

Footprint Finder tool for sublevel caving

T Diering *Dassault Systèmes, Canada*

F Breed *De Beers Group of Companies, South Africa*

Abstract

This paper discusses a new software tool developed specifically for use in a sublevel caving project/mine. The tool includes a number of new algorithms or technologies which when combined provide a very efficient way to evaluate a complete sublevel layout in a matter of seconds.

The core technologies imbedded in the tool are simulation of material mixing for dilution forecasting using Template Mixing algorithm, optimisation of the draw factors or extraction percentages per ring, and rapid sequencing/scheduling of the resulting ore tonnages.

Examples are presented showing how the rapid analysis can allow sensitivity studies on tunnel spacing, mining rate and economics, face shapes and draw or extraction strategies. A real-world example from the De Beers Venetia Underground project is presented.

Concluding remarks will discuss how this new tool fits in with the existing industry tools for open pit, open stope, block cave and sublevel cave operations optimisation.

Keywords: *orebody evaluation, SLC mining method evaluation, SLC footprint assessment*

1 Introduction

Mining engineers within mining companies and consulting groups are continually looking for rapid ways to evaluate projects for different mining methods. This is demonstrated by the success of existing tools within the industry such as open pit optimisers, stope optimisers and the Footprint Finder tool within GEOVIA PCBC. (Isabel 2013). Any of these tools would look to identify a volume of rock to be mined and then a sequence in which the blocks within that volume will be mined. By applying a mining rate, an approximate mining production schedule can then be evaluated.

To date, very little has been published for such a tool applied to sublevel cave (SLC) mining projects. Many consultants will have used Excel sheets or adaptations of other optimisation tools. However, for caving operations, the addition of significant dilution which is strongly affected by mining sequence complicates the process. This paper describes a new tool, Footprint Finder for SLC (FFSLC), which tackles this problem.

For SLC, the dilution calculation process is made particularly complicated by the fact that the material extraction history from one level affects the material available in future deeper levels. Thus, the overall problem is quite non-linear and not well-suited to solutions using linear programming technologies. With open pits and open stopes, the usual assumption is that 100% of the material blasted will be mined (as close as possible taking into account relatively low levels of dilution). However, with SLC, the fractions mined (referred to as extraction percentages) in the paper are variable from as low as 40% to as high as 500% in some cases. In general, low extraction percentages are used near the top of the orebody in order to establish the caving process and limit early entry of dilution while higher extraction percentages are used deeper down to attempt to improve overall recovery of the orebody and any overlying material with reasonable grades. This adds another level of complexity in the overall sequencing and scheduling process, namely the estimation of the extraction percent to apply to each mining point (ring).

In most other mining methods, the detailed mining geometry (e.g. toes and crests in an open pit or draw points in a block cave) are simplified down to the optimisation or processing of a regular block model. A similar approach has been adopted for FFSLC. Material attributes within each block are used as inputs and

a relatively simple process can be used to estimate an undiscounted dollar value for each of these blocks. What one then seeks is an improved or potentially optimal way in which to mine the blocks (or fractions of blocks). However, in any realistic evaluation, a large number of sensitivity runs are required to assess mining economics, mining sequence, production rate, etc., so set-up time and execution speed are very important.

1.1 Basic steps in the SLC evaluation process

FFSLC was derived from the Footprint Finder tool in GEOVIA PCBC. Taking into account the special adaptations required for SLC, the following key steps have been implemented:

- Input of key control/economic/geometric parameters.
- Input of block model attributes.
- Assembly of blocks into proposed mining levels. (For example, we may have a 6 m block size but want to mine to 25 or 30 m spaced mining levels).
- Re-blocking rows and columns in the block model to convert from block resolution to a reasonable mining (selective mining unit) type resolution. This results in a number of vertical columns each consisting of groups of blocks from the block model which can then be evaluated.
- Evaluation of target columns which are able to be mined.
 - Assessment of the extraction percent to be taken from each level for each of the vertical columns.
 - Simulating the material flow/dilution/mixing which takes place as the material is mined from each level of each vertical column. This is the heart of the overall process.
- Defining a sequence in which each component of each vertical slice will be mined.
- Evaluating the tons and grades to be mined in a number of time periods given a target production rate, relevant mining constraints and mining sequence.
- Presenting results to users for rapid evaluation and comparison.

Key steps in the above workflow are described in the following sections after which examples are presented. The examples also give an idea as to the typical runtimes which can be expected.

2 Core algorithms with FFSLC

The primary technologies within FFSLC are as follows:

- Geological block to mining unit conversion.
- Rapid simulation of vertical mixing within vertical columns.
- Optimisation strategy for extraction percent optimisation.
- Sequencing and evaluation of post-mixing mining units.
- Other adjustments and refinements.

2.1 Geological block to mining unit conversion

As computers get faster, block models tend to get larger (with smaller blocks). Conversely, mining companies are seeking to mine using ever-larger equipment enabling larger mining dimensions. Other considerations are geological resolution and orebody variation. However, for this paper there is a need to manipulate blocks from a block model into larger blocks representing reasonable 'mining blocks' which could simulate a group of, say, 5–10 rings per tunnel. Thus, the mining blocks are typically 5–20 m horizontally by 15–30 m vertically.

Preparation of mining units is relatively simple but does need to take account of some complicating factors such as variable mining level spacing (e.g. 25 m for five levels and then 30 m for the remaining deeper levels), separate mining of development and production tonnes (from tunnels and rings respectively), internal waste, ore blankets, overlying dilution or allowance for failure of open pit material and others.

2.2 Rapid simulation of vertical mixing within vertical columns

The algorithm described here is a stochastic method and relies on rules and probabilities. They do not have a foundation on physical principles. The tool should be used and calibrated accordingly.

Suppose we have a vertical column consisting of a number of slices (i) equivalent to one slice for each mining level. If we extract material from slice i , then we will extract material from that slice as well as all the slices (j) above i back to the top level 1 in the model. Thus the material mined from slice i (m_i) will be a summation of the contributions from each of the above slices s_{ij} . In normalised terms, the material mined m_i is related to the extraction percent E_i for that level.

$$E_i = \sum (s_{ij}) \text{ for } j = 1, 2, 3, \dots, i \quad (1)$$

Thus, if we can compute all the slice fractions (s_{ij}), then we can evaluate the material to be extracted from each slice (E_i). The problem is that s_{ij} cannot be pre-computed. As soon as we modify the amount we want to extract from any slice i , the s_{ij} values for all the remaining slices will change and require to be updated.

Initial attempts to try to simulate this within Excel from which an algorithm could be coded up were not successful. The problem is that the sum of the s_{ij} values cannot exceed 1 for any row or column in the overall ij matrix. A constraint is that you cannot mine more material from a slice than what exists in that slice. Getting something to work when extraction percentages were around 100% was relatively simple, but as soon as high extraction percentages (e.g. 200–500%) were tested, the Excel models failed.

The next approach was to consider the Template Mixing algorithm (Diering 2007) which has been successfully used for both block and sublevel cave projects in full three dimensions. The basic algorithm was simplified and reduced to a one-dimensional version denoted TM1D. This is a recursive algorithm which starts at the top of the column and requests material from the current slice. It then works recursively up the other slices to replenish material from the higher slices into the lower slices using pre-assigned block fractions (similar to the probabilities used in conventional cellular automaton programs). Special conditions are needed at the topmost slice to decide whether the slice is interacting with more external dilution or something else such as the base of an open pit.

Most material mixing algorithms are number crunching intensive, so efficient coding of the above approach was essential. A typical project might contain 1,000 vertical columns, each of which might require to be evaluated from 100 to 10,000 times for optimisation iterations and each column could have up to 20 slices from which variable depletion would be required. Each depletion would require from 1 to 20 depletions into the overlying slices. Further, each slice needs to be depleted in smaller increments so that we might require from 10 to 100 increments of depletion per slice. This could result in a total of a few billion depletions. This demonstrates the need for an efficient internal mixing algorithm.

The amounts extracted from each slice are conveniently presented in Excel array format as shown in Figure 1. The slices are numbered from 1 to 7 from bottom to top. The extraction percentages per slice (where mining takes place) are shown in the bottom row (1.000, 1.000, ...). Total extracted from each slice is shown in the right column (0.476, 0.636, etc). Look at column 2 and row 2 (top) as an example. This shows that 100% of slice 2 has been extracted which is made up of 48, 16, 8, 5, 4 and 21% from slices 2, 3, 4, 5, 6, 7. Or one can say that 48% of slice 2 was recovered from slice 2. Sixteen percent of slice 3 was recovered from slice 2, etc. From this, we can perform direct comparison with experimental marker trials from Ridgeway (Villa 2012; Power 2004). We can see good agreement. The primary recovery from each slice is further controlled by two key template mixing variables, namely percent frozen and erosion rate.

Percent frozen variable will ‘freeze’ a portion of a slice preventing it from being mined as primary recovery. However, a small portion of the frozen portion may be ‘eroded’. This is demonstrated in Figure 1, column 1 for slice 1 (at the bottom). By increasing the extraction percent from 1 (100%) to 1.4 (140%) (Figure 2), the primary recovery increased from 48 to 49%. What this shows, and has been demonstrated in marker studies and real mining, is that mining much more from a ring does not necessarily result in extracting much more ore (Power 2004). The simulated dilution in this example jumps from 19 to 40%.

Figure 2 shows the results of increasing the extraction percent for slices 1 and 2 from 100 to 120%, and 140% respectively. Overall recovery of the target slices (ore) increases, but so does the total simulated dilution (slice 1).

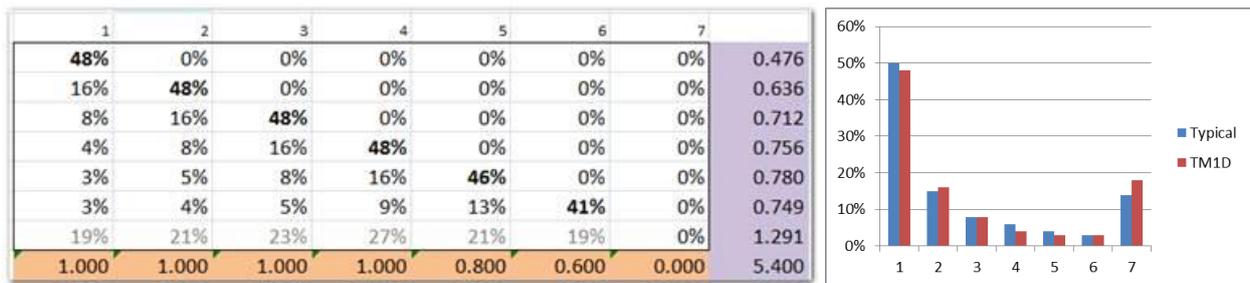


Figure 1 Slice contributions to overall extraction percent (100% down to 60%)

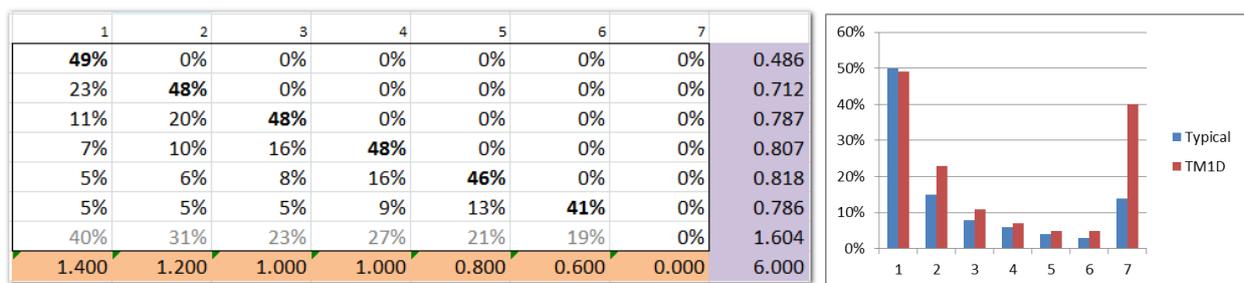


Figure 2 Slice contributions to overall extraction percent (140% down to 60%)

The above and related testing showed that the TM1D algorithm was performing very well and giving results in agreement with typical marker behaviour and which conserved material.

2.3 Optimisation strategy for extraction percent optimisation

The next key component of the overall workflow is the optimisation of the extraction percentages from each slice of each vertical column. In general, it is to be expected that each vertical column will show slightly different behaviour. In order to perform the optimisation, some constraints are required for each slice. Thus, for each slice we have to determine the order of extraction of the slice as primary, secondary, tertiary, quaternary, etc. (referred to as PSTQ). Next, minimum and maximum values for the extraction percent can be set up. For example, for primary extraction of the top level, we might typically set a target of 50% extraction with a minimum of 40% and a maximum of 60%. A lower level might have a target of 140% extraction with a minimum of 100% and a maximum of 200% (Figure 3). This leads the algorithm to do a constrained optimisation of the value of each column where value is related to a discounted or non-discounted evaluation of a dollar value (depending on the user application).

	A	B	C	D
1	Slice	Ext_per	Ext_min	Ext_max
2	17	0%	0%	0%
3	16	0%	0%	0%
4	15	0%	0%	0%
5	14	60%	30%	120%
6	13	80%	40%	160%
7	12	100%	50%	200%
8	11	100%	50%	200%
9	10	100%	50%	200%
10	9	100%	50%	200%
11	8	100%	50%	200%
12	7	100%	50%	200%
13	6	100%	50%	200%
14	5	120%	60%	240%
15	4	120%	60%	240%
16	3	120%	60%	240%
17	2	120%	60%	240%
18	1	120%	60%	240%

Figure 3 Example of initial, minimum and maximum extraction percent values

What values of extraction percent will yield the maximum total dollar value?

An informal background in non-linear optimisation techniques suggested the following overall approach for each vertical column:

- Start from an initial estimate for E_i for each slice i .
- Make a new guess by changing one of the E_i estimates. The choice of i and whether to increase or decrease the value can be done in a pseudo random manner. Don't make changes outside of the specified constraints.
- Re-evaluate the mixing array based on the updated E_i values and then evaluate the updated dollar value. Keep the new estimate if this is better than the previous estimate.
- Continue the above for the number of iterations set by the user (typically 500 to 10,000).
- Periodically introduce a much larger change to the E_i value. This is loosely equivalent to a genetic 'mutation' used on genetic algorithms. It is aimed to see whether a better value can be obtained by nudging the solution away from a local maximum in search of a larger value as part of a different 'hill'.

There are no mathematical guarantees that the above algorithm will find the global solution for every column, but all work to date has shown that the algorithm works well. In general, if a user suspects that a better strategy might be available, then this can be achieved by changing the initial estimate. For example, if we do one run starting from all minimal values and a second from all maximal values and the result is the same, that is likely the true maximum.

An early example showing how the extraction percent profile improves or changes versus dollar value is shown in Figure 4.

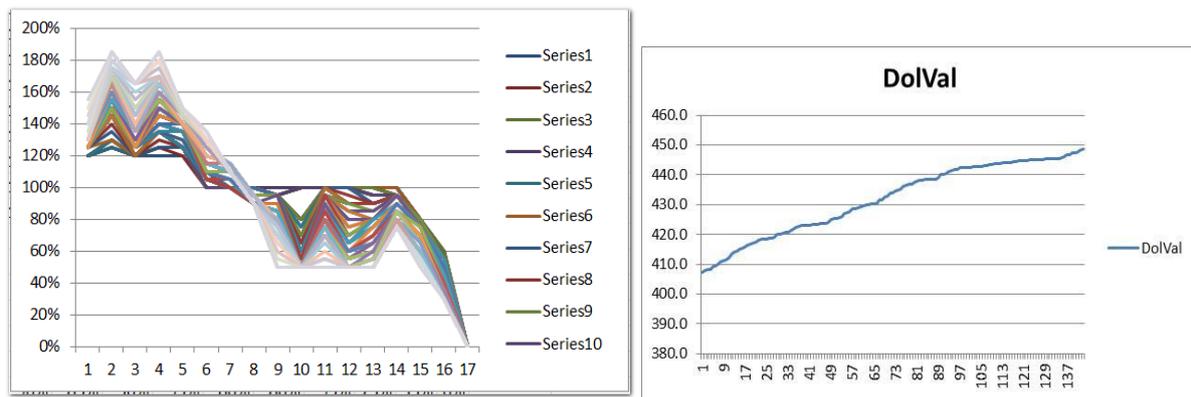


Figure 4 Extraction percent profile and dollar value improvement with iterations

2.4 Sequencing and evaluation of post-mixing mining units

Mining sequence can have a huge impact on the overall schedule results. A good sequence will result in good productivity via many active tunnels on each level, good access to the higher grade material and not all tunnels on a level needing to start up simultaneously. Similarly, the appropriate lead/lag between levels can ensure smooth transition when old levels end and new levels start.

The sequencing tool in FFSLC uses a baseline distance for each mining unit to order them appropriately. Mining units can be further divided into sectors which can be sequenced together or sequentially as required. The baselines are determined from user input mining shapes as shown, for example, in Figure 5. The shape is shown to the left and the result of mining using this shape from northeast towards southwest is shown to the right for four levels.

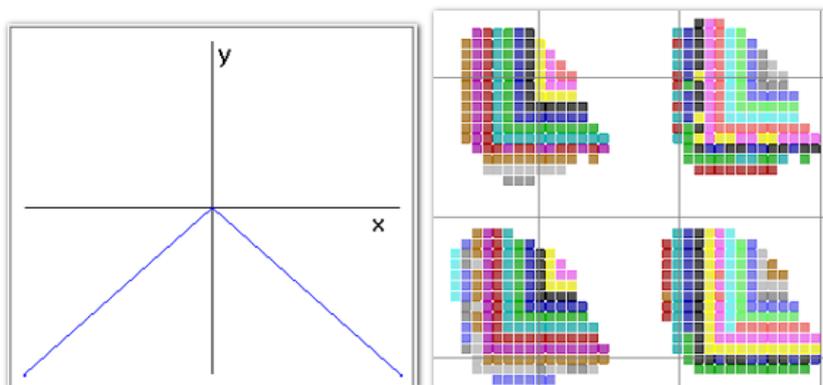


Figure 5 Mining sequence management using a simple XY curve

The origin and direction of mining are set independently.

Once the direction has been set and the material mixing evaluated, the blocks mined are simply grouped by mining periods to yield the final mining schedule.

2.5 Other adjustments and refinements

The previous sections have described the overall algorithm. But, as the saying goes, the devil is in the detail. Some of the refinements required to solve real-world problems are listed below.

- Variable level spacing.
- Treatment of internal low grade/dilution zones.
- Treatment of higher dilution zones or air (in the case of open pit exposure).

- Mining shut-off versus total shut-off for footprint delineation/extraction percent calculations.
- Simulation of development tonnages pre-mined from development tunnels.
- Simulation of various horizontal tunnel spacings.
- Limiting maximum mining levels or sectors.
- Changing strategies from ore recovery to time of recovery to dilution management.

3 Some initial test cases

Initial testing for the new developments focuses around the extraction percent optimisation, calibration of mixing to industry reported primary, secondary and tertiary level recoveries and comparison to other software results.

3.1 Time of mining versus dilution versus ore extraction

During the early testing and development of the program, specifically the optimisation of the vertical columns, the behaviour of the extraction percentages was really interesting.

The main findings are as follows:

- The tool would be basically useless without the inclusion of time effects in the dollar value calculations.
- With no discounting of dollar value for deeper levels versus upper levels, the program will simply maximise the extraction of the lower levels and minimise the extraction of the higher levels. It effectively tries to turn the SLC into a block cave.
- With extreme high discounting of dollar value, the program will mine as much as possible from each level to the incremental or mining shut-off. This is similar to a strategy where dilution is brought as close to the active mining area as possible. While this approach is frequently used and is attractive from an operational standpoint, it is unlikely to be optimal in most cases.
- If dilution material has reasonable value, then much higher extraction percentages are generated.
- In general, it is very unlikely that any human would be able to guess the right solution.

Figure 6 shows five examples. The specific input values for each run are not important here. What is more important is that the combination of time discounting and variable value of the overlying dilution material can have a significant impact on the optimal extraction profile result.

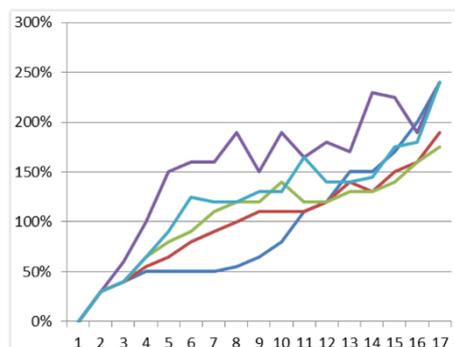


Figure 6 Variability of extraction percent profiles with changing dilution and discounting assumptions

4 Application areas for the new tool

Once a basic project has been set up, it is very easy and quick to perform sensitivity analysis on the following:

- Metal price versus mining costs (repeat various runs by changing economic profile settings).
- Simulate tunnel spacing (vertical). This will change the way the block model levels (vertical) are grouped into slices for the model.
- Mining rate sensitivity. This is done by mining more or less blocks in any period.
- Change constraints and targets. Allows for different extraction percent strategies to be compared.
- Discount rate. As the discount rate goes up, the deferring mining of material gets more expensive, so the program will try to mine more material sooner. For a low discount rate, mining of material is deferred as late as possible so as to minimise dilution entry.
- Ore blanked thickness study. It is easy to change the simulated location of the top mining level, thereby looking at different sized ore blanket.
- Horizontal mining rate sensitivity. The number of rings blasted per tunnel per month can be converted into a horizontal maximum mining rate which in turn can be used to constrain mining rate.
- Number of levels active sensitivity. Usually, having more levels active at once will enable a higher production rate.
- Mining direction. This simply changes the sequence in which blocks will be mined. This can be done per sector.
- Sequence shape (leads/lags). This is managed via simple changes to a single XY curve (or shape definition).
- Extraction percent optimisation. Look at ways to increase project value by extracting the right tonnes from the right places within each vertical column.
- Mining constraints such as maximum production from a level or sector can be added as constraints.
- Do sensitivity on horizontal tunnel spacing. This is only an indirect simulation and is achieved by varying the percent frozen and erosion rate parameters. As horizontal tunnel spacing increases, percent frozen will also increase, so the primary recovery of material from rings will decrease. If there are enough active levels, this primary material (which is not recovered) might be able to be recovered from deeper levels.

5 Case studies

Unfortunately, getting real data on which to conduct real studies seems to be getting ever more difficult as mining companies seek to keep data confidential against general publication. For this reason, detailed results are limited. To date, however, the tool has been used on around 20 different projects at different levels of complexity and this has exposed a number of different geometries and mining conditions from scattered, to near-vertical to inclined deposits with marginal economics, to complete mining history for completed projects.

The following examples are presented here:

- Pocket Aces (fictitious simulated diamond deposit).
- Walker Bay (fictitious simulated Au/Cu deposit).
- Venetia Mine underground project.

5.1 Pocket Aces simulated diamond deposit

This project has the following characteristics. It is near-vertical, well-shaped and well-suited to either block cave or SLC depending on simulated rock strength, grades, etc. The example consists of six levels with three additional layers of marginal material above, with one for ‘external’ zero grade dilution. In this example, overall recovery would be expected to be quite good since there are multiple opportunities to recover material from lower levels.

Figure 7 shows a typical sequence used in the evaluation. Six levels are plotted in ‘offset’ mode so that each level is clearly visible without being obscured by higher levels.

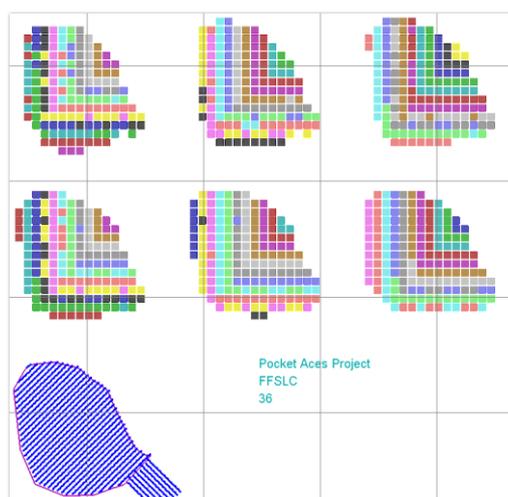


Figure 7 Six levels and overall layout for Pocket Aces example

Extraction percentages are then optimised (in less than 1 min in this example) to give an alternative extraction scenario. The results are shown in Figure 8.

Level	Tons In-Situ	CPHT	Level	Tons 1 Iteration)	CPHT	LevelN	Tons (300)	CPHT
L370	2,303,190	49.43	L370	2,735,344	45.95	L370	4,061,480	38.25
L395	2,333,975	49.99	L395	2,348,685	46.96	L395	3,928,733	41.35
L420	2,441,170	48.95	L420	2,451,212	46.83	L420	3,592,197	44.20
L445	2,472,213	49.45	L445	2,477,302	48.20	L445	3,126,039	47.65
L470	2,503,330	51.08	L470	2,040,075	49.88	L470	2,486,993	50.20
L495	2,568,323	52.34	L495	1,567,770	51.38	L495	1,529,748	52.19
B520	-	0.00	B520	-	0.00	B520	-	0.00
B545	-	0.00	B545	-	0.00	B545	-	0.00
B570	-	0.00	B570	-	0.00	B570	-	0.00
X595	-	0.00	X595	-	0.00	X595	-	0.00
Total	14,622,200	50.24	Total	13,620,388	47.91	Total	18,725,190	44.34

(a)
(b)
(c)

Figure 8 Tons and grades for in situ versus one iteration versus 300 iterations for extraction percentages

Figure 8(a) shows in situ tons without mixing for 14.6 Mt at 50.24 carats per hundred tonnes (cpht). Figure 8(b) shows the results with one iteration (not optimised) using the simple constant extraction percentages typically used in current projects. This shows 13.6 Mt at 47.91 cpht. Figure 8(c) shows the results with 300 iterations to optimise each vertical column resulting in 18.7 Mt at 44.34 cpht. The simulated discounted dollar value for the two runs goes from \$696 M to \$820 M. Although somewhat artificial, the example does clearly show that the optimisation process is effective.

5.2 Walker Bay fictitious deposit

A similar exercise was conducted for another fictitious deposit. The example is included here simply to show visually the impact of the changes to the extraction percent profile before and after optimisation. Figure 9 shows blocks from the block model (used to construct the vertical columns) before and after optimisation. The overall resulting strategy to keep extraction percentages near the upper ore/dilution interface low and increase extraction near the base of the deposit is clear. Remember that this strategy is heavily dependent on the discount rate strategy which affects the value of blocks mined sooner versus later.

This example also shows how the blocks were re-blocked both horizontally and vertically to simulate the mining levels and geometry.

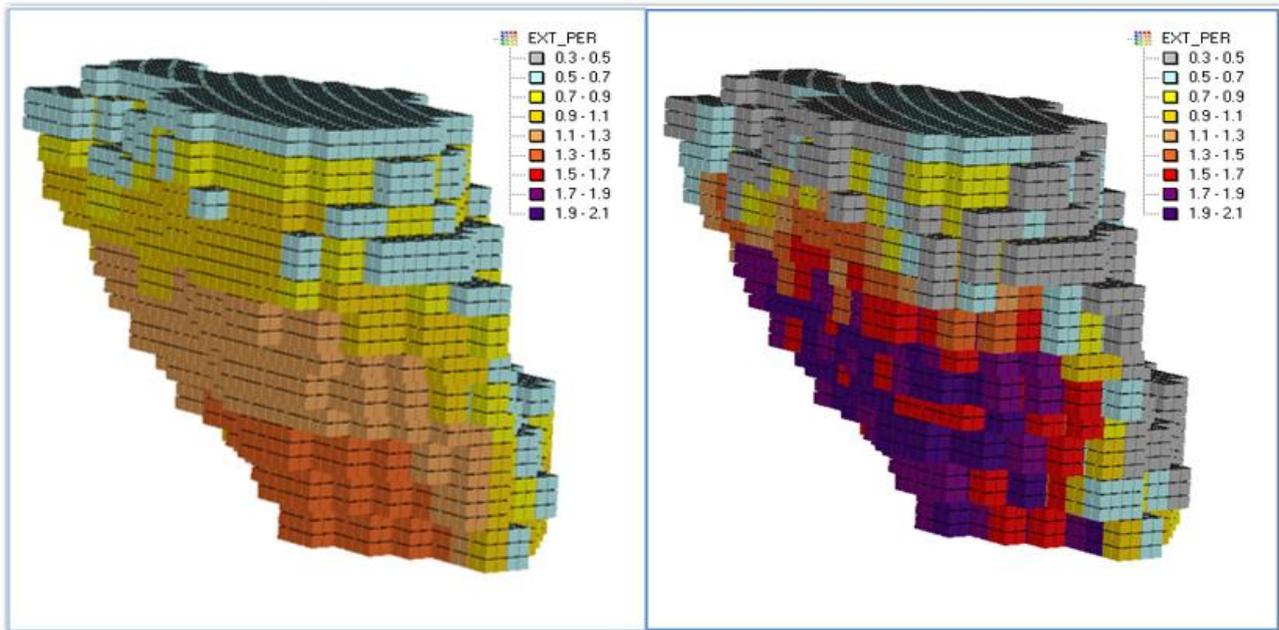


Figure 9 Extraction percentage of blocks from the block model before and after optimisation

The optimised case increased in value by nearly 6% (Table 1), while also increasing the average gold and copper grade by 4.64% and 1.68% respectively.

Table 1 Tons and grade comparison for base and optimised cases

Scenario	Tonnes	Volume	Dollars	Discounted dollars	Au (g/t)	Cu (%)	Time per run
Base	42,224,475	15,275,921	2,362,426,484	1,531,961,250	1.57	0.68	22 s
Optimised	42,437,498	15,365,461	2,501,376,824	1,621,805,654	1.65	0.69	3 min 20 s
Difference	0.50%	0.59%	5.88%	5.865%	4.64%	1.685	

5.3 Venetia Mine underground project

Venetia Mine open pit started in 1992 and currently mines approximately 5.4 Mt/annum from the K01 orebody. On completion of the open pit, the mine will transition to an underground operation with 4 Mt/annum mined from the K01 orebody 1.9 Mt/annum mined from K02 orebody using the SLC mining method. K02 is capable of producing 40 Mt/annum but is constrained by the process plant which is capable of a throughput of 5.9 Mt/annum when processing the underground run-of-mine material.

During 2017, the Venetia underground project undertook a project to optimise the production output from K01 and K02. The objective was to optimise the ring extraction percentage for K01 and K02 and optimise the K02 layout.

Although the focus of this case study is on the K02 orebody, Footprint Finder was effectively used in K01. The optimisation focused on the local extraction percentage where applied and this resulted in a 3% improvement in recovered carats.

The K02 orebody optimisation process was driven by a strategic initiative to attempt to help with the open pit transition to underground mining by accelerating K02 production. An iterative process for the optimisation was developed and the process was as follows:

1. Define extraction strategy.
2. Model extraction scenarios.
3. Evaluate results and conclusion.

The next steps were to use the SLC Footprint Finder on the K02 orebody to test and compare the results from the iterative process described above.

The base case scenario for K02 was an eight-level SLC with a high extraction percentage applied to lower levels to draw the 200 m lift cave.

5.3.1 Define extraction strategy

Analyses of the block model showed that there were three distinct grade groups in the K02 orebody that needed to be considered. The low grade (LG), medium grade (MG) and high grade (HG) areas. These were identified and are indicated in Figure 10. In the east–west view, which is the direction of mining, the grade varies significantly. However, the grade along the tunnel does not vary, as can be seen in the north–south view.



Figure 10 Grade variability for K02

Research (Villa 2012) shows that a recovery curve can be used to calibrate the production scheduling in the SLC. K02 optimisation used a recovery curve to focus on post-primary recovery to identify the areas where high or low extraction needs to be applied. Villa (2012) proposed that primary, secondary, tertiary, quaternary and quinary recovery levels are represented as in Figure 11. The figure shows that 40–50% of the extracted tonnes in a ring could come from levels above the current mining horizon.

Various extraction scenarios were created using the grade zone to determine what the ideal extraction percentage should be depending on whether the post-primary recovery material is zero grade, low grade, medium grade or high grade. The conclusion was that if there was higher grade material available for post-primary recovery, the extraction percentage was increased, and if lower grade material was available, the extraction percentage was decreased.

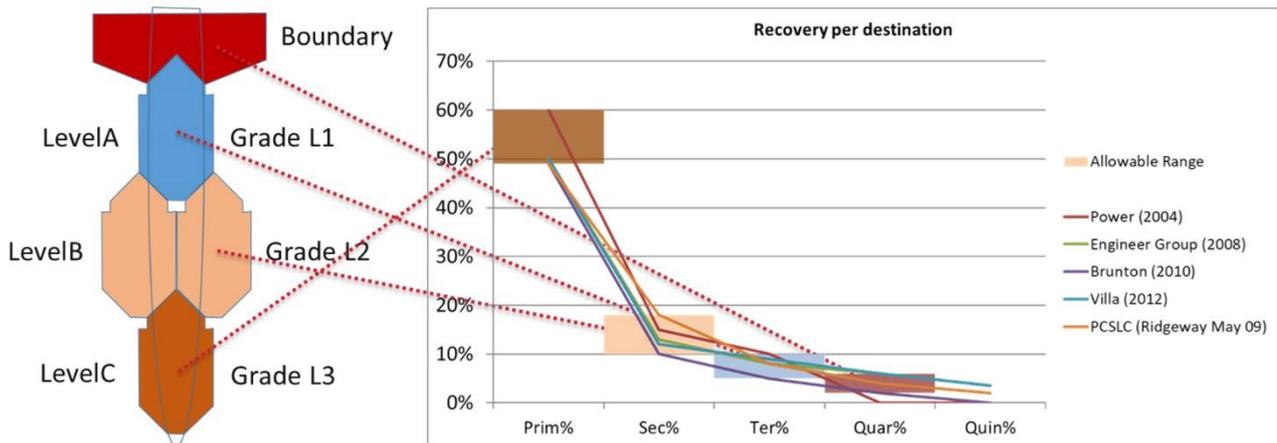


Figure 11 Recovery curves applied in the SLC calibration (Villa 2012)

5.3.2 Model extraction scenarios and deplete

Figure 12 represents the extraction models built and evaluated on recovery, grade and financials. The aim was to first test the effect of removing the LG zone from the plan, then the MG zone.

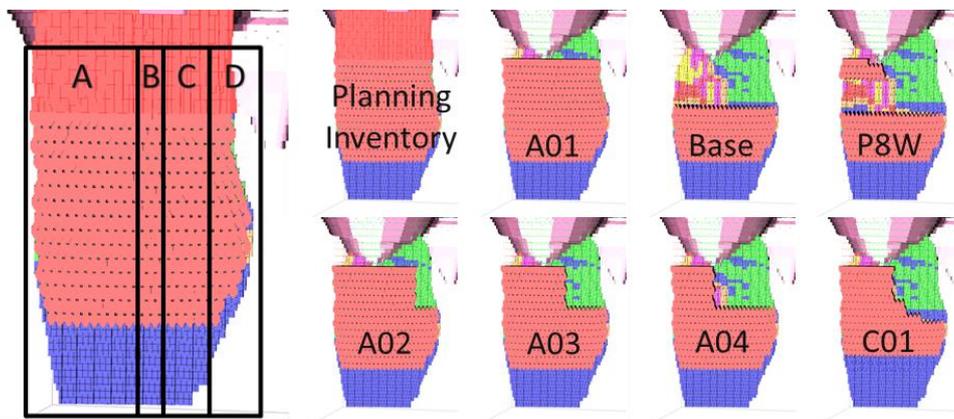


Figure 12 Graphical representation of the scenario evaluated

5.3.3 Evaluate results and conclusion

For the financial modelling, an operating cost differential was applied for caved tonnes versus blasted tonnes which benefitted the higher extraction percentage scenarios. The scenarios were modelled and the financial results were evaluated. Figure 13 shows the NPV, run-of-mine grade and dilution percentage for each of the above scenarios.

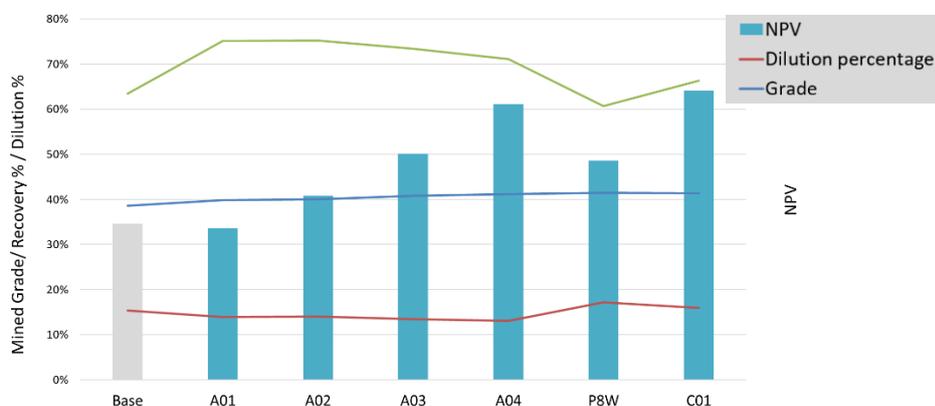


Figure 13 Financial analysis for iterative optimisation process

It was concluded that scenario C01 was the best scenario to apply due to its improved NPV and recovered grade. Improved results were achieved when the LG zones were excluded from the plan, and where planned rings were located below the LG zone. The results showed that the planned rings extraction percentages should be decreased. Scenario A04 had a higher capital requirement as the number of levels increased significantly and it also showed no benefit for lower operating cost with a high extraction percentage.

The total duration of the project was three weeks.

5.3.4 Test and compare results from iterative process using Footprint Finder

With a good understanding of the resource, it was very easy to setup Footprint Finder to start the optimisation. The setup process described in Section 1.1 was done and the first evaluation run was completed. A couple of exploratory evaluation runs were completed to test the functionality and sensitivity to the inputs. The revenue parameter had the greatest impact and it was good that negative revenue could be applied to dilution, to penalise the dilution entry, as this has a major impact on processing plant capacity. The optimisation quickly started to plateau and it was found that around 1,000 iterations were enough for the K02 evaluation.

The flexibility around the range of extraction percentages to apply to the planned rings can limit the extraction to just removing swell or even apply a zero extraction percentage. The process quickly identified areas of the orebody that did not contribute to the NPV and could effectively be removed from the mine plan. Comparing the results from Footprint Finder to the long iterative process, the general layout of the mining level was similar and thus the C01 layout was kept unchanged to continue the optimisation process.

The next step in the process was to limit Footprint Finder to only the targeted areas. By excluding all material outside of the layout, the expected extracted tonnage profile was more accurate. During this optimisation run, the number of iterations was increased to 10,000 and the incremental change reduced from 5 to 1%. This increased the running time but the optimisation results were more accurate.

Figure 14 shows the rings on the horizontal axis and the extraction percentages on the vertical axis. The comparison clearly shows where there was an under-extraction in high-grade areas and over-extraction in lower-grade areas. These findings highlighted the importance of setting the maximum limit as the algorithm will continue mining ignoring over-extraction risks, e.g. creating an airgap. Using global percentage (per level) is not optimal. It is important to apply local extraction percentages (per ring) as this will ensure an improved value proposition for the orebody.

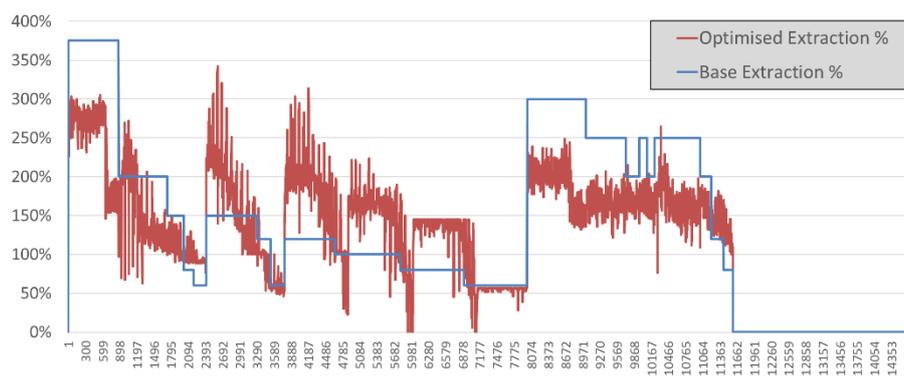


Figure 14 Optimised extraction percentage versus global planned extraction percentages

The objective to optimise the ring extraction percentages was achieved and the result was a 13% increase in recovered carats from the K02 orebody.

The financial impact of this increase is clearly shown in Figure 15 where the NPV of the Footprint Finder scenario exceeds all previous scenarios. The optimisation delivered the best results in operating cost, grade, capital efficiency and dilution.

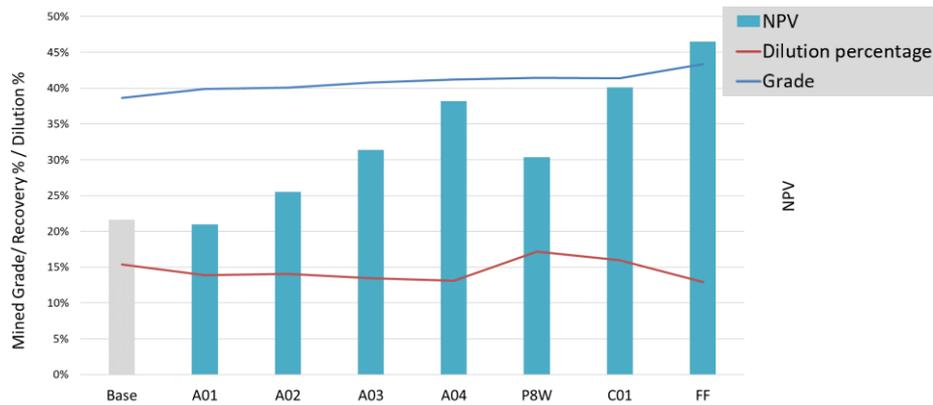


Figure 15 Financial analysis for iterative optimisation process, including the Footprint Finder result

The most noticeable improvement compared to the iterative process is in duration of the evaluation period. The full evaluation took two days and this included training, bug fixes and multiple runs with multiple parameters.

6 Concluding remarks

Footprint Finder is a quick and effective tool to evaluate a sublevel cave mining method.

This tool can be rolled out at any level of study and operation. In early study phases, it can evaluate the expected tonnage and grade profile by only using the block model. As an optimisation tool, it can be used in later study phases to optimise layout and extraction percentages. By using the past tonnes function, the optimisation can also be applied to any operating sublevel cave.

The purpose of an optimisation tool is to help the engineer, or planner, select the best input to produce a detailed plan. An optimisation tool should be able to derive a solution quickly and easily. The processing time and total duration of the Footprint Finder project show an effective tool that produces quality results in a fast and repeatable process.

The new Footprint Finder tool for sublevel caving fills an important gap not covered by other tools such as open pit, block caving and stope optimisers.

Acknowledgement

The authors acknowledge and thank Dassault Systèmes, Canada and De Beers Consolidated Mines for the financial assistance to conduct the research and development leading to the development of FFSLC. We also thank De Beers Venetia mine for permission to use selected results in this paper.

References

- Diering, T 2007, 'Template mixing: an alternative depletion engine for block scheduling', *Proceedings of the 33rd International Symposium on the Application of Computers and Operations Research in the Mineral Industry*, Gecamin Ltda, Santiago, pp. 313–320.
- Isabel, A 2013, *Efficient Evaluation of Block Cave Footprints for A Range of Elevations*, Gemcom white paper, Gemcom Software International Inc., Vancouver.
- Power, G 2004, 'Full scale SLC draw trials at Ridgeway Gold Mine', in A Karzulovic & MA Alfaro (eds), *Proceedings of MassMin 2004*, Instituto de Ingenieros de Chile, Santiago, pp. 225–230.
- Villa, D 2012, *Calibration of a Mixing Model for Sublevel Caving*, MSc thesis, The University of British Columbia, Vancouver.