Les Cahiers du GERAD ISSN: 0711–2440

Solving a large SIP model for production scheduling at a gold mine with multiple processing streams and uncertain geology

M. de Freitas Silva, A. Lamghari,
R. Dimitrakopoulos

G–2014–61

August 2014

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Solving a large SIP model for production scheduling at a gold mine with multiple processing streams and uncertain geology

Mario de Freitas Silva

Amina Lamghari

Roussos Dimitrakopoulos

GERAD & COSMO – Stochastic Mine Planning Laboratory, Department of Mining and Materials Engineering, McGill University, Montréal (Québec) Canada, H3A 2A7

mario.silva@mail.mcgill.ca
amina.lamghari@mcgill.ca
roussos.dimitrakopoulos@mcgill.ca

August 2014

Les Cahiers du GERAD

G–2014–61

Abstract: The optimization of open-pit mine production scheduling (OPMPS) is an intricate process due to its size and uncertainty of key input parameters. Over the last decade, substantial effort has been made towards the development of new stochastic frameworks that incorporate uncertainty into the decision process. However, due to the intrinsic complexity of the mathematical programming formulation and the large size of mineral deposits, finding an exact solution for OPMPS is likely intractable, motivating the development of new computationally efficient solution approaches.

In this paper, an efficient heuristic solution approach is applied and tested to the stochastic mine production scheduling of a relatively large gold deposit containing about 120 thousand blocks and considering a set of 15 geological scenarios generated stochastically. The case study addresses multiple processing streams and a 'grade' stockpile, which adds flexibility to the specific operation by advancing the processing of highly valuable material. The solution approach first generates an initial feasible solution by sequentially solving the stochastic OPMPS period by period, and then a network flow algorithm is used to sequentially search for improvements. In this network graph, the nodes identify candidate blocks which might have their extraction postponed or advanced, aiming for new schedules with higher value and lower risks. The results show that production schedules with low deviations from production expectations can be generated in a reasonable time for an actual mining environment.

Key Words: Open-pit mine production scheduling, multiple processors, stockpile, geologic uncertainty, network flow algorithms.

1 Introduction

Open-pit mine production scheduling (OPMPS) generates the optimal sequence of extraction of mining units over the life-of-mine (LOM), given a set of physical and technical constraints. Such a decision process needs to be made under conditions of uncertainty, however, conventional approaches for optimizing OPMPS (e.g., Johnson, 1969; Dagdelen and Jonhson, 1986; Gershon, 1987; Whittle, 1988; Tolwinski and Underwood, 1996; Cacceta and Hill, 2003; Hustrulid and Kuchta, 2006) tend to assume that parameter inputs are fully known, ignoring potential risks and opportunities that might arise from the different sources of uncertainty (Ravenscroft, 1992; Dowd, 1994,1997). An example in Dimitrakopoulos et al. (2002) shows that the results in key performance indicators of a conventional mine design are misleading in the presence of geological uncertainty, highlighting the limits of deterministic optimization techniques.

Spatial uncertainty of geological attributes can be modelled through stochastic simulation techniques which are able to provide a series of equally probable scenarios of the orebody (Goovaerts, 1997; David, 1988). The availability of these models leads to the development of stochastic optimization frameworks that are able to integrate uncertainty into the decision process, minimizing downside risks and maximizing potential upsides. During the last decade, a substantial focus has been given for the development of new models and solution approaches for the stochastic version of the OPMPS. For example, Godoy (2003) introduces a stochastic framework where multiple schedules derived from each geological scenario are firstly joined up. Thereafter, a combinatorial optimization problem is solved by an algorithm based on simulated annealing in order to provide a single schedule with a higher NPV (improvements of 28%) and substantially lower deviations from production targets, when compared with the results reported by the conventional schedule. Similar conclusions are drawn in Leite and Dimitrakopoulos (2007) for an application of the framework in a copper deposit. Albor and Dimitrakopoulos (2009) show for a specific case study that the application of this stochastic framework leads to a larger ultimate pit with an NPV 10% larger than the one obtained by constraining the schedule with a conventional pit limit.

Menabde et al. (2007) develop a mathematical formulation to maximize the expected NPV over several scenarios while minimizing deviations from production targets in an average sense. Dimitrakopoulos and Ramazan (2008) bring a stochastic integer programming (SIP) formulation which maximizes the expected net present value (NPV) and incorporates recourse actions to tackle the uncertainty modelled through stochastic simulations, by minimizing possible deviations from production targets over the life-of-mine. Ramazan and Dimitrakopoulos (2013) extend this SIP formulation to introduce a stockpile option, reporting an increase of 10% in the NPV if compared to the economic performance reported by a conventional schedule. In addition, the method provides more realistic schedules that minimize the chance of deviating from production targets, regarding geological uncertainty. These results highlight the ability of stochastic schedules on simultaneously maximising economic returns and driving the mining sequence through zones where the risk of not achieving the target ore production is minimised. Other variants and applications of SIP have also shown significant improvements over the deterministic OPMPS: Leite and Dimitrakopoulos (2014) show through an application to a porphyry copper deposit that, even for a low grade variability deposit, the NPV can be increased by 29%; Dimitrakopoulos and Jewbali (2013) incorporate in the SIP model simulated future data information, outperforming the NPV of the conventional mine design of a gold mine; Benndorf and Dimitrakopoulos (2013) extend the model to account for several elements of an iron-ore operation, showing that the capability of the stochastic approach to controlling risks of deviating from production targets for critical quality-defining elements. Boland et al. (2008) incorporate metal uncertainty via a multistage stochastic programming approach in such a way that, decisions made in later time periods might depend on observations of the properties of the material mined in earlier periods.

The stochastic models proposed by Ramazan and Dimitrakopoulos (2012), Menabde et al. (2007) and Boland et al. (2008), are all solved using a mixed integer programming solver such as CPLEX (ILOG, 1998), which limits their practical application to instances of relative small sizes, typically accounting for less than 20 thousands blocks (Lamghari and Dimitrakopoulos, 2012). As a result, over the past few years, several authors have been seeking the development of new solution approaches, which can efficiently tackle large instances of the stochastic OPMPS. Lamghari and Dimitrakopoulos (2012) introduce a metaheuristic approach based on Tabu search for solving large-scale SIP models within a few minutes up to few hours (while a commercial solver would take days for some instances), with a deviation of less than 4% from optimality for most of their runs. Comparable results are obtained in Lamghari et al. (2013b) who use two variants of a variable neighbourhood decent algorithm and average deviations of less than 3% from optimality for several instances.

The present paper focuses on an application of a heuristic approach introduced by Lamghari et al. (2013a) which incorporates geological uncertainty, multiple processors, stockpiles, and is capable of solving large-size mining schedule problems in a reasonable time. The solution approach can be seen as a very large-scale neighbourhood search method (Ahuja et al., 2002) and it basically involves two stages: (i) the generation of an initial solution and (ii) the application of an improvement algorithm based on network flow. In the following sections, the SIP formulation and the solution approach are revisited, followed by the application at a gold mine employing two processing streams and one 'grade' stockpile. Discussions and conclusions follow.

2 Stochastic integer formulation revisited

The stochastic integer formulation proposed by Lamghari et al. (2013a) is briefly outlined in this section. The following notation is used:

- N is the total number of blocks and i is the block index with $i = 1,\ldots,$ N.
- T is the total number of periods and t is the period index with $t = 1,\ldots,$ T.
- S is the total number of scenarios used to model the geological uncertainty and s is the scenario index with $s = 1,\ldots,$ S.
- P is the total number of processing streams and p is the processing index with $p = 1,\ldots,$ P (e.g., a mill and a leaching facility).
- *Pred(i)* is the set of predecessors for a given block i, which means that all blocks in this set must be exploited before i in order to satisfy the slope constraints.
- dr is the economic discount rate over the time basis being considered.
- d_{ips} is a parameter indicating the most profitable destination for a block i under scenario s. Therefore, comparing the block grade g_{is} to the cut-off policy adopted, d_{ips} is equal to one for its most profitable destination stream for scenario s and it is equal to zero for all other destinations.
- w_i is the total tonnage of a given block i.
- $E[BEV_i]$ is the expected block economic value (BEV) of a given block i. This value is calculated for each geological scenario, considering the best destination of the block accordingly to the cut-off policy of the project, which is given by the d_{ips}.
- SC_P^t and RC_P^t are both undiscounted costs related to stockpile activities for a given process p during period t. The former cost stands for sending material to the stockpile and the latter for reclaiming material from the stockpile.
- $\tilde{r}_{ps}^{t}$ is the discounted revenue returned, if a tonne of ore under a given scenario s is reclaimed from the stockpile and sent to process p during production period t.
- W^t and Θ_p^t are the maximum mining and processing capacities (for each processing option p) respectively, for a given period t.
- I_p is the initial amount of material in the stockpile of processor p.
- Binary variables (x_i^t) for each block i and period t. It is considered that x_i^t is equal to one if the block i is already mined by period t, otherwise it assumes the value of zero. It means that, if a block I is mined in period t^*, $x_i^t = 0$ for all $t = 1 \ldots t^* - 1$ and $x_i^t = 1$ for all $t = t^* \ldots$ *last period*. In consequence, $x_i^t - x_i^{t-1}$ only assumes the value of one for the period t when the block is mined. To simplify the notation, a set of N dummy decision variables $x_i^0 (i = 1 \ldots N)$ are introduced, all having a fixed value equal to zero.
- Linear variables (y_{ps}) related to processing streams. In the model proposed, y_{ps}^{t+} and y_{ps}^{t-} represent the surplus and shortage of material in a given period t, for a process destination p, regarding a specific scenario s. These variables are used to model the stockpile streams related to each process p. In case of surplus under a given scenario, y_{ps}^{t+} is the amount of material that must be stockpiled in order not

to violate the processing capacity available. In case of shortage, y_{ps}^{t-} accounts for the amount reclaimed from the stockpile to fulfill the processing capacity. Finally, the variables y_{ps}^{t} denote the amount of ore in the stockpile at the end of period t.

The mathematical model aims to maximize the discounted cash flow (Eq. 1) given some physical and technical constraints related to the mining operation (Eqs. 2 to 10) as summarized below:

$$\max \sum_{t=1}^{T} \frac{1}{(1+dr)^{t}} \left\{ \sum_{i=1}^{N} E[BEV_i](x_i^t - x_i^{t-1}) + \frac{1}{S} \sum_{s=1}^{S} \left[-\sum_{p=1}^{P} (\tilde{r}_{ps}^{\,t} + SC_P^t) y_{ps}^{t+} + \sum_{p=1}^{P} (\tilde{r}_{ps}^{\,t} - RC_P^t) y_{ps}^{t-} \right] \right\} \quad (1)$$

Subject to:

$$x_i^{t-1} \leq x_i^t \qquad \forall\, i, t \quad (2)$$

$$x_i^t \leq x_j^t \qquad \forall\, i, j \in Pred(i), t \quad (3)$$

$$\sum_{i=1}^{N} w_i (x_i^t - x_i^{t-1}) \leq W^t \qquad \forall\, t \quad (4)$$

$$\sum_{i=1}^{N} d_{ips} w_i (x_i^t - x_i^{t-1}) - y_{ps}^{t+} + y_{ps}^{t-} \leq \Theta_p^t \qquad \forall\, t, p, s \quad (5)$$

$$y_{ps}^{t-1} + y_{ps}^{t+} - y_{ps}^{t-} = y_{ps}^{t} \qquad \forall\, t, p, s \quad (6)$$

$$x_i^t = 0 \; or \; 1 \qquad \forall\, i, t \quad (7)$$

$$x_i^0 = 0 \qquad \forall\, i \quad (8)$$

$$y_{ps}^{t+}, y_{ps}^{t-}, y_{ps}^{t} \geq 0 \qquad \forall\, t, p, s \quad (9)$$

$$y_{ps}^{0} = I_p \qquad \forall\, p, s \quad (10)$$

As per Eq. 1, the objective function can be separated in two major terms: the first one refers to the mining decisions, without having access to full information about the material that is underground (scenario independent); the remaining is associated to scenario dependent variables (stockpile actions), because once a block is mined, the operation can take the most suitable decision about where to send a given mined block, leading to different stockpile actions under each scenario. The first part of the stockpiling term refers to the total approximated undiscounted cost related to send exploited material from the mine to the stockpile of processor p in period t under scenario s; and the second part refers to the total approximated undiscounted net revenue after reclaiming material from the stockpile of processor p in period t under scenario s. As one may note, the option of using the stockpile incurs additional costs in the objective function. Thus, in an optimal solution the use of the stockpile is minimized, which means that the risks of overproduction regarding all geological simulations are also minimized.

The addition of a stockpile to the mathematical model adds the number of quadratic terms to the formulation of the OPMPS. The calculation of the parameters $\tilde{r}_{ps}^{\,t}$ in the objective function (1) depends on the knowledge of the average grades ($\tilde{G}_{ps}^{t}$) of the material that is sent/reclaimed in/from the stockpile. For a realistic approximation of the average grade of the stockpile in a given period, one should track the average grade of the material being extracted and sent to the stockpile up to this period. However, the notion of which blocks ought to be extracted are related to decision variables of the model, resulting in a cross relationship with the revenue generated by processing the material from the stockpile, and giving rise to a non-linear term in the objective function. To maintain the linearity of the model, Lamghari et al. (2013) use a single initial approximation for the average grade of the stockpile associated to each processor, which is fixed for all the scenarios and over all the periods. In this paper, the set of average grades $\tilde{G}_{ps}^{t}$ is iteratively approximated and may vary from period to period and scenario to scenario. This iterative approximation is performed in the following way: first, the schedule is solved using a approximated average grade, which might come from the average grade of all blocks in the deposit which are candidates to go to the destination related to the stockpile, or the average grade of materials within the cut-off between its processor and the low-grade processor. After solving the OPMPS with this approximation, the optimizer outputs the amount of material going in and out of the stockpile (respectively given by the linear variables y^{t+}_{ps} and y^{t-}_{ps}), but it does not track which blocks specifically are being stockpiled. Since no blending constraints are considered, it is assumed that in each period, from the set of blocks scheduled to be sent to a given destination, the ones that go to the stockpile are the ones with the lowest grades, because in an optimal solution, due to the time value of money, the low-value material is stockpiled in order to leave room for the processing of high-value material. By doing such an analysis, it is possible to calculate the "expected" average grade of the stockpile for that given schedule. These average grades by period are then fed as input to the optimizer to generate a new schedule. The same process of approximating the grades and rerunning the solver keeps looping until the difference between the input grade and the "expected" one is less than a threshold ε (e.g., 10%). A similar approach is used by Sarker and Gunn (1997) to solve nonlinear problems, where the authors iteratively solve multiple linear programming problems approximating the quality of the blended material at different locations in terms of sulfur, ash and BTU content. The same authors show that, not only it is a simple and fast way of dealing with nonlinear problems, it is able to provide solutions near optimality after few iterations. A comparison between the approximation for the average grade of the stockpile shown herein and the one in Lamghari et al. (2013) is shown in the Appendix.

Constraints in Eq. 2 are the *reserve constraints*, which guarantee that each block is mined at most once. Constraints in Eq. 3 are the *slope constraints*, which entails that to access a given block, a set of predecessors must be mined before, assuring the slope angles are predefined. Constraints in Eq. 4 are the *mining constraints* which enforce that the total amount of material mined in a given period t cannot be higher than the mining capacity available for that period. Constraints in Eq. 5 are the *processing constraints*, which imposes an upper bound for the total material sent for a given process, in period t and under scenario s. Constraints in Eq. 6 are the *stockpiling constraints* which balance the mass flows of each stockpile. At the end of each period t, the stockpile of a processor p under scenario s must contain the amount of material initially available at the end of period $t-1$, plus the amount of material stockpiled minus what was reclaimed during period t.

It is noteworthy that, although the model does not consider explicitly a lower bound capacity for the processing streams in order to better control the ore feeding, the optimizer always tries to use all the capacity available, mostly in earlier periods as an attempt to increase the NPV of the project. From this constraint one may also note that, in an optimal solution, either stockpiling or reclaiming is active, since both incur costs in the objective function. Thus, the surplus variable (y^{t+}_{ps}) will assume positive values when it is worthy to mine and send to a given processor more material than it can handle during that period, which is an attempt to reach as fast as possible the high-grade zones of the deposit, stockpiling low-grade material that had to be mined previously. On the other hand, the shortage variable (y^{t-}_{ps}) will only assume a positive value either if it is better to reclaim material from the stockpile than fulfilling the processor capacity with material from the ground, or when the material mined is not enough to fulfill the processor capacity. Moreover, inasmuch as the stockpile actions incur costs to the objective function (both associated to rehandling and opportunity cost of postponing the process valuable material in the stockpile), in an optimal solution the use of the stockpile is minimized, which means that the risks of overproduction regarding all geological simulations are also minimized. These features are expected to drive the optimizer to maximize value and minimize geological risk throughout the life-of-mine.

3 A review of the solution approach

For solving the OPMPS model introduced in the previous section, a multistage heuristic algorithm described in Lamghari et al. (2013a) is used. It comprises two major steps: generation of an initial feasible solution and then its improvement by using a network flow based algorithm which efficiently searches for improving solutions over a large neighbourhood.

3.1 Generating an initial feasible solution

Two heuristics methodologies are used to test different initial solutions. Both of them are based on the "divide and conquer" principle, by solving the model period by period, and thus, each period composes a reduced sub-problem. As soon as an earlier period is solved, the mining blocks scheduled to this period are taken out from the model to reduce the problem's size. The later periods are sequentially solved in a similar way. After this sequential process, the solutions found are merged, providing an initial feasible solution.

The differences existing between the two heuristic methods used are basically in the way each one solves the sub-problems. In the first method, the solutions are given by an exact mathematical programming method implemented in CPLEX. The second method is a greedy heuristic procedure (GH) which at each iteration tries to include in the set of mining blocks scheduled for a given production period t, a set of blocks represented by a base block (i) and its predecessors (*Pred(i)*) not mined yet, in such a way as to maximize the objective function of the sub-problem model, respecting the mining capacity constraint and at the same time postponing the extraction of waste and advancing the extraction of ore, thus, deferring costs and advancing profits to earlier periods. This greedy heuristic incorporates a *look ahead* feature, since it looks after blocks with all their unmined predecessors instead of treating blocks separately one by one. In both methods for generating initial feasible solutions, blocks that are not included in the sets of mined blocks in each period until the last one (T) are left behind. To represent these unmined blocks, they are included in a set corresponding to a fictitious period (T+1).

3.2 Improving the initial solution with a network flow algorithm

It is well known that sequentially solving the mine production schedule does not lead to an optimal solution of the long-term production schedule (Gershon, 1983). Therefore, in a second stage, the goal is to improve the initial solution generated by any of the two heuristic approaches explained above, providing a new schedule with a higher NPV. To achieve this, the improvement algorithm proposed basically tries to postpone the extraction of blocks responsible to decrease the objective function (1) and advance those which improve it.

The algorithm is based on a network-flow structure, where each problem is defined on a graph G = (V, E) (V is the set of nodes and E is the set of arcs). Different graphs are built according to the problem being solved: delaying (*backward pass*) or advancing (*forward pass*) extraction of blocks. Only the construction of the first case is shown henceforth, since the formulation of the *forward structure* is straightforward. Thus, for the *backward structure*, the set of nodes represent blocks which matches the following characteristics:

- the total expected economic value of a given block and all its successors scheduled for the same period t is negative, since NPV increases as the costs are deferred, and;
- the total tonnage of this same group of blocks, summed to the total tonnage already scheduled for the next period $(t+1)$, minus the total tonnage of a candidate group of blocks scheduled to period $(t+1)$ which can be postponed to $(t+2)$, must not exceed the mining capacity W^{t+1}. This condition ensures that the mining capacity is not violated when blocks are moved from one period to another.

Each node of the graph is associated to a block and its predecessors scheduled for the same period, respecting the conditions stated above. To complete the network, an additional node is added to the fictitious period T+1 which represents the set of blocks that will not be extracted; for each period, one extra node is added for fictitious blocks with neither weight nor costs, representing a path through where no modifications are done to the current schedule. In addition, two extra nodes must be added to the network referring to its source and sink. In this formulated graph, the set of arcs E involves all possible connection between two nodes currently scheduled in consecutive periods t and $t+1$. In addition, some arcs are added connecting the source to the nodes belonging to the first period and one more arc is connected from the fictitious node in T+1 to the sink. As a result, each path from the source to the sink, passing through nodes in consecutive periods, represents a new solution to the stochastic mine production schedule, where a given mining block and its successors represented in a node has their extraction delayed to the next period and so on. Blocks at the head of the arc are moved to the following period and the blocks represented by the node at the tail of the arc are mined in their place. Figure 1 shows a simplified illustration of graph G.

Thus, the goal is to find a single feasible path which improves the value of the objective function as in Eq. (1). If no such path is found, the solution given by the algorithm is the path which includes the set of fictitious nodes introduced before, and no block would be moved from one period to the other and the value of the objective function remains the same. To identify the feasible path which increases the value of

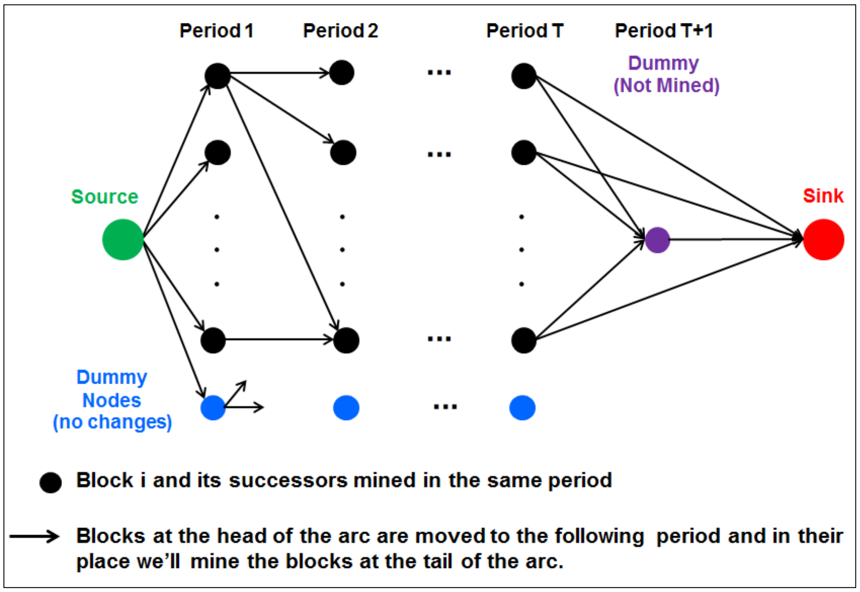

Figure 1: Ilustration of the graph built for the network-flow improvement algorithm (*Backward Case*).

the objective function the most, each arc is weighted accordingly to the feasibility of the delaying movement and the gain it brings to the objective function. After weighting each arc the model becomes a *longest path problem*, which consists in finding the simple path of maximum length, where the length in this case is represented by the sum of the weighted arcs. The weights (p_{ij}) of each arc $(i, j) \in$ E are defined as follows:

- Arcs connecting nodes associated to real blocks in each period $t = 1, \ldots,$ T+1, represent movements of blocks associated in node i from their period t to $t+1$ and movements of blocks associated to node j from their period $t+1$ to $t+2$. So, if this move results in a violation of mining capacity at $(t+1)$, the weight associated to this arc (i, j) is a large negative number, which flags its infeasibility. Otherwise, a weight p_{ij} is associated to the arc, indicating the delta change in the two first parts of the objective function (Eq. 1). The last part, which accounts for the value of reclaimed material from the stockpile, is not considered because it cannot be evaluated in the network structure, since one needs to know what is available in the stockpile before knowing how much can be reclaimed;
- Arcs connecting the source to nodes in the first period are always feasible, (the upper bound of the mining capacity is always satisfied in period 1, because only delaying action can be taken). Thus, the p_{ij} weights are simply associated to the modifications in the two first terms of the objective function;
- A $\mathrm{p}_{ij} = 0$ is associated to the arcs between the fictitious nodes created for each period. A null weight is also assigned to the arc from the dummy node in T+1 to the sink.

As mentioned earlier, once the graph is built, solutions are generated by solving the *longest path problem*, associating each arc to a binary variable and sending a unitary flow from the source to the sink, which guarantees that the solution provided is always a simple path. It is interesting to note that the constraint matrix (nodes-arcs incidence matrix) of this integer programming problem is unimodular. This property indicates that the integrality constraints can be dropped and only restricts $\mathrm{z}_{ij} \in [0,1]\ \forall\, (i, j) \in E$. Subsequently, the problem can be efficiently solved using linear programming or network-flow techniques.

In summary, the algorithm works in an interactive way such as the following: first, it performs a *backward pass* (initial mode), trying to delay the extraction of blocks. If the solution changes, a new network is built for the new current schedule and another *backward pass* is carried out. Otherwise, if the *longest path* found identifies the set of fictitious nodes, meaning that no improvement can be made, the problem is switched to a *forward pass*, and the algorithm looks for blocks to advance their extraction. In the same way as in the first, several passes are performed until no improvement is achieved, and then the problem switches its mode again. The algorithm stops when it executes two consecutive modes, that is, *backward* and *forward passes* (and vice-versa) without any improvement in the value of the objective function.

4 Case study at a gold mine

To demonstrate the application related aspects of the method previously described, a case study at gold mine is presented here. The deposit being mined consists of an intensely mineralized shear system localized in mainly steeply dipping, NNW to NW striking lodes. Gold lodes can be up to 1,800 m (5,900 ft) long, have vertical extents of 1,200 m (3,900 ft) and be up to 10 m (33 ft) wide. The mine accounts with two processing streams, a mill and a leaching facility, with the first having an associated stockpile. Fixed stockpiling/reclaiming costs are used throughout the LOM and no material is in stock for the first production period.

A set of 15 stochastic simulations, discretized in about 120 thousand blocks of $20{\times}20{\times}20\text{m}^3$ and generated by direct block simulation (Godoy, 2003), are used to model the spatial uncertainty of grades thought the deposit. This number of scenarios is used because past works, such as in Albor and Dimitrakopoulos (2009) and Leite (2008), indicates that after about such number of representations of an orebody, the stochastic schedules tend to converge to a stable final schedule and to provide stable forecasts of production performance. Such results are not surprising because, despite the spatial uncertainty modeled over blocks with few cubic meters, a production schedule of a mine represents a grouping of several hundreds to thousands of these blocks in one mining period under different constraints. Thus, with this significant increase of support (from blocks to production in mining periods), the stochastic schedules tend to be less sensitive to additional scenarios after a relatively small number of scenarios.

The general parameters for the stochastic mine production schedules are summarized in Table 1.

Table 1: Technical and economic parameters for OPMPS

Mining Cost	\$1.80/t	**Mining Capacity**	90 Mta
Metal Price	\$730/oz	**Selling Price**	\$5.0/oz
Discount Rate	8.0%	**Slope Angle**	45°
Mill - High Grade			
Recovery	90%	**Proc. Cost**	\$9.50/t
Stockpiling Cost	\$0.50/t	**Reclaiming Cost**	\$0.85/t
Proc. Capacity	22.0 Mta	**Stockpile Capacity**	20 Mt
Leaching - Low Grade			
Recovery	50%	**Proc. Cost**	\$5.00/t
Proc. Capacity	1.0 Mta		

The case study is split in two subsections in order to show the differences obtained when using *branch-and-cut* (Wolsey, 1998), an exact mathematical programming method implemented in CPLEX (ILOG, 2008) or a greedy heuristic to generate the initial solution. The computations are performed in a Intel Xeon 5650 (2.66GHz) with 24GB RAM. In both case studies, CPLEX is used to solve the *longest path problem* over the network during the improvement stage of the algorithm.

4.1 Stochastic schedules

Two different schedules are generated, each respectively using CPLEX and the greedy heuristic (GH) to generate the initial feasible solutions. The risk profiles for the ore throughput for the mill and the material stockpiled by the end of each period are respectively shown in Figure 2 and Figure 3. In these graphs, the blue and red solid lines refer to the expected ore input to the mill in the schedules generated by respectively using CPLEX and GH as initial solutions. The dashed lines represent the percentiles P10 and P90 for the ore throughput over the different geological scenarios .

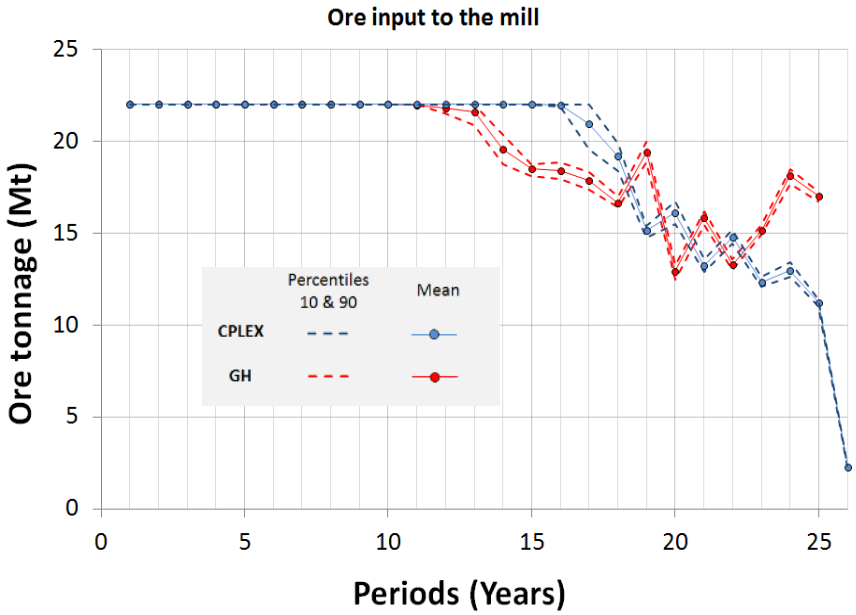

Figure 2: Expected ore tonnage throughput for the mill and related risk profiles, using CPLEX (blue) and GH (red) to generate initial solutions.

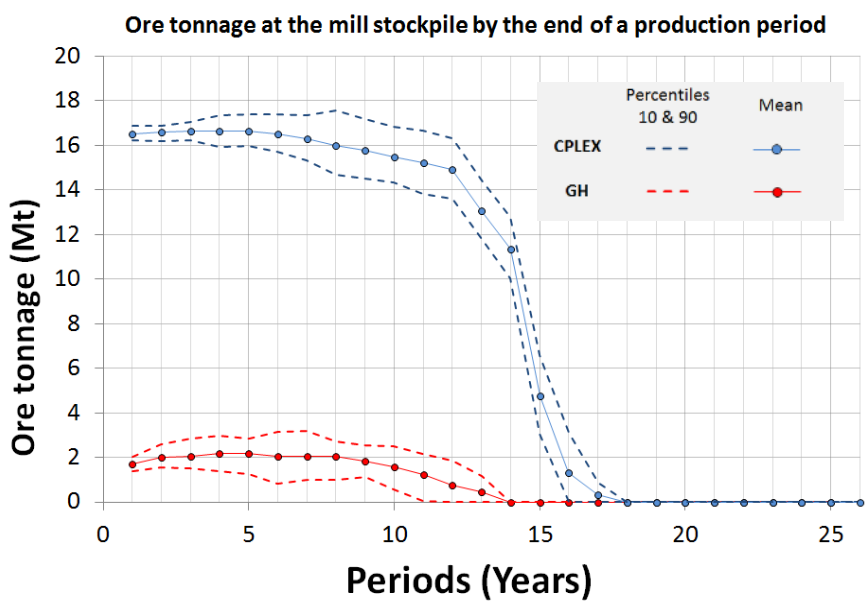

Figure 3: Expected ore tonnage at the mill stockpile and related risk profiles, using CPLEX (blue) and GH (red) to generate initial solutions.

The schedule using CPLEX as initial solution considers an additional year to the LOM, shown in blue (Figure 2) and an ultimate pit 1.1% than the schedule using GH as initial solution. As seen in Figure 2, the range of variability about the expected value of throughput for the mill is quite low, which suggests that this process is likely to operate with low uncertainty for the expected throughput. For the schedule obtained used the initial solution from CPLEX, the mill will potentially work at full capacity (22Mt) during the first sixteen years, while for the OPMPS using the GH as initial solution, this period is shortened to eleven years. During these time spans, the mill potentially works with almost no risks of over/under production. This occurs because during those periods, the tonnage uncertainty is somehow “shifted” to the stockpile, since for each scenario, the overproduction is sent to the stockpile and in case of shortages, material can be reclaimed from the stockpile. During the years for which the mill works below its capacity (Figure 2), the mine operates at full mine capacity (90Mt) and not enough material is available in the stockpiles under all geological scenarios (Figure 3). These factors lead the optimizer to work below the mill's maximum capacity, since the mining rate entails in a bottleneck for the operation and no penalties are incurred for underproduction in the SIP formulation presented in a previous section. A way of dealing with this would be to explicitly incorporate penalties for idle capacity (shortage in production) in the formulation, in such a way that they do not compete with the reclaiming variables, or allow a flexibility to the mine to increase its capacity during later periods, through the acquisition of mining equipment.

Figure 3 shows that, for both schedules generated, in the first period is when most of the material is sent to the stockpile, which allows the mill to advance the metal production, by working with a high grade material as shown in Fig 4. These results show, as expected, the flexibility added to the project by the use of a stockpile: (i) it allows the operation to reach high grade material earlier during the LOM and (ii) 'buffers' the risks of oversupply of ore and/or having idle processing capacity, with respect to geological uncertainty.

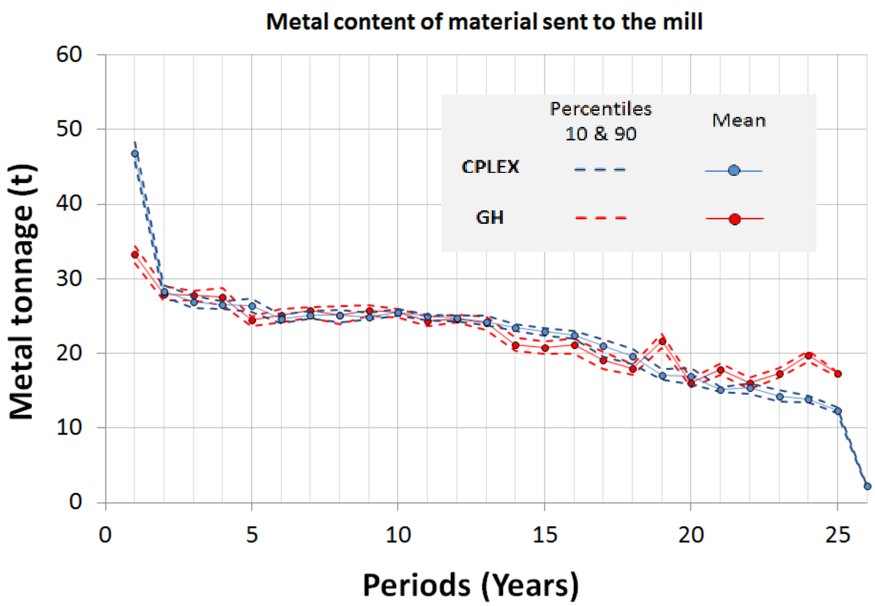

Figure 4: Expected metal input to the mill and related risk profiles, using CPLEX (blue) and GH (red) to generate initial solutions.

Regarding the differences between the two schedules generated, Figure 2 shows that the OPMPS using CPLEX to generate the initial solution is able to advance the production of ore (*see* years 13 to 18) and to reach high grade areas during earlier periods if compared to the solution by using the GH as initial solution. The metal content for the ore input to this processor during the first year is about 40% larger in the first production schedule than in the second. Figure 3 shows that these differences are mostly related to the fact that, in the first solution approach the greater use of the stockpile provides a larger flexibility to the operation.

In contrast to the behaviour seen for the mill, Figure 5 shows that the leaching process will potentially work under its nominal capacity of 1Mt and with a much more variable throughput. Such a result is expected because the SIP formulation used in this paper, controls the geological risks exclusively through the use of a stockpile associated to the process, which is not the case for leaching. Figure 6 shows the risk profile for the metal production of this same processing destination.

Regarding the economic performance of the project, the risk profiles of the cumulative NPV are shown in Figure 7. These curves show a very low uncertainty about the expected NPVs for the project (less than 3% of upper/lower deviations regarding the P10 and P90). In addition, Figure 6 shows that, the OPMPS using

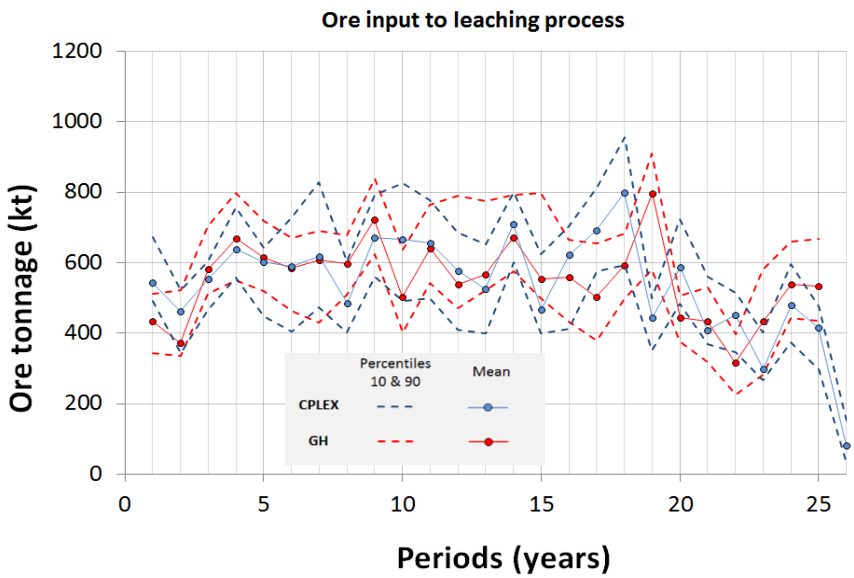

Figure 5: Expected ore tonnage throughput for the leaching and related risk profiles, using CPLEX (blue) and GH (red) to generate initial solutions.

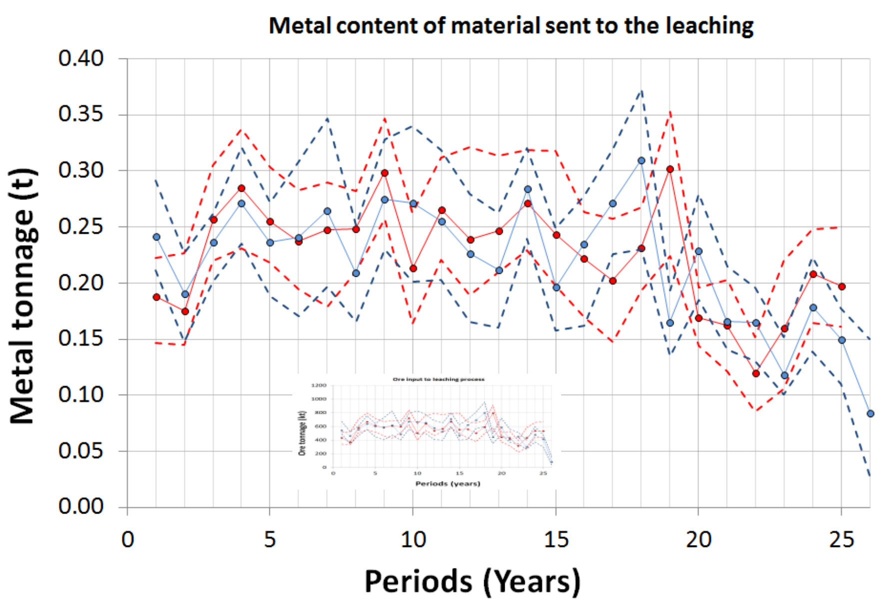

Figure 6: Expected metal input to the leaching and related risk profiles, using CPLEX (blue) and GH (red) to generate initial solutions.

the GH as initial solution has an overall NPV 7.9% (M$215) lower than the one obtained by using CPLEX as initial solution. In this specific case study, this difference is mostly related to the ability of the latter solution to produce a larger amount of metal during the first period. In this year, its NPV is 56% (M$245) higher than the one achieved by the mine production scheduling obtained by using GH as initial solution.

While CPLEX takes hours to generate an initial solution, the GH takes only seconds. In addition, the final OPMPS using CPLEX as initial solution took a total time of 32 hours against the 38 hours required for the generation of the final solution by the approach using the GH as initial solution. This excessive time

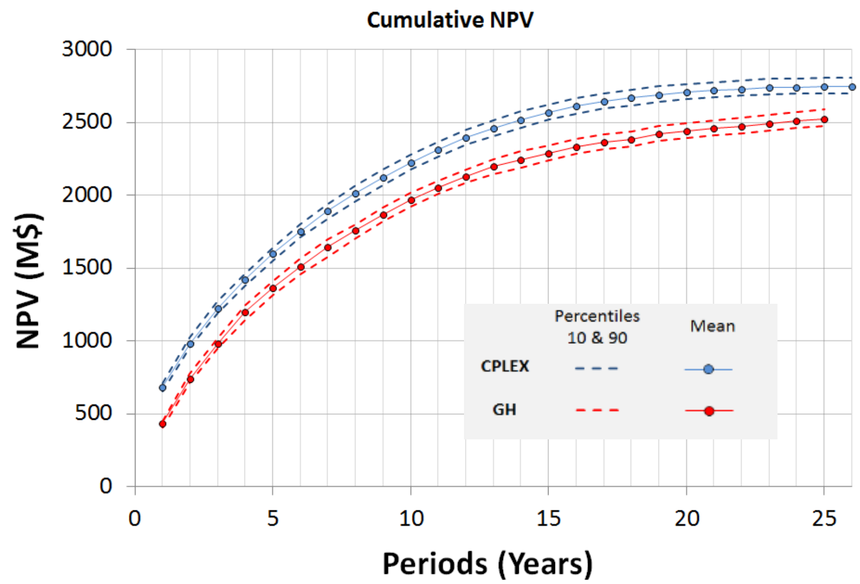

Figure 7: Expected cumulative NPV and related risk profiles, using CPLEX (blue) and GH (red) to generate initial solutions.

reported by this last approach is related to the size of the neighbourhood found in each iteration when it tries to make a *backward move*. In many of these iterations, the graph built has contained more than 2.7 thousands of nodes and 22 millions of arcs. This represents a very large linear programming problem, requiring more than 30 minutes to be solved. Thus, as one may observe, for this case study, besides providing a higher NPV, the final schedule generated using CPLEX's initial solution also demands a smaller computational time than the approach using GH to generate an initial solution.

Figure 8 brings South-North cross sections of the schedules, illustrating the differences in their physical sequence of extraction. Using CPLEX as initial solution produces a less smooth sequencing pattern than the one provided by employing GH as initial solution. For this case study, this latter approach tends to maintain the "clustered" structure intrinsic from its formulation

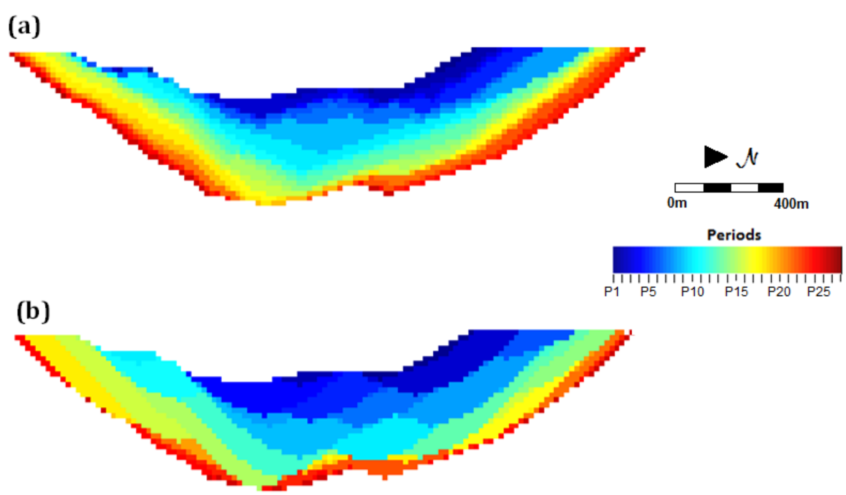

Figure 8: South-North vertical cross-section of the physical sequences of extraction for the schedules using different initial solutions (a) CPLEX and (b) GH.

5 Conclusions

The present study highlights the practical aspects and performance of a neighborhood search method based on a network flow algorithm, developed to solve a stochastic version of the open-pit mine production scheduling. A case study was performed at a relatively large gold mine comprising about 120 thousand blocks, two processing streams and a stockpile. This consists of a very large mathematical programming model, with about 3 million integer variables. Two different ways of generating initial feasible solutions to be input to the network flow algorithm were tested. The first uses CPLEX and the second a greedy heuristic to sequentially solve the mine production schedule period-by-period. For the specific case study, although the greedy heuristic was able to find the initial solution in a few seconds and the exact method demanded hours for the same task, the improvement stage was much longer when using the greedy heuristic solution. This latter approach took 38 hours to generate a final schedule, against 32 hours required by the optimizer when the CPLEX initial solution is used. This behavior is different to the common trend observed in previous tests (Lamghari et al, 2013). Note that, when CPLEX was used to produce the linear relaxation of the stochastic integer programming model of this case study, it could not provide an optimal solution after two weeks, highlighting the advantages of looking for computationally efficient solutions, such as the one used in this paper.

In this case study, the production schedules generated showed that by using the initial solution from CPLEX, a better final solution can be achieved in terms of NPV (7.9% higher than starting from the initial solution generated by the GH). All results have shown that the stochastic mine production schedules have controlled deviations in ore production for the processor with a stockpile associated to it, since the SIP formulation used in this paper, considers that the recourse actions to control the geological risk are incorporated in the stockpiling actions. The overproduction under any scenario is sent to the stockpile and shortages are only controlled if there is material available in the stockpile. These actions imply costs associated to rehandling of material and the opportunity cost of leaving valuable material in the stockpile, and therefore, penalizing deviations related to uncertainty. The shortcoming is that, if in a giving production period, no material is available at the stockpile, shortages are not explicitly penalized.

These observations highlight that the heuristic method tested in this paper is able to tackle large SIP formulations for realistic mine environments, producing mine production schedules with low deviations about expected production rates.

Appendix A – Testing different approximations for the stockpile's average grades

To illustrate the differences of using an iterative approach to approximate the average grade of the stockpile, a third schedule is generated and shown in this Appendix. This schedule is generated using a single initial approximation for the stockpile average grades, which is the average grade of all resources that will be potentially sent to the mill accordingly to the marginal cut-off grade of the operation. For this schedule, CPLEX is also used to generate an initial feasible solution. The results have shown that, the biggest differences are related to the management of the stockpile (Figure 9), given that for iterative approach, the OPMPS considers a more extensive use of the stockpile than the schedule obtained from using a single approximation. In this specific case study, this lead to a NPV of only 1.3% (M$35.66) larger for the first schedule than the one obtained by the second approach tested herein.

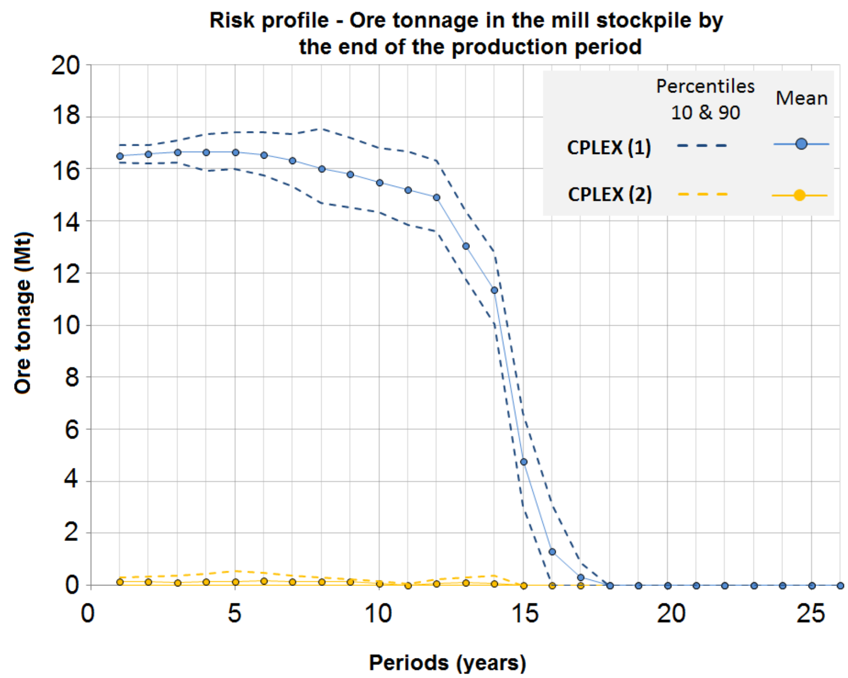

Figure 9: Risk profiles for tonnage at mill stockpile - Comparing Schedules generated by iterative approximation (CPLEX(1)) or by a single approximation (CPLEX(2)) of the stockpile grade.

References

Ahuja, R.K., Ergun, Ö., Orlin, J.B., and Punnen, A.P. 2002. A survey of very large-scale neighborhood search techniques. Discrete Applied Mathematics, 123(1), 75–102.

Albor Consuegra, F.R. and Dimitrakopoulos, R. 2009. Stochastic mine design optimisation based on simulated annealing: pit limits, production schedules, multiple orebody scenarios and sensitivity analysis, Min. Technol. (Trans. Inst. Min. Metall. A), 118(2), A80–A91.

Benndorf, J., and Dimitrakopoulos, R. 2013. Stochastic long-term production scheduling of iron ore deposits: Integrating joint multi-element geological uncertainty, J. Mining Science, 49(1), 68–81.

Boland, N., Dumitrescu, I., Froyland, G., and Gleixner, A.M. 2008. LP-based disaggregation approaches to solving the open pit mining production scheduling problem with block processing selectivity, Computers & Operations Research, 36(4), 1064–1089.

Caccetta, L., and Hill, S.P. 2003. An application of Branch and Cut to open pit mine scheduling, Journal of Global Optimization, 27(2–3), 349–365.

Dagdelen, K., and Johnson, T.B. 1986. Optimum open pit mine production scheduling by lagrangian parameterization, in Proc. 19th APCOM Symposium of the Society of Mining Engineers, 127–142, New York, AIME.

David, M. 1988. Handbook of applied advanced geostatistical ore reserve estimation, Amsterdam, Elsevier.

Dimitrakopoulos, R., Farrelly, C.T. and Godoy, M. 2002. Movingforward from traditional optimization: grade uncertainty and risk effects in open pit design, Min. Technol. (Trans. Inst. Min. Metall. A), 111(1), 82–88.

Dimitrakopoulos, R. and Ramazan, S. 2008. Stochastic integer programming for optimizing long-term production schedules of open pit mines: methods, application and value of stochastic solutions, Min. Technol. (Trans. Inst. Min. Metall. A), 117(4), 155–160.

Dimitrakopoulos, R., and Jewbali, A. 2013. Joint stochastic optimization of short- and long- term mine production planning: Method and application in a large operating gold mine. Min. Technol. (Trans. Inst. Min. Metall. A), 122(2), 110–123.

Dowd, P. 1994. Risk assessment in reserve estimation and open pit planning, Min. Technol. (Trans. Inst. Min. Metall. A), 103, 148–154.

Dowd, P. 1997. Risk in minerals projects: analysis, perception and management, Min. Technol. (Trans. Inst. Min. Metall. A), 106, 9–18.

Gershon, M. 1983. Optimal mine production scheduling: evaluation of large scale mathematical programming approaches, Geotechnical and Geological Engineering, 1(4), 315–329.

Gershon, M. 1987. Heuristic approaches for mine planning and production scheduling, International Journal of Mining and Geological Engineering, 5(1), 1–13.

Godoy, M. 2003. The effective management of geological risk in long-term production scheduling of open pit mines, PhD thesis (unpublished), University of Queensland, Brisbane, Australia.

Goodfellow, R., and Dimitrakopoulos, R. 2013. Algorithmic integration of geological uncertainty in pushback designs for complex multi-process open pit mines, Min. Technol. (Trans. Inst. Min. Metall. A), 122(2), 67–77.

Goovaerts, P. 1997. Geostatistics for natural resources evaluation, New York, Oxford University Press.

Hustrulid, W. A., and Kuchta, M. 2006. Open pit mine planning & design, 2nd ed., London, Taylor and Francis.

ILOG, 2008. ILOG CPLEX 11.2 User's Manual, Armonk, IBM.

Johnson, T. 1969. Optimum production scheduling, in Proc. of the 8th international symposium on computers and operations research, 539–562, Salt Lake City.

Lamghari, A., and Dimitrakopoulos, R. 2012. A diversified tabu search approach for the open-pit mine production scheduling problem with metal uncertainty, European Journal of Operational Research, 222, 642–652.

Lamghari, A., and Dimitrakopoulos, R. 2013a. A network-flow based algorithm for scheduling production in multi-processor open pit mines accounting for metal uncertainty. [online]. Available at: https://www.gerad.ca/en/papers/G-2013-63 [Accessed 1 April 2014].

Lamghari, A., Dimitrakopoulos, R., and Ferland, J. 2013b. A variable neighborhood descent algorithm for the open-pit mine production scheduling problem with metal uncertainty, J. Operational Research Society, doi:10.1057/jors.2013.81.

Leite, A. 2008. Application of stochastic mine design and optimization methods, M. Eng. Thesis, McGill University, Montreal, Canada.

Leite, A., and Dimitrakopoulos, R. 2007. Stochastic optimization model for open pit mine planning: application and risk analysis at a copper deposit. Min. Technol. (Trans. Inst. Min. Metall. A), 116(3), 109–118.

Leite, A., and Dimitrakopoulos, R. 2014. Stochastic optimization of mine production scheduling with uncertain ore/metal/waste supply: application at a copper deposit, J. Min. Sci. Technol., In Press.

Menabde, M., Froyland, G., Stone, P., and Yeates, G. 2007. Mining schedule optimisation for conditionally simulated orebodies, in Proc. Orebody Modelling and Strategic Mine Planning Spectrum Series, 14, (ed. R. Dimitrakopoulos), 379–384, Burwood, The Australasian Institute of Mining and Metallurgy.

Montiel, L., and Dimitrakopoulos, R. 2013. An extended stochastic optimization method for multi-Process mining complexes. [online] Available at: https://www.gerad.ca/en/papers/G-2013-56 [Accessed 1 April 2014].

Ramazan, S., and Dimitrakopoulos, R. 2004. Traditional and new MIP models for production scheduling with in-situ grade variability. International Journal of Surface Mining, Reclamation and Environment, 18(2), 85–98.

Ramazan, S. and Dimitrakopoulos, R. 2013. Production scheduling with uncertain supply: A new solution to the open pit mining problem, Optim. and Eng., 14(2), 361–380.

Ravenscroft, P., 1992. Risk analysis for mine scheduling by conditional simulation. Min. Technol. (Trans. Inst. Min. Metall. A), 101, 104–108.

Sarker, R.A., and Gunn, E.A. 1997. A simple SLP algorithm for solving a class of nonlinear programs. European Journal of Operational Research, 101(1), 140–154.

Tolwinski, B., and Underwood, R. 1996. A scheduling algorithm for open pit mines. IMA Journal of Mathematics Applied in Business & Industry, 7, 247–270.

Whittle, J. 1988. Beyond optimisation in open pit design, in Proc. Computer Applications in the Mineral Industries, (ed. K. Fytas), 331–337, Rotterdam, Balkema.

Whittle, J. 1999. A decade of open pit mine planning and optimization — The craft of turning algorithms into packages, in Proc. 28th APCOM Symposium, 15-24, Golden, SME/AIME.

Wolsey, L.A. 1998. Integer programming, New York, Wiley.