Pseudoflow Explained

A discussion of Deswik Pseudoflow Pit Optimization in comparison to Whittle LG Pit Optimization

Julian Poniewierski
JMPstart Mine Technical Service

1. Introduction

The purpose of this document is to inform the user about Deswik Pseudoflow, within the context of the mining industry's most accepted pit optimization process that uses Whittle software based on the Lerchs-Grossman (LG) algorithm.

In summary, both are variations of network flow algorithms that achieve the same result, with the pseudoflow algorithm being a computationally more efficient algorithm developed some 35 years after the original Lerchs-Grossman algorithm (1965).

It took 20 years from the formulation of the LG algorithm for it to be incorporated in the first commercially available software (Whittle Three-D), with another 10 years before it became a mainstream approach to open pit optimization.

In the past 30 years, Whittle and the LG algorithm have become synonymous with the optimization of open pits, and now suffers from having become a generic term for the process of pit optimization– similar to the genericization of the Hoover brand with vacuuming in the UK (such that in the UK, people hoover their carpets).

It is now 15 years since the formulation of the pseudoflow algorithm, and there are at least three commercial implementations available, including Deswik's implementation.

Just as "hoovering" does not need to be a Hoover vacuum cleaner – indeed the Dyson cyclonic vacuum is recognized as far more efficient at vacuuming– the pseudoflow algorithm should now be used to replace the LG algorithm in "whittling" your pit optimization.

It should be noted that the Deswik implementation is not constrained (nor is it aided) by the extensive set-up parameter tables and inputs provided in the Whittle software for cost and revenue calculations. For Deswik's pseudoflow implementation, the user is required to calculate the revenues and costs for each block in the block model used, and is required to do their own block regularization within the Deswik.CAD environment. The user is thus in full control of how the costs and revenues are calculated and assigned. This does however require the user to be fully familiar with their block model, cost structures and revenue parameters (which we believe is a "good thing"). This enables the set-up to be as flexible as required by the user (unconstrained by template set-up dialogs).

2. History of pit optimization

2.1. MANUAL PROCESS

Prior to the development of computerized methods of pit optimization and pit design, mining engineers used manual interpretation methods with evaluation on manually drawn cross-sections (on paper, linens, or film), and then a manual pit design.

In the manual method, a simple optimization of economic pit depth was usually carried out by hand calculator (or slide rule) for regular shaped orebodies using incremental cross-sectional areas, for ore and waste, and an overall pit slope. The incremental stripping ratio (the ratio of the tonnage of waste that must be moved to access the next mass increment of ore that will be accessed) on each cross-section was compared against the break-even stripping ratio for the estimated ore grade and appropriate revenue and cost inputs.

The final pit shell was then produced by drawing increasingly larger pit shells on cross section such that the last increment had a strip ratio equal to the design maximum.

This was a very labor intensive approach and could only ever approximate the optimal pit. The design had to be done on a large number of cross sections and was still inaccurate because it treated the problem in only two dimensions. In cases of highly variable grade the problem became extremely complex, and relied heavily on the "gut feel" of an experienced designer using trial and error.

2.2. FLOATING CONE

Pana (1965) introduced an algorithm called Moving (or Floating) Cone. The method was developed at Kennecott Copper Corporation during the early 1960s (McCarthy, 1993) and was the first computerized attempt at pit optimization, requiring a three dimensional computerized block model of the mineral deposit.

The projected ultimate pit limits are developed by using a technique of a moving "cone" (or rather a frustum of an inverted cone – that is, the "pointy" end has been cut to a minimum mining area). The cone is moved around in the block model space to generate a series of interlocking frustum shaped removal increments.

However, the shortcoming of this approach is that it creates overlapping cones, and it is incapable of examining all combinations of adjacent blocks.

For this reason, the algorithm fails to consistently give realistic results.

Mintec/MineSight (being a US based company and supplier of early solution to Kennecott) were an early implementer of the floating cone algorithm (and may still offer it in their solution suite).

2.3. LERCHS - GROSSMAN

It was also in 1965 that Lerchs and Grossmann published a paper that introduced two modeling approaches to solving the open pit optimization problem. The Lerchs-Grossman (LG) algorithm is well documented in the technical literature (Lerchs and Grossman, 1965; Zhao and Kim 1992; Seymour, 1995; Hustrulid and Kuchta 2006).

The LG method was based on a mathematical technique which was unusable in practice until a practical optimization program called Whittle Three-D was developed by Jeff Whittle of Whittle Programming Pty Ltd in the mid-1980s.

Two methods to the solution of open pit optimization were detailed by Lerchs and Grossmann, being a Graph Theory algorithm, which is a heuristic approach, and a Dynamic Programming algorithm, which is an application of an operations research technique. Both methods gave an optimum pit limit for an undiscounted cash flow – based on an economic block model of an ore body and its surrounding waste, and determined which blocks should be mined to obtain the maximum dollar value from the pit.

The LG methods took into account two types of information:

1. The required mining slopes. For each block in the model, the LG method needs details of what other blocks must be removed to uncover it. This information is stored as "arcs" between the blocks ("nodes").
2. The value in dollars of each block once it has been uncovered. In the case of a waste block this will be negative and will be the cost of blasting, digging and haulage. In the case of an ore block, the removal cost will be offset by the value of the recovered ore, less any processing, sales, and other associated costs. Any block which can, during mining, be separated into waste and ore is given a value which reflects this.

Given the block values (positives and negatives) and the structure arcs, the LG method progressively builds up a list of related blocks in the form of branches of a tree (called a "graph" in mathematics). Branches are flagged as 'strong' if the total of their block values is positive. Such branches are worth mining if they are uncovered. Other branches with negative total values are flagged as 'weak'.

The LG method then searches for structure arcs, which indicate that some part of a strong branch lies below a weak branch. When such a case is found, the two branches are restructured so as to remove the conflict. This may involve combining the two branches into one (which may be strong or weak) or breaking a 'twig' off one branch and adding it to the other branch.

The checking continues until there is no structure arc which goes from a strong branch to a weak branch. At this point the blocks in all the strong branches taken together constitute and define the optimum pit. The blocks in the weak branches are those which will be left behind when mining is complete.

In effect, what the LG algorithm has done is to find the maximum closure of a weighted directed graph; in this case the vertices represent the blocks in the model, the weights represent the net profit of the block, and the arcs represent the mining (usually slope) constraints. As such the LG algorithm provides a mathematically optimum solution to the problem of maximizing the pit value (importantly, note that this is for an undiscounted cash flow value).

It should be noted that it is a mathematical solution. Except for the information given by the arcs, the LG algorithm "knows" nothing about the positions of the blocks – nor indeed about mining. The LG algorithm works only with a list of vertices and a list of arcs. Whether these are laid out in one, two or three dimensions and how many arcs per block are used is immaterial to the logic of the method, which is purely mathematical.

Also note that it took some 20 years between the publication of the LG method (1965, which was also the year that the floating cone method was computerized) and the first commercial available adoption of the LG method (Whittle's Three-D).

The basic LG algorithm has now been used for over 30 years on many feasibility studies and for many producing mines.

2.4. NETWORK FLOW SOLUTIONS

"In their 1965 paper, Lerchs and Grossmann indicated that the ultimate-pit problem could be expressed as a maximum closure network flow problem but recommended their direct approach, possibly due to computer memory constraints at the time. The LG algorithm was therefore a method of solving a special case of a network flow problem" (Deutsch, et al, 2015).

In 1976, Picard "provided a mathematical proof that a "maximum closure" network flow problem (of which the open cut optimization problem is one) were reducible to a "minimum cut" network flow problem, hence solvable by any efficient maximum flow algorithm. As a consequence, sophisticated network flow algorithms could therefore be used in place of the LG algorithm, and they can calculate identical results in a fraction of the time." (Deutsch, et al, 2015).

One of the first efficient maximum flow algorithms used in solving the open pit optimization problem was the "push-relabel" algorithm (Goldberg and Tarjan, 1988; King et al., 1992; Goldfarb and Chen, 1997).

"Hochbaum and Chen's study (2000) showed that the push-relabel algorithm outperformed the LG algorithm in nearly all cases. When the number of vertices is large, greater than a million, network flow algorithms perform orders of magnitude faster and compute precisely the same results." (Deutsch, et al, 2015).

Numerous authors implemented the push-relabel algorithm, and various heuristics and techniques were developed to maximize its performance. This was the algorithm that MineMax implemented in their first pit optimizer software offering.

Development of more efficient network flow algorithms have continued. The generally accepted most efficient algorithm currently available are the various pseudoflow algorithms developed by Professor Dorit Hochbaumn and her colleagues at University of California, Berkeley (Hochbaum, 2002, 2001; Hochbaum and Chen, 2000).

Pseudoflow methods give new life to the LG pit optimization. The "highest label" method implementation of the pseudoflow algorithm in particular is consistently faster than the generic LG methods and is also usually faster than the alternative "lowest label" method implementation of the pseudoflow algorithm. The increase in speed can be from two to 50 times faster than the LG methods, and theoretically much faster for larger problems (Muir, 2005).

3. Algorithm performance comparisons

Muir (2005) gave the most comprehensive analysis of the pseudoflow algorithm performance and a practical example of the identical results achieved in comparison to the LG algorithm in solving a pit optimization. These analyses and results were presented to the mainstream mining industry in the 2005 AusIMM Spectrum series publication: Orebody Modelling and Strategic Mine Planning. Key results of Muir's analysis are reproduced herein.

It should be noted that the code written by Muir (2005) is the underlying calculation engine that has been implemented in Deswik Pseudoflow.

As a check of the correct implementation of that code, the results for the Deswik implementation were compared against four publicly available test data sets from Minelib[1] (Espinoza et al, 2012). The specific data sets the results were checked against were for Marvin, McLaughlin, KD and P4HD. The Pseudoflow results were identical to the published results at Minelib.

Table 1 (from Muir, 2005) shows the relative run-times for several variants of both the LG and pseudoflow algorithms. It can be seen from these results that the "highest label pseudoflow priority queue" (HLPQ) implementation took just under 2% of the time it took for the standard LG algorithm to solve a 38 bench pit optimization problem.

Table 2 (from Muir, 2005) shows that the number of blocks and the profit value for the HLPQ solution was identical to the LG solution of the same 38 bench pit optimization problem.

The relative solution times shown in Table 1 are shown plotted in Figure 1.

In addition to Muir's paper, there are a couple of other known examples of published comparisons between the LG algorithm and flow network solutions to the pit optimization problem.

Jiang (2015) stated that the final pit limits from using a pseudoflow algorithm implementation versus the Whittle LG implementation have always been found to be materially the same, with any minor differences that are observed always being due to how the various implementations compute the slope angle constraints.

The push-relabel algorithm implemented by MineMax was compared to the LG algorithm by SRK (Kentwell, 2002) and was found to produce "the same results for the actual optimal pit calculations" (to within less than 0.01% - with the differences appearing to be due to block coarseness and slopes).

[1] http://mansci-web.uai.cl/minelib/Datasets.xhtml

Table 1 – Optimization times (seconds) to various pit levels for 220 x 119 x 38 profit matrix

BENCH	LG	LGS	LLP	LLPS	LLPQ	HLPQ
26	285	56	186	91	23	9
28	398	94	247	107	35	13
30	632	130	327	125	54	17
32	878	176	410	145	83	28
34	1157	243	480	152	107	27
36	1387	478	541	157	116	28
38	1527	628	556	160	126	29

LG	Normal Lerchs-Grossmann
LGS	Subset Lerchs-Grossmann
LLP	Lowest Label Pseudoflow (no priority queue)
LLPS	Subset Lowest Label Pseudoflow (no priority queue)
LLPQ	Lowest Label Pseudoflow (priority queue)
HLPQ	Highest Label Pseudoflow (priority queue)

(after Muir, 2005)

Table 2 – Statistics for level 38 for 220 x 119 x 38 profit matrix

	LG	LLPQ	HLPQ
Profit value	57 118 058	57 118 058	57 118 058
Blocks removed	95 228	95 228	95 228
Blocks remaining	830 754	830 754	830 754
Branches relinked	950 175	329 599	420 244
Branches pruned	1 703 036	459 454	638 088
Time (seconds)	1527	126	29

Figure 1 – Solution times for four pit optimization algorithms for different bench number pit problems

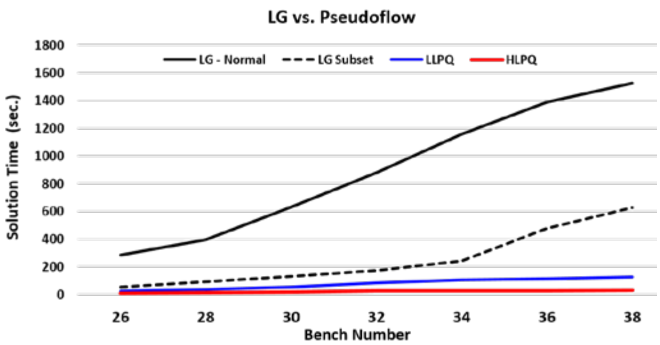

4. Modeling issues to note

Having shown that it has been proved that the pseudoflow algorithm will give identical results to the LG algorithm, it is appropriate to also point out that no algorithmic solution will provide the exact "true" optimization solution. There are a large number of approximations inbuilt into the algorithmic solution to the pit optimization problem as well as a number of common errors and uncertain assumptions used in the process.

The huge efforts dedicated to the development of sophisticated optimization algorithms is usually not matched by similar attention being paid to improving the correctness and reliability of the data used in the modeling exercise, as well as the correct use of the results of the modeling.

Some of the numerous sources of error, uncertainty and approximations in the process of pit optimization that need to be recognized are discussed below.

In summary, be aware that the process of pit optimization is based on coarse and uncertain estimated input parameters. Deswik therefore recommend that the user concentrate on the overall picture and getting as accurate as possible estimates of the "big ticket" items. And remember: "don't sweat the small stuff".

Deswik also advise to design for risk minimization of the downside in your assumptions, as per the scenario strategies advocated by Hall (2014) and Whittle (2009), but to also check the optimistic case upside scenario to determine infrastructure boundaries.

4.1. SOLUTION APPROXIMATION "ERRORS"

a. The effect of using blocks with vertical sides to represent a solution (a pit design) that has non-vertical sides. It is possible to output a smoothed shell through the block centroids, but note that this will not give the same tonnes and grade result as the block based optimization when the surface is cut against the resource model blocks.

b. Slope accuracy representation. The accuracy of the overall slope created in the modeling process with respect to slope desired to be modeled will depend upon the height and number of dependencies (arcs) used to define the slope. This will always need to be checked for suitability. Larger blocks will generally give less slope accuracy, and more, smaller blocks that allow greater accuracy will require more modeled arcs (block precedencies) and will slow the processing down. An accuracy tolerance of around 1° average error is usually considered acceptable.

c. Changes in converting a shell to a pit design. A difference of 5% in tonnes is quite common during this process. This is due to the approximation of the overall slope with respect to the actual design and effects of placement of haul roads on that overall slope.

d. Effect of minimum mining width on the bottom of a shell. Many pit optimizations are undertaken without consideration of the minimum mining width at the bottom of each shell – even when the package used provides such a facility. This will change the value of the selected shell used for design. At present Deswik's implementation of Pseudoflow does not have a tool to consider minimum mining width – but this is in the future development plans.

e. Effect of stockpiling. The pit optimization algorithms – both Whittle LG and Deswik Pseudoflow assume the value generated is the value that occurs at time of mining, and stockpiling delays the recovery of that value. Stockpiling for 10 or more years will mean that the time value of the block of ore stockpiled can be a fraction of the value used in the pit optimization. Mines with significant amounts of marginal stockpiled ore will suffer a significant oversize effect from the difference in when the algorithm values the block and when the value is actually generated.

f. If an elevated cut-off grade policy is used in the scheduling of the pit early in the pit's life, as a means of maximizing the NPV (Lane, 1988), then the tonnage stockpiled will be increased, and the time related differences in value between when the pit optimization assigns the value and when the value is actually realized in the plan increases further.

4.2. COMMON INPUT/OUTPUT ERRORS AND ISSUES

a. Errors in block model regularization and the assumed Smallest Mining Unit (SMU). If a block model is used that features grade estimated blocks at smaller than the SMU size, then unrealistic mining selectivity will be built into the result. If a model is regularized to a larger than SMU size for purposes of processing speed, then the ore/waste tonnage classifications and grades at the SMU size need to be maintained and not smoothed out to the larger regularized block size. Not considering the over-selectivity can easily result in pits with an expectation of double the value of a pit selected from a block model with an appropriately sized SMU.

b. Using Revenue Factor (RF) =1 shells for final pit design. The pit limits that maximize the undiscounted cashflow for a given project will not maximize the NPV of the project.

   As discussed by Whittle (2009) when the time value of money is taken into account, the outer shells of the RF = 1 pit can be shown to reduce value, due to the fact that the cost of waste stripping precedes the margins derived from ore ultimately obtained. The effect of discounted cash flow means the discounted costs outweigh the more heavily discounted revenues. The optimal pit from a Net Present Value (NPV) viewpoint can be between revenue factor 0.65 and 0.95, depending on the deposit's structure and the mining constraints (minimum mining width, maximum vertical advancement per year, and limit on total movement) and processing capacity. This can be seen where the peak of the discounted cash flow of the specified case is at a lower overall tonnage than the peak of the undiscounted total cash curve.

Despite the fact that this aspect is well discussed in the technical literature, the selection of RF=1 shell is still commonly seen in the industry for ore reserves work and project feasibility studies.

Additionally, the curve of discounted cash value versus tonnage tends to be flat at the top. For example, it is common for the last third of the life-of-mine to be quite marginal. Whilst it is worth maintaining the option to operate during this period and in this part of the deposit in case prices, costs, or technology improve, this part of the resource should not be regarded as a core part of and driver of a project (Whittle 2009).

c. Processing Plant Performance Parameters. Aside from price, the other big factor with significant uncertainty and used in the calculation of revenue received for a block of ore is the processing plant recovery. Variations in recovery for grade, mineralogy and hardness can be expected compared to the recovery used in the model. The commonly used constant recovery will almost always be wrong (either because it is optimistically over-estimated, or because there is a fixed tail component not be taken into account).

   Additionally, it should also be noted that project value can often be increased by sacrificing metal recovery to pursue lower cost, higher throughput – as discussed by Wooller (1999).

d. Cut-Off Grade. If blocks with extremely small values (cents per tonne of positive value) are left within the block model used (effectively the use of a marginal cut-off grade value of zero), then a lot of ore will be processed in the project for very little value. Effectively, a significant percentage of the ore is being mined and processed for little more than practice – as discussed in Poniewierski (2016).

   Deswik suggest that in order to avoid this situation that a value cut-off greater than zero be applied. It is suggested that a suitable value would be the minimum desired percentage margin on the processing and sales costs employed. Such blocks would have their revenue value set to zero, so they do not influence the optimal shell selection. Once the final shell has been selected and the ultimate pit designed, the marginal material in that pit can be reconsidered for inclusion in ore reserves and stockpiling if so desired.

   It should also be noted that for NPV maximization, a variable cut-off grade or cut-off value policy should be adopted (as per Lane 1988).

4.3. INPUT UNCERTAINTIES

a. Geological uncertainty. This is one of the biggest sources of error in a pit optimization, as the pit optimization results ultimately depend on the accuracy of the model and the competence of the geologist interpreting all the available geological data. The block model has been created from sparse imperfect data that makes assumptions and estimations on mineralization limits, mineralization grades modeling, fault interpretation and lithology interpretation.

   In the author's experience, many resource models have contained metal errors of at least 10% or more (model over-call) and up to 30% has been seen. Cases of under-call do also occur, and will predominate in the literature as no-one likes to discuss the bad outcomes publically. In the authors experience 70 to 80% of all resource models suffer from overcall to some degree.

   In addition to the grade uncertainty, there is also density uncertainty and in-situ moisture uncertainty.

b. Effect of Inferred resources. Should these be included or not included? If included, these can easily be in error by 50% or more. If not included, the design will change when these are converted to Indicated or Measured status.

c. Geotechnical uncertainty. While a lot of focus can be spent on ensuring the desired overall angles are modeled accurately, in many cases the slopes provided for use may be little more than a geotechnical engineer's guesstimate based on very little rock mass quality data, sparse and imperfect knowledge of faulting, jointing, bedding and hydrology. Even in operating pits, the geotechnical conditions can change quickly from that currently being used.

d. Dilution and loss are nearly always "guesses" – except for sites with a number of years of operating experience and a good reconciliation system that allows for assessment of the dilution and loss (which is not all that common).

e. Economic uncertainty. This is also one of the major sources of "error" with pit optimization. In the analysis of costs and revenues, we have to make assumptions about the macro-economic environment such as commodity prices, exchange rates, interest rates, inflation, fuel and power costs, sustaining capital costs, contractor costs and labor costs. For the commodity price in particular, we can confidently state that the price used will be 100% wrong for the life-of-the mine (it will never be one static value).

f. Costs. Except for operating mines with a good understanding of their detailed cost driver history, there is usually a great deal of uncertainty on the costs being used in the pit optimization. Many parameters used to estimate costs such as equipment selection, annual production rate, plant capacity and requirements, etc. are just estimates. There is usually an imperfect understanding of fixed and variable costs that do not truly reflect the changes in costs as the pits being assessed change in size.

g. In addition, it needs to be noted that fixed costs (or time period costs) need to be applied on the basis of the mine/mill system bottle-neck. As a general rule this is often the SAG mill (with power rather than tonnage being the limit).

5. Summary

Both the Lerchs-Grossmann and pseudoflow algorithms are variations of network flow algorithms that achieve the same result. Pseudoflow is however a computationally more efficient algorithm developed some 35 years after the original Lerchs-Grossman algorithm (1965), and has been available for use for some 15 years, with the first implementation for mining discussed in 2005 (Muir, 2005).

If differences are seen between a Whittle LG result and a Deswik Pseudoflow result, it will be a difference in the set-up used. There are numerous set-up factors and parameters that can cause differences in pit optimization results, and the user should be aware of all of these to avoid falling into common error traps.

It should be noted that the Deswik implementation is not constrained (nor is it aided) by the pre-defined template inputs provided in the Whittle software for cost and revenue calculations (these templates can be restrictive for both very simple set-ups or complex set-ups not catered for).

For Deswik's Pseudoflow implementation, the user is required to calculate the revenues and costs for each block in the block model used, and is required to do their own block regularization within the Deswik.CAD environment.

The user is thus in full control of how the costs and revenues are calculated and assigned, but this does require the user to be fully familiar with their block model, cost structures and revenue parameters (which we believe is a "good thing"). This enables the cost and revenue calculations to be as simple or complex as required by the user (unconstrained by template set-up dialogs).

REFERENCES

Alford C G and Whittle J, 1986. Application of Lerchs– Grossmann pit optimization to the design of open pit mines, In Large Open Pit Mining Conference, AusIMM–IE Aust Newman Combined Group, 1986, 201–207.

Carlson, T R; Erickson, J D, O'Brain D T and Pana, M T, 1966. Computer techniques in mine planning, Mining Engineering, Vol. 18, No. 5, p.p. 53-56.

Chandran, B G and Hochbaum, D S, 2009. A computational study of the pseudoflow and push-relabel algorithms for the maximum flow problem, Operations Research, 57(2): 358-376.

Deutsch, M., González, E. & Williams, M. 2015. Using simulation to quantify uncertainty in ultimate-pit limits and inform infrastructure placement, Mining Engineering, 67 (12), p.49-55.

Dagdelen, K, 2005. Open pit optimization — strategies for improving economics of mining projects through mine planning, in Orebody Modelling and Strategic Mine Planning, Spectrum Series No 14, pp 125-128 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Espinoza, D, Goycoolea, M, Moreno, E and Newman, A N, 2012. MineLib: A Library of Open Pit Mining Problems, Ann. Oper. Res. 206(1), 91-114.

Francois-Bongarcon, D M and Guibal, D, 1984. Parameterization of optimal design of an open pit -beginning of a new phase of research, Transactions of Society of Mining Engineers, AIME, Vol. 274, pp. 1801-1805

Goldberg, A and Tarjan, R E, 1988. A new approach to the maximum flow problem, Journal of the Association for Computing Machinery, 35, 921-940.

Goldfarb, D and Chen, W, 1997. On strongly polynomial dual algorithms for the maximum flow problem, Special Issue of Mathematical Programming B, 78(2):159-168.

Hall, B, 2014. Cut-off Grades and Optimising the Strategic Mine Plan, Spectrum Series 20, 311 p (The Australasian Institute of Mining and Metallurgy: Melbourne).

Hochbaum, D S, 2008. A pseudoflow algorithm: a new algorithm for the maximum-flow problem, Operations Research, 56(4): 992-1009.

Hochbaum, D S and Chen, A, 2000. Performance analysis and best implementations of old and new algorithms for the open-pit mining problem, Operations Research, 48(6): 894-914.

Hochbaum, D S, 2001. A new-old algorithm for minimum-cut and maximum-flow in closure graphs, Networks, 37(4): 171-193.

Jiang, Y, D. 2015, "Impact of Reblocking on Pit Optimization" https://www.linkedin.com/pulse/effect-reblocking-pit-optimization-yaohong-d-jiang

Kentwell, D, 2002. MineMax Planner vs Whittle Four-X - an open pit optimization software evaluation and comparison, MineMax Whitepaper available from https://www.minemax.com/downloads/Minemax-Planner-vs-FourX.pdf

Kim, Y C, 1978. Ultimate pit design methodologies using computer models the state of the art, Mining Engineering, Vol. 30, pp. 1454 1459.

King, V, Rao, S and Tarjan, R, 1992. A faster deterministic maximum flow algorithm, Proceedings of the Third Annual ACM-SIAM Symposium on Discrete Algorithms. Academic Press, Orlando, FL, USA, 157-164.

Lane, K F, 1988. The Economic Definition of Ore: Cut-off Grades in Theory and Practice (Mining Journal Books: London).

Lerchs, H and Grossmann, I F, 1965. Optimum design of open pit mines, The Canadian Mining and Metallurgical Bulletin, Vol. 58, January, pp.47-54.

McCarthy, P L, 1993. "Pit Optimization" internal paper for AMC and whitepaper on the AMC website, available at http://www.amcconsultants.com.au/library

Muir, D C W, 2005. Pseudoflow, new life for Lerchs-Grossmann pit optimization, in Orebody Modelling and Strategic Mine Planning, Spectrum Series No 14, pp 97-104 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Pana, M, 1965. The simulation approach to open-pit design, In J. Dotson and W. Peters, editors, Short Course and Symposium on Computers and Computer Applications in Mining and Exploration, College of Mines, University of Arizona, Tuscon, Arizona. pp. ZZ–1 – ZZ–24.

Picard, J, 1976. Maximal closure of a graph and applications to combinatorial problems, Management Science, Vol. 22, No. 11, pp. 1268–1272.

Poniewierski, J, 2016, Negatively geared ore reserves – a major peril of the break-even cut-off grade, Proc. AusIMM Project Evaluation Conference, Adelaide, 8-9 March 2016. pp236-247.

Whittle, G, 2009. Misguided objectives that destroy value, in Proceedings Orebody Modelling and Strategic Mine Planning, pp 97-101 (The Australasian Institute of Mining and Metallurgy: Melbourne).

Wooller, R, 1999. Cut-off grades beyond the mine – optimising mill throughput, in Proceedings Third Biennial Conference on Strategic Mine Planning. pp 217-230 (Whittle Programming: Melbourne).