Integrated Strategy Optimisation for Complex Operations

B King[1]

ABSTRACT

While large mining operations frequently provide enormous value for their shareholders, they also contain enormous challenges for those determining the operating strategies that maximise the net present value (NPV). Experience shows that several percentage points of additional value lay waiting to be released though additional or 'second order' optimisation of the extraction policies. For operations constantly looking for ways to trim costs and add value, this analysis may help initiate a step change in the project's NPV.

Determining the best operating policy is often limited by the analysis time and the availability of skilled engineers to appropriately utilise various planning tools. For example, the size, shape and timing of even a single pushback may have thousands of valid alternatives. For a very small operation, these alternatives may be evaluated to determine which one results in the best project value. Mines with durations of more than five years often have so many valid alternatives that the number of neutrons in the universe appears small by comparison!

Optimisation algorithms implemented in commercial software tools for maximising project value provide guidance for some key parts of this process. The question then becomes how we can answer questions that are not explicitly optimised by the algorithms. These questions may include what mining and processing capacity should be installed, the size and timing of a processing expansion, the timing of extracting resource from nearby mines, and when a resource should be mined from underground rather than from the surface.

This paper outlines a framework for optimising many of these policies in large, and often very complex, mineral resources. Examples are presented from experiences at major operations in Australia, Chile, Peru and the USA. It is hoped that this will assist engineers to better exploit the finite resources in our care.

INTRODUCTION

The first part of this paper focuses on the wider context of mine planning and the objectives of planning finite resources. It aims to clarify some of the key parameters of an optimisation study including the business objectives, the need to break the problem into manageable tasks, taxation and how to proportion costs for different studies.

While there are several decisions that can be made with the guidance of commercially available optimisation algorithms, the second part of the paper focuses on how to determine the best choice for a policy that is *not* optimised by these tools. We do not have to search far in large operations to find strategies that are not optimised and yet may release substantial value prizes for appropriately made decisions.

PART I – STRATEGIC PLANNING CONTEXT

Business objectives

While small mining ventures have ridden the tides of metal prices and market conditions for short-term profit and unconstrained resource high grading, the following objectives of a modern mining company are more commonly stated:

- to act responsibly as a steward of the resources in its care so that they benefit both the countries in which they are found and the world at large which depends on them; and
- to create long-term wealth for its own shareholders.

1. Director, Strategy Optimisation Systems Pty Ltd. E-mail: Brett.King@stops.com.au

These objectives are believed to be in harmony with each other, and both are a vital part of the mining industry. In determining realistic policies, environmental, safety and political constraints must all be considered. There is clearly little point investing time and effort in developing plans that cannot be implemented for failing to obey these constraints.

The present value of a business with a finite life can be calculated by discounting a series of cash flows. In order that the *net* present value may be calculated, the cash flows must include *all* costs and revenues of the project. The net present value is often used as the criterion for increasing long term shareholder.

The following formula shows how the net present value (*NPV)* is calculated from a series of *cash flows* using a constant discount rate, δ, over the *life* of the operation:

$$NPV = \sum_{year=1}^{life} Cashflow_{year} \times (1+\delta)^{-year} \quad (1)$$

This formula discounts each cash flow to its value at the beginning of the first year. This equation assumes that the cash flow occurs at the end of each year.

Parameters that impact the net present value

As can be seen from Equation (1), there are only two parameters that can impact the net present value of a deposit: the discount rate and the annual cash flows.

The discount rate is the reward demanded by investors for accepting delayed payment, also referred to as *rate of return, hurdle rate* or the *opportunity cost of capital*. It is an opportunity cost because it is the return foregone by investing in the project rather than investing in other securities of similar risk. The discount rate incorporates a 'safe' ('risk free') component (often compared with USA government debt) and various technical, country, political and other risks.

This discount rate should be confirmed as being appropriate for the type of analysis being undertaken. For example, we would expect the discount rate of a pre-feasibility project to be different to that of the same project after operating for several years due to the experience in mining and processing risks.

The annual cash flows provide the primary controllable means of modifying the NPV of a deposit. The cash flow components may be a function of a large number of properties including:

- price at which the products are sold (eg gold and copper prices);
- quantities of rock processed for ore and sent to waste;
- recovered proportion of processed products;
- operating costs of mining and processing equipment;
- environmental rehabilitation and decommissioning costs;
- equipment purchase, replacement and maintenance costs;
- engineering, consulting and administration costs; and
- royalties and government taxes.

Although this list is not intended to be comprehensive, it is clear that there are many decisions that can impact the cash flow and, hence, the NPV. The number of parameters are often further complicated by monthly or daily changes to these parameters and geological variations.

Although the ultimate objective is to maximise these discounted cash flows, they are manipulated through parameters such as production rate and cut-off grade (Roman, 1973). The parameters may be divided into two categories:

- *Uncontrollable* – A few of the parameters that impact on the annual cash flow are not under the control of the mine. These may include government royalties and the resource grade distributions. If there is no control over these parameters then they are fixed for the study.
- *Controllable* – Normally there are many parameters whose impact on the cash flow and resulting NPV are controllable throughout the mine life. These parameters may include the production rates, excavation sequencing, cut-off grade and, to some degree, product recoveries. Controllable parameters are the focus of optimisation studies.

The definition of what issues are controllable may need to be defined for each study. For example, commodity prices are often considered 'uncontrollable'. In studies centred on a relatively small copper producer, this may be a valid assumption. When a relatively large proportion of market production for a particular commodity is studied, the influence of the operating strategies on prices may be significant and to some extent may be considered 'controllable'.

Optimal mine design

Ideally, all possible decisions that could influence a mine's value should be considered to achieve designs which will result in the maximum NPV. The number of combinations of these parameters over the life of the mine is overwhelming for any global optimisation technique unless several assumptions are made.

Until the entire mine design problems can be solved with one integrated algorithm, smaller components of the process are often worked on sequentially, as indicated below:

1. *Resource estimation* – typically using a variety of geostatistical techniques with assumptions about mine cut-off grades, spatial grade relationships and selective mining unit (SMU) dimensions.
2. *Ultimate pit limits* – often using the Lerchs and Grossmann algorithm for determining the ultimate pit shell requires assumptions of ore and waste classification (and therefore cut-off grades), costs and grades that are independent of the block schedule (Lerchs and Grossmann, 1965; Whittle, 1988, 1989).
3. *Pushback design* – incorporating ramps, geotechnical slope constraints, geometry and access constraints.
4. *Scheduling timing of pushbacks* – Figure 2 - Cut-off grade policies for surface to underground transition, normally with fixed cut-off grade, processing and mining policies.
5. *Cut-off grade and mining rates* – for a fixed sequence of mining as described by Lane (1964, 1988).
6. *Comminution and process policies* – such as grind size, reagent consumption, residence time.
7. *Financial analysis* – based on various cost and revenue time-dependent drivers.

These subproblems normally form a sequential mine planning process that can be repeated iteratively. Higher value results are expected as more of the subproblems are simultaneously considered in a single optimisation. Efficient algorithms are required to ensure the solution times do not explode as the complexity increases. These may be very expensive to look for and, if found, costly to turn into practical tools for the mining industry.

For example, the dynamic programming based algorithms used in the COMET software integrate pushback timing, cut-off grade, processing policies and financial analysis from the categories above (stages four to six above). This allows the interaction between the policies to be exploited to maximise the project value. Further discussion of COMET is provided by Wooller (2004) and King (2004).

Because any 'optimising' tool only works on a limited model of reality, results should be reviewed to check that they are reasonable. *When breaking the problem down into components, the resulting designs lose any guarantee of finding the maximum present value, they simply hope to be close to the maximum.* While this may be somewhat disappointing to management or investors, it is a sobering reminder of the complexity of mine planning and the need to appropriately resource this vital work.

Building model complexity

Recognising that we are always working on a model representation of reality, the quality of a scheduling model is paramount to the quality of the results. The amount of complexity that is warranted should be scoped with respect to the questions to be answered and the time available. Short time frames available for a study can have substantial constraints on the model complexity available. Even some simple models may be good predictors of NPV and suitable for some relative ranking of strategic options.

When a complex model is required, it is preferable to build it up in stages. Start with the main constraints (say annual mining, processing and market capacities) and a simple model of cash flow (say simple costs per unit mass mined and processed, an average recovery and single metal price).

By looking at the main components of the cash flow (ie costs and revenue streams), you may then decide to model part of the operation in more detail. Ultimately, you are trying to produce a good estimate of cash flow from the activities associated with mining an increment of rock. There may be several ways to get to this position, some with very complex models, and some with quite simple models. By reviewing the *accuracy* of cash flow results you may see a reducing benefit by making a model more *precise*.

Once you have a working model, you can do preliminary sensitivity studies to see which are the strongest value drivers and then focus further work in these areas. Sensitivity studies at parameter limits may help eliminate weak value drivers. For example, you may try halving and doubling costs to see what impact each has on the project NPV and operating policies.

Taxation and cost considerations

Government taxation requires a distinction between '*Capital*' and '*Operating*' costs but does not require a breakdown of '*Fixed*' and '*Variable*' costs. 'Capital' costs are normally treated differently for depreciation, a significant component of the taxation system. Tax calculations have been demonstrated to impact operating policies such as cut-off grade (Schaap, 1981; Dowd and Xu, 1995, 1999). Despite this theoretical result, the substantial complexity of calculating after-tax NPVs is beyond the capability of many optimisation algorithms and commercial tools.

The simplification often assumed is that taxation makes all cash flows proportionally smaller in a way that does not impact on operating policies. This means that operating policies are initially decided using NPVs based on cash flows without taxation, depreciation and loan repayments. Although this is a welcome simplification to the mine planning process, plans should be reviewed to see if this is a good approximation for the business. The review process may utilise a more comprehensive financial model (with tax considerations) to analyse selected schedules.

While many inputs are required to calculate the cash flow and subsequent NPV, experience reveals costs are often poorly defined for optimisation studies. Part of the reason for this may lie in the different reporting required by government (primarily for taxation) than is appropriate for optimisation. Optimisation studies often try to use the same accounting systems that generate the accounting information for the government reports. It is worth remembering that these systems are normally designed to track expenditure, not predict future costs and demands.

To help clarify the distinction between the various costs, the following definitions are suggested to better define costs for optimising a project. The terminology has deliberately not used commonly misunderstood words such as fixed, variable, capital and operating.

- *Foundation costs* – There are some costs that are incurred at specific times regardless of how much material is processed and when the operation shuts down. Although these irreversible outflows may be significant and impact the NPV of the project, they do not alter the decisions of how best to operate the business. Although many of these costs are incurred prior to a planning study (eg feasibility study costs and purchased equipment), they may also be in the form of future commitments (eg to clean a contaminated area over the next five years). Exploration expenditure may also fall into this category when not related to improving the confidence of existing reserves.
- *Activity costs* – These costs include all elements that change as the amount of material mined, processed, metal productions (most of the traditional 'variable' costs). These costs are best described per truck hour operated ($/hr), mass mined ($/t), processed ($/t), concentrated ($/t) and metal shipped ($/oz). When maintenance and replacement costs are strongly correlated to activity (such as truck replacement costs), these costs should also be expressed as activity costs.
- *Reserve timing costs* – There are several costs that are incurred as reserves are mined and may include infrastructure construction and decommissioning costs. For example, closure costs are dependent on the time when the operation processes the last ore. Typically these costs will include decommissioning the mine workings, processing plant, heap leach pads and infrastructure. These costs may be incurred before the mine actually shuts down and may continue many years afterwards. These costs are best expressed in $M and may require discounting to bring them to a cost at the time of closure. Dewatering costs may be appropriately modelled in this category.
- *Period costs* – These costs are incurred annually while the project operates. These costs do not vary if more (or less) material is mined in a year, more material is treated in a year, or more metals are produced. These costs stop as the last ore is treated and mining ceases. The units of these costs are best described as $/year and are incurred regardless of the production rates until the mine is closed.

Although many general and administration costs are put in the 'Period' costs, each cost should be reviewed in light of the time frame which is being planned. Staffing levels fixed for most daily plans are considered largely 'Activity' costs for strategic analysis. Equipment replacement 'Capital' is normally dependent on production and so should be included with the mining and processing 'activity costs'.

PART II – OPTIMISING COMPLEX POLICIES WITHOUT OPTIMISATION ALGORITHMS

Large operations with multiple orebodies, mining areas and processing alternatives are often so complex that the feasible solutions make the number of neutrons in the universe seem small (10^{128}). This next part of the paper is about using available tools to solve some of these problems found in many large operations. The problems selected are those that are important to many large operations and which don't have integrated optimisation algorithms to solve them.

Generic process

The question to answer is: '*How can I use the available tools to optimise policies that are not optimised by these tools?*'

Let us assume we have policies A, B and C to determine, though only policies B and C can be simultaneously optimised with available algorithms. This paper offers the following generic steps for approximating the optimal choice of policy A:

1. Choose a tool that:
   a. Uses an objective that accurately reflects your business objectives (NPV is assumed in this paper).
   b. Simultaneously optimises as many of the other key policies (policies B and C) as possible. (Sometimes multiple tools are required to model the process; COMET software has been used for the problems identified in this paper.)
2. Identify the broad options for policy A.
3. For each policy A alternative, optimise the remaining policies (B and C).
4. Choose the highest value options for some more detailed analyses if the schedule values are close (within the accuracy of the estimate).

The option for policy A that gives the highest value is chosen to optimise this policy. This process often contains a substantial manual component, which is both expensive and time-consuming. If this decision is to be evaluated often, then some automation may be justified.

The above process can also be used when several policies must be chosen though not able to be simultaneously optimised with the available technology. Policy A can be a complex policy, a combination of several policies. If we have independent policies X and Y, both with three options (X_1, X_2 and Y_1, Y_2). Policy A can be defined as a set of four policies of X and Y (X_1Y_1, X_1Y_2, X_2Y_1, X_2Y_y). The number of options can rapidly approach the number of neutrons in the universe again so some judgement may be needed to consider only reasonably likely options.

While the above approach may be deduced by common sense, the key issue to recognise is that all combinations of policies A, B and C do *not* have to be searched in order to find the highest value path. Only one path for each policy A alternative needs to be valued. This can have enormous time saving benefits when applied to many complex operations.

The following examples have been used to illustrate how this process may be applied to a number of problems found in large complex mining projects.

Surface to underground interface

Many large resources mined from the surface also have a potential resource that can be extracted from underground. Although surface mining methods may be used to extract much of the resource, the highest value for the project should consider both underground and surface options.

There are many factors that impact the ideal transition from surface to underground operations. Some of these issues are listed below:

- Surface
  - cut-off grades,
  - waste stripping, and
  - stockpile generation and reclaim.
- Underground
  - access to higher grades,
  - dilution,
  - proportion of resource extracted (due sterilisation associated with the mining method),
  - production costs and capacities, and
  - capital requirements.
- Combined
  - tailings capacity, and
  - closure cost implications.

There is currently no algorithm (and therefore software product) for determining the best transition between surface and underground mining that will take into account all of the above issues. Where currently available software tools attempt to answer these questions, only a few of these aspects are considered. It is not the objective of this paper to list the limitations of commercially available software tools, many of which can still be profitably used despite these shortcomings.

The question becomes, *'How can we use the available tools to optimise the transition between surface and underground mining?'*

By applying the generic process suggested above the best transition from surface mining to underground can be evaluated. For the example illustrated in Figure 1 there are three different options to evaluate. The underground mining options each have a different designs, production schedules, capital requirements, life and of course value. The open pit designs have several pushbacks that extract different portions of the resource.

These underground alternatives impact on the opportunity costs for processing surface material since every day spent processing surface ore could alternatively be spent processing underground ore. The surface policies, such as cut-off grade and ultimate pit limits, are dependent on the value of the remaining underground resource. If the underground is considered without reference to the open pit, Option A (large underground) is chosen. If the open pit is evaluated without the underground impact, Option C (large open pit) is selected.

For example, without an underground the open pit cut-off grades will normally drop down close to break-even as the last material is mined. With a highly profitable underground that cannot start until the surface operation is complete, open pit cut-off grades will generally increase to bring forward the value from the underground resource. Although specific policies are very dependent on the particular project, Figure 2 shows the change in cut-off grade for the three cases shown in Figure 1. Each of these schedules was optimised using the COMET software using a successive approximation dynamic programming algorithm.

The changes in policies are very dependent on the constraints, economics and resource mined. Figure 2 shows that the shorter open pit options generally utilise lower cut-off grade strategies that have the result of extracting more value from the earlier open pit phases. As is usually the case with optimised cut-off grade policies, the policies rise as more high-grade material is reached and then generally decline with time. While the cut-off grade policies are interesting, the most important number is the NPV presented in Table 1. The highest value schedule was the second option B with the medium surface and underground designs.

Of interest in Table 1 is the mine life, increasing with pit size. The primary reason for this is the lower grade material that was processed in the larger pit options. The underground costs do not justify the removal of all of this material and so larger underground designs have smaller reserves.

A second point to note in Table 1 is that all three cases yielded positive NPVs, some were just a little more positive than the others! This should serve as a reminder that a high value schedule does not necessarily mean that an even higher value schedule is not possible with a little more effort. Further information from the best case (Option B) is presented in Figure 3.

TABLE 1

Surface to underground transition summary results.

Option	A	B	C
Open pit size	Small	Medium	Large
Underground size	Large	Medium	Small
NPV ($M)	2287	2425	2410
Life (years)	21	22	27

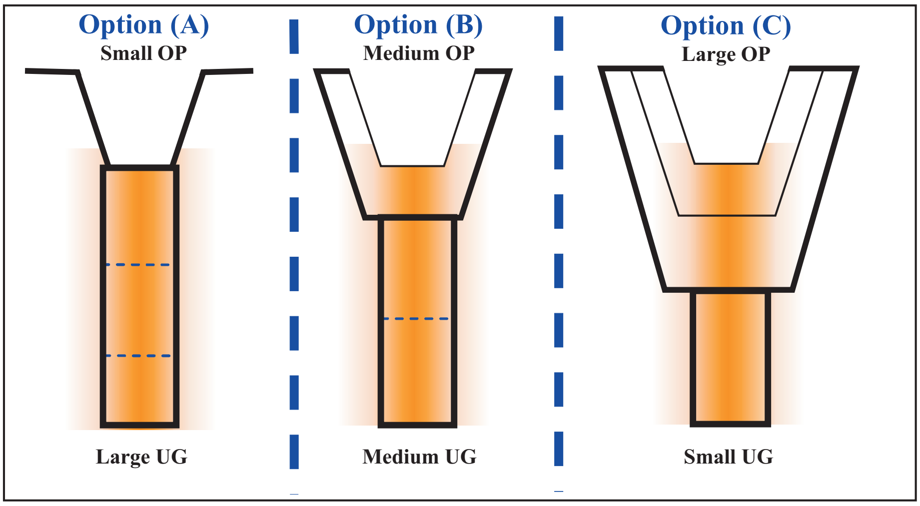

FIG 1 - Surface to underground transition options.

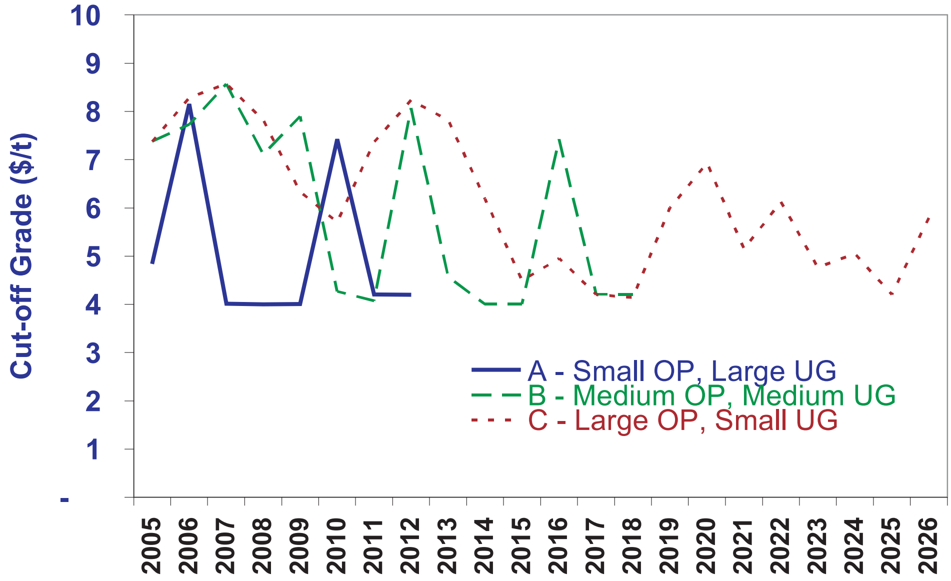

FIG 2 - Cut-off grade policies for surface to underground transition.

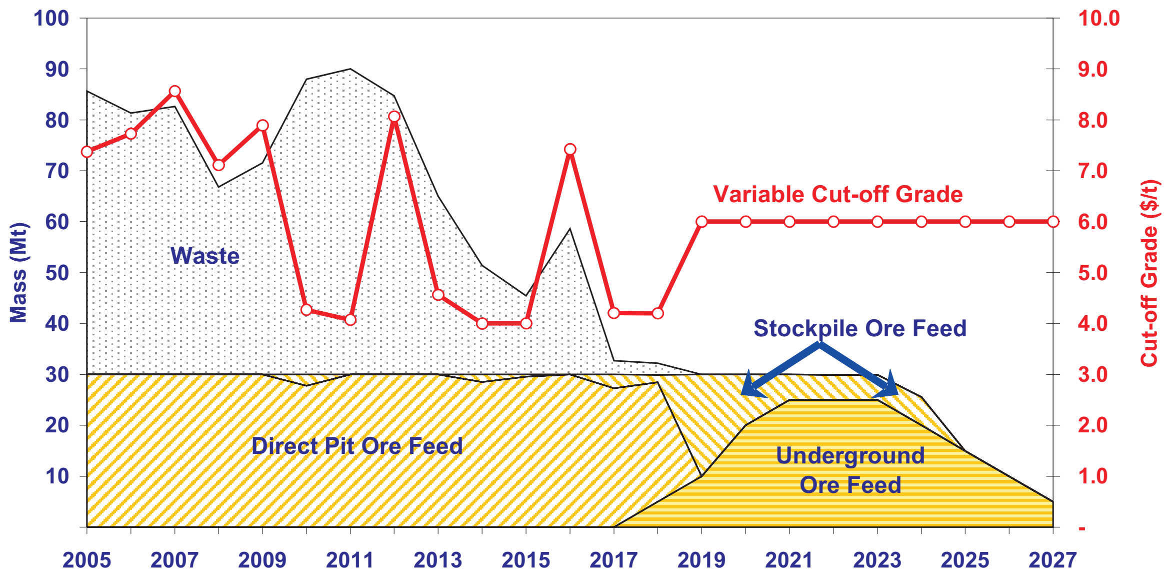

FIG 3 - Highest value integrated surface and underground schedule (Option B, $2425M) with variable cut-off grade and pushback sequencing.

The low-grade stockpile is reclaimed once the open pit ore runs out and during the underground mining period. The underground ore was higher grade than the stockpile but was not able to be mined at a sufficient rate to use the full mill capacity. Ideally, the mill throughput/recovery would be modelled to add further value.

A new set of schedules was therefore undertaken to exploit a time varying grind policy to maximise the project value. For the purpose of this paper, a grind relationship was used in which the mill could process up to ten per cent more material with the loss of five per cent in recovery, or processes 20 per cent less material and realise ten per cent higher recoveries. Figure 4 shows the schedule when optimised with a variable throughput/recovery policy optimised simultaneously with the cut-off grade and pushback sequencing.

A substantial increase in value (+4 per cent from $2425M to $2523M) was realised by simultaneously optimising the grind policy (throughput/recovery relationship) with the other policies (cut-off grade and pushback sequencing). An outline of the theory of how to optimise multiple policies like these was presented by King (2001). The same surface to underground transition was found to produce the highest value and all schedules had higher values.

The addition of grind to the optimisation is an example of adding complexity as a model of the project is developed. As time is spent analysing and understanding the project value drivers, some areas are obvious candidates for greater model accuracy.

It is important to review the sensitivity of these decisions to price, cost and constraint variation. A low reserve schedule that has the highest value at a low price may well be less than the best schedule at a higher price since more reserves can utilise the higher prices). It is also important to recognise the different risk profiles of the resulting schedules. The risk is often a more difficult property to measure; however, there are normally some parameters that reflect this risk.

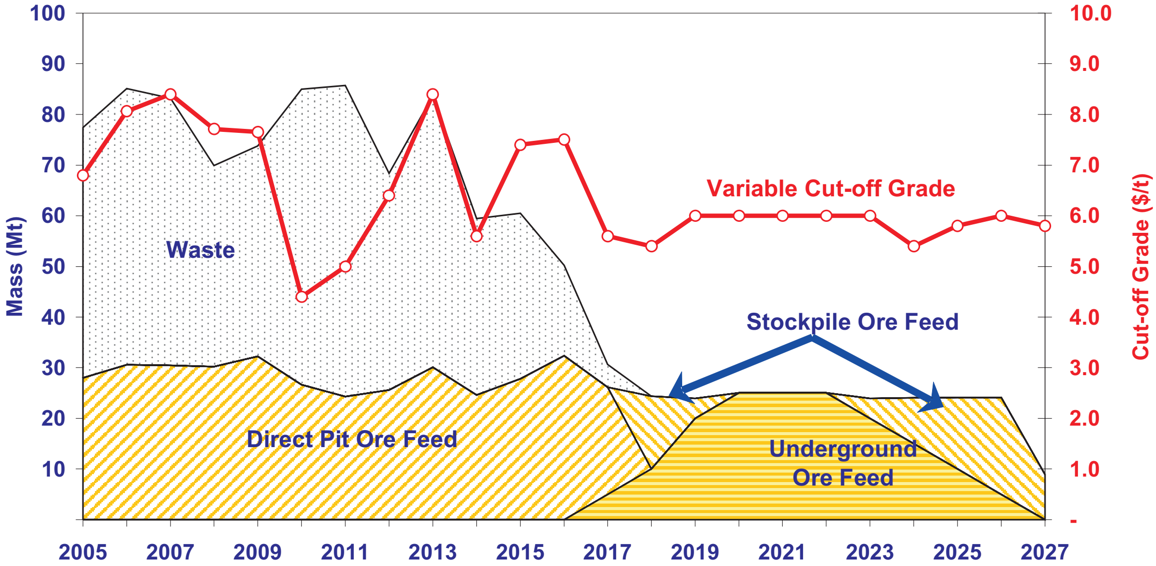

FIG 4 - Best integrated surface and underground schedule (Option B, $2523M) using variable grind (throughput/recovery), cut-off grade and pushback sequencing.

Additional complex policies

Pushback designs

Designing realistic pushbacks is a fundamental part of planning a large surface mining operation. Many engineers use tools based on the ultimate pit algorithms (LG or cone algorithms) to provide guidelines for creating intermediate pushbacks. These tools are often run at lower than expected metal prices to determine a nested set of shells. While this approach does provide useful guides, there are a number of issues that arise in large operations that limit the usefulness of these guides:

- mining and processing capacities are not considered;
- time dependent properties including prices and costs are not considered;
- operating policies such as cut-off grade are not considered;
- interaction between material mined and processed is not considered;
- shells may be much smaller or larger than can be practically mined; and
- ramp locations and some geotechnical constraints (such as stress unloading) are not considered.

The above issues may provide substantial uncertainty of the best shape for intermediate pushbacks. Several options may need to be manually designed and scheduled to find the best designs and maximum value.

The ultimate pit size is also subject to the same assumptions and therefore limitations as described above. For example, the location of the final pushback may need to be confirmed by grouping shells into a realistic width and scheduling with all other policies (such as cut-off grade and stockpiling) optimised.

Mine and process expansion optimisation

Mine and processing expansions may provide the keys to unlock substantial additional project value. These capacities are not automatically optimised by the currently available algorithms and software tools so we ask, '*What is the optimum mining and processing capacity for the project?*' Although a simple question, the answer can involve a complex combination of policies throughout the business.

For example, increasing the flotation capacity would also require increasing the SAG capacity, crushing and grinding capacities, concentrate handling capacities and quite possibly the tailings capacity. Once the entire processing system has been upgraded and new cost and recovery functions implemented you may still see a negligible increase in value. The reason could well be due to the operation being constrained by the mining equipment. When mining constrained, cut-off grades drop to breakeven grades and very marginal material is processed. In order to reveal the full value of a processing expansion it should be coupled with a mining capacity expansion.

Although schedule optimisation tools may not directly provide the optimum choice of mining or processing capacity, by scheduling several options an engineer can rapidly determine the optimum choice of both mine equipment fleets and process capacities.

To evaluate all the possible options of just truck and shovel fleets would be an enormous and unnecessary task. Most of the options are able to be discarded as unlikely to achieve higher value. For example, expanded truck fleets without associated shovel fleets are unlikely to reveal any further value unless the operation was already truck constrained. By applying sensible boundaries to the options and reviewing results as they are generated, options for analysis can be greatly reduced.

CONCLUSIONS

Strategic and long-term plans set the context for shorter term decision making. Substantial value is realised by ensuring that strategic and long-term planning follows the corporate objectives, normally defined using the net present value.

Many optimisation algorithms have been developed to solve parts of the planning problem. There are still important problems that are not able to be automatically optimised with these algorithms. This paper demonstrates that by using an efficient schedule optimisation tool, many of these policies can be optimised to add substantial value to a project.

ACKNOWLEDGEMENTS

I am grateful for the opportunity and support to spend the last decade researching optimal mine scheduling and developing tools useful for engineers around the globe. Many professional engineers at Rio Tinto Technical Services and BHP Billiton have willingly given time and ideas to make sure that what was being developed would be useful to real world mining problems. Richard Wooller has been a tremendous sounding-board and debating partner on many issues. Thanks also to Brian Baird for his support and help getting these ideas implemented into major operations around the world.

REFERENCES

Dowd, P A and Xu, C, 1995. Financial evaluation of mining projects, in *Proceedings APCOM XXV*, pp 247-254 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Dowd, P A and Xu, C, 1999. The financial evaluation of polymetallic mining projects, in *Proceedings APCOM '99 Computer Applications in the Minerals Industries, 28th International Symposium*, (Ed: K Dagdelen), pp 385-392 (Colorado School of Mines).

King, B M, 2001. Optimal mine scheduling policies, Royal School of Mines PhD, London University, United Kingdom.

King, B M, 2004. http://www.stops.com.au/Product_Comet.htm COMET Optimisation Services, Strategy Optimisation Systems Pty Ltd.

Lane, K F, 1964. Choosing the optimum cut-off grade, *Colorado School of Mines Quarterly*, 59(4):811-829.

Lane, K F, 1988. *The Economic Definition of Ore* (Mining Journal Books Limited).

Lerchs, H and Grossmann, I F, 1965. Optimum design of open-pit mines, *Trans CIM*, LXVIII:17-24.

Roman, R J, 1973. The use of dynamic programming for determining mine-mill production schedules, in *Proceedings 10th Symposium on Application of Computer Methods in the Mining Industry*, pp 165-169 (South African Institute of Mining and Metallurgy: Johannesburg).

Schaap, W, 1981. A theory of compound decisions on mill-dump cutoff grades, *Geologie en Mijnbouw*, 60:237-244.

Whittle, J, 1988. Beyond optimization in open pit design, in *Proceedings First Canadian Conference on Computer Applications in the Mineral Industry*, (Eds: K Fytas *et al*), pp 331-337 (Balkema: Rotterdam).

Whittle, J, 1989. *The Facts and Fallacies of Open Pit Optimisation* (Whittle Programming Pty Ltd: Melbourne).

Wooller, R, 2004. Optimising multiple operating policies for exploiting complex resources – an overview of the COMET scheduler, in *Proceedings Orebody Modelling and Strategic Mine Planning*, pp 77-90 (The Australasian Institute of Mining and Metallurgy: Melbourne).