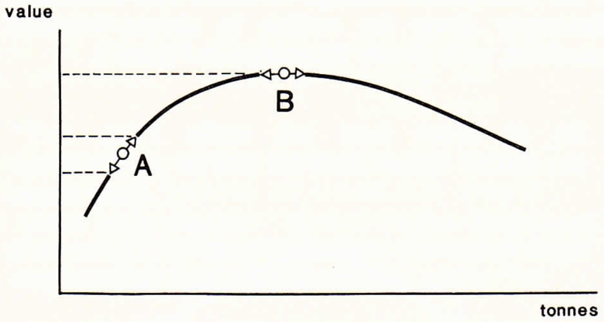

FIGURE 5

If detailed design starts from a point (A) which is not optimal, then small changes to the design can make significant changes to the value of the pit. This is the normal case, and pit designers over the years have spent days, or even weeks, assessing the effects of detailed changes to designs.

On the other hand, if detailed design starts from the optimal outline, point (B), where the curve is horizontal, small changes have almost no effect on the value of the pit. Frankly the precise detail of where you put this haul road, or that safety berm, has no significant effect on value, and you might as well put them where it is convenient.

Provided you follow the 'spirit' of the optimal outline, without tidying it up too much, the resulting design will be so close to optimal that any difference will be irrelevant.

This makes design very easy.

Block Size

Optimization works from a block value model and true three-dimensional optimization finds the optimal list of blocks to mine. That is, not a single block in the list can be changed without reducing the total value or breaking the slope constraints. There can only be one optimal list.

What size of block should you use?

We can identify four different block sizes which are relevant in pit design.

1. There is the block size used for ore body modelling. It has to be small enough to outline the ore body clearly. This can lead to block models consisting of millions of blocks. Some systems reduce this number by allowing sub-division of blocks as required. Some systems use outlines directly rather than blocks.

2. There is a block size which represents the minimum volume which can be selectively mined. This is the size at which block values should be worked out.

3. There is a block size which is suitable for pit optimization for design.

   We now have considerable experience regarding the effect of block size on optimization for design and we find that if

   (a) the final pit occupies a substantial part of the full (rectangular) model,

   and

   (b) the final pit outline (not the ore body) is reasonably regular (i.e. not very convoluted),

   then the optimization of a full model of 100,000 blocks is usually sufficient for design work.

   A model of this size can be optimized overnight on a PC AT.

The reason you can manage with so few blocks lies in the shape of the curve in Figure 2, which is reasonably typical of real cases. This indicates that, providing you get reasonably close to the optimal pit volume, you will get very close to the optimal value.

   (When re-blocking for optimization purposes, you must do it by adding together the dollar values of the smaller component blocks to get the value of each larger block. If you merely average the grades of the smaller blocks, the result will be distorted because the grades will be smoothed out.)

4. Finally there is a block size which is suitable for doing sensitivity studies using optimization.

   Where you are merely concerned with plotting graphs of tonnage, value etc. against variables such as product price or slope, then a model of 10,000 to 20,000 blocks is sufficient. Although the outline it gives may be somewhat irregular, the deviations will be random and will largely average out.

   A model of this size can be optimized in under half an hour on a PC AT.

Note that, for sensitivity work, there is no need to do any detailed designs. This is because the difference in tonnage or value between two optimal block outlines will be very close to the difference between the two corresponding detailed designs.

Costs

When calculating block values prior to pit optimization, there are two rules to follow:

1. The block must be given a value on the assumption that it has already been uncovered and will be mined.

   Thus, stripping ratios, and break even cut-offs based on assumed stripping ratios, are irrelevant. Ore body modelling which is based on such considerations is almost certainly too conservative.

   The only cut-off which is relevant is that used in deciding whether ore which is to be mined is of high enough grade to repay the processing costs and any extra handling costs.

2. Apart from obvious incremental costs, any overhead expense which would stop if mining were stopped must be included in the cost calculations.

   This is because, when an optimization program is 'deciding' whether to mine a block it is, in effect, deciding whether to extend the life of the mine, and must therefore take account of the costs involved in that.

   If overheads are not included in the costs, the program will probably design a pit which will not support the overheads in its final years.

   Overheads should be divided between mining and processing. It can distort the pit if all overheads are added entirely to processing costs or entirely to mining costs.

The Steps Involved in Designing a Pit

To design a pit which will make as much profit as possible, you should carry out the following steps:-

1. Develop a model of the ore-body, the waste and the air.

2. Produce a block model with blocks no smaller than the selective mining unit.

3. Carry out the sensitivity work.

   For a range of economic conditions:-

   (a) Generate a value model.

   (b) If necessary, re-block it to reduce the number of blocks to the range of 10,000 to 20,000.

   (c) Optimize to get the pit tonnage, pit value and approximate shape. (These runs are very quick.)

   During this work, sort out the general scale of mining and hence the costing details. Also, decide approximately where the haul roads are to go and adjust the slopes in those places to reflect the average slope required.

4. Having fixed the economic parameters, generate a value model.

5. If necessary, re-block to reduce the number of blocks to about 100,000 (more if the pit outline will be very convoluted).

6. Carry out the optimization using the required slopes.

7. If you think that the mining sequence you propose will significantly affect the size of the optimal pit:-

   (a) Repeat steps 4, 5 and 6 with progressive reductions in product price so as to produce a set of nested pits. Two or three should suffice.

   (b) Evaluate each pit subject to your mining sequence and select the one with the highest value.

8. Do the detailed design using the highest-valued pit as your guide.

   Note that only one detailed design is done, and that is easy because the shape and size of the pit have been established.

The Whittle Programming **Three-D** package will do steps 3(c) and 6 for you.

The Whittle Programming **Four-D** package will do steps 3, 4, 5, 6 and 7(a) for you, and offers flat and incremental mining sequences for 7(b).

THE FALLACIES OF OPEN PIT OPTIMIZATION

1. "You need a mainframe for pit optimization."

   Not any more - hundreds of open pit optimizations have been done on PCs using our software.

2. "You can't allow for haul roads in pit optimization, so that the final design carries too much waste."

   You can if, during optimization, you use average slopes rather than ultimate slopes.

3. "Pit optimization sometimes omits profitable ore."

   Floating cone programs frequently do this. Programs such as ours, which are based on graph theory, never do.

   Floating cone programs, which are easy to program but give erratic results, have given pit optimization a bad name.

4. "Pit optimization may give you a pit worth a million dollars when a pit half the size exists which is worth only a dollar less."

   True, in theory, but very rare in practice, and you can easily spot the problem during the sensitivity runs.

5. "Pit optimization will produce pits which don't pay their way in their final years."

   Only if you calculate the costs incorrectly.

6. "Pit optimization takes no account of minimum mining widths."

   In many cases our packages can take account of minimum mining widths through the use of additional constraints. In most practical cases, however, this is not necessary.

7. "Pit optimization can only cope with slopes which are one block up and one across."

   That is 2D Lerchs-Grossmann, not 3D. In our packages, slopes are completely independent of block proportions and can vary with both direction and position.

8. "Because there is so much uncertainty about the ore body, there is no point in using optimization."

   That is a bit like saying that, because your golf swing isn't perfect, there is no point in aiming at the flag.

9. "Pit optimization programs can only handle simple ore bodies."

   Quite the reverse - complexity of the ore body has very little effect on run times.

10. "You can only optimize initial designs, it is no use for re-designs once mining has started."

    The only thing you need to do which is different in a re-design is to change the blocks which have already been mined to air blocks. These have a value of zero.

11. "Optimization cannot take account of such things as site boundaries and immovable plant."

    It is quite easy to handle such things by putting blocks with very high negative values into the model at the appropriate places. The optimization package can never justify mining these blocks so they are left in place.

CONCLUSION

The aim of open pit design is to create profitable and practical pits, and we have shown how optimization helps you to create more profitable designs in less time and with no compromise whatever as to practicability.

As a company, Whittle Programming produces and markets open pit optimization products exclusively and, although we have not gone into details of our two optimization packages, **Three-D** and **Four-D**, we can say that they are well-established packages and are recognised for their quality and reliability. They will operate on any computer, alongside any generalized mining package.

The packages are constantly being enhanced, and the enhancements are passed on to our users. A great advantage of being a small and specialized company is that we know our software intimately and can offer excellent support.

Prices and details readily available on request from us or from any of our agents.

ABOUT THE AUTHOR

Jeff Whittle, B.Sc., M.A.C.S., was born and educated in England. He worked as an Experimental Physicist for some years before moving to Australia, where he became involved in computing in 1962.

Since then, he has worked on a very wide variety of computing projects as a programmer and/or as a project leader. All his work has been in the area of technical (as distinct from commercial) computing, and he has developed a strong interest in optimization and operations research.

In 1977, he became a Consultant and, since 1979, has worked exclusively on computing for the mining industry. Apart from the general areas of data collection and manipulation, he has developed advanced geostatistical programs, and done extensive work on scheduling involving multiple pits using linear and mixed integer programming techniques.

Jeff Whittle developed **Three-D**, his first open pit optimization package, in 1985, and his second and much more comprehensive package, **Four-D**, was released in October, 1987.

His experience has mainly been in gold mining, but he has also had exposure to iron ore, copper and diamond mining through the use of his optimization packages.

Jeff Whittle's expertise in open pit optimization techniques is unique in the world.