

A Scalable Approach to Optimal Block Scheduling

Jorge Amaya

Center for Mathematical Modeling and DIM, Universidad de Chile

Daniel Espinoza

Department of Industrial Engineering, Universidad de Chile

Marcos Goycoolea

School of Business, Universidad Adolfo Ibañez

Eduardo Moreno

School of Engineering, Universidad Adolfo Ibañez

Thomas Prevost

Network Mining

Enrique Rubio

Department of Mining Engineering, Universidad de Chile

Abstract

A critical aspect of long-term open-pit mine planning consists in computing a production schedule based upon a block sequencing strategy. Such a schedule should specify when and if blocks should be extracted in such a way as to maximize NPV, while satisfying wall-slope and production capacity constraints. It is well known that this problem can be modeled with integer programming (IP). However, integer programming is not used in practice because the size of typical block models makes such problems intractable to standard IP solvers. In this article we describe a scalable IP-based methodology for solving very large (millions of blocks) instances of this problem. We show that embedding standard IP technologies in a local-search based algorithm we are able to obtain near-optimal solutions to large problems in reasonable time. This methodology has been tested in several mine wide block models.

Biography

Marcos Goycoolea obtained his PhD from the School of Industrial and Systems Engineering (ISYE) of Georgia Tech in 2006, and is currently an Assistant Professor at the Universidad Adolfo Ibañez School of Business in Santiago, Chile. He teaches Operations Management and Optimization to MBA and Graduate students and conducts research on integer programming methodologies and their application to natural resource management. His focus of research is on developing computational methods and software for large-scale optimization-based decision support systems. For more information visit his website at <http://mgoycool.uai.cl>.

Introduction

The long term planning of a mine operation consists of defining the life of the mine, the mining reserves, the capital requirements and the production capacity at which the value of the project will be maximized. These decisions enable analysts to make a decision of whether or not to invest to carry on the mining endeavor. The main steps to overcome in open pit mining are the computation of a regular block model,

the delineation of the final pit, the definition of a pushback sequence and the computation of a production schedule.

Traditionally the open pit planning problem has been solved in several steps in order to have a solution in a reasonable time frame for mine operations. The steps are described as follows:

- Block value model: this step consists of defining a block economic value as a function of the metals content, prices, costs. This step defines a priori the final destination of the block, which in a trivial operation would be the definition of ore and waste
- Final pit limit: this step consists of applying the Lerchs and Grossman (1965) algorithm to define the economic boundary that maximizes cumulative profit subject to the required slope angles.
- Pushback sequencing: this step consists of iteratively changing the original block values obtaining a sequence of nested pit limits. This is then used to define the different volumes of the ore body that would be available in time to feed a production schedule.
- Cut off grades: the cut off grades are computed to define the destination of the blocks, i.e., mill, leach pad, waste dump, that are available in the different pushbacks previously computed.

See Hustrulid and Kuchta (2006) for a more thorough introduction to open-pit mine planning.

To our knowledge, the first effort to formally describe a mathematical model to solve this problem in an integrated way is the work of Johnson (1968). Given the large size and inherent difficulty of Johnson's model, most articles in the academic literature consider a simplified version of the problem with a single methodology for block extraction. In these models it is decided a-priori what will be done with each block before it is even extracted. More specifically, using economic or fixed grade cut-off criteria, it is decided if the block will go to a mill for processing, or to a waste dump, regardless of what other blocks may be extracted, and how they are processed.

We formally describe this model as follows. Let B represent the set of blocks, t_{max} the number of time periods, and r_{max} the number of different resource types (for example, milling capacity, or trucking capacity). For each block b and each time period t define a variable,

$$x_{b,t} = \begin{cases} 1 & \text{if block } b \text{ has been extracted by time } t, \\ 0 & \text{otherwise.} \end{cases}$$

We define three important families of constraints for this model. The first constraint is for consistency of the variable definition. For every block b and every (non-final) time period t define constraint,

$$x_{b,t} \leq x_{b,t+1} \quad (1)$$

The second family of constraints imposes the wall-slope (or precedence) condition. That is, if block a is immediately above block b , or, if block a should be extracted before block b in order to ensure that the resulting wall-slope of the pit is not too steep, for every time period t define constraint,

$$x_{b,t} \leq x_{a,t} \quad (2)$$

Finally, for each resource r and each time period t , there should be a constraint of the form,

$$\sum_{b \in B} q_{r,b} (x_{b,t} - x_{b,t-1}) \leq c_{r,t} \quad (3)$$

where $q_{r,b}$ represents the amount of resource r consumed by extracting block b , and $c_{r,t}$ denotes the amount of resource r available in time period t . The objective function of the problem consists in maximizing,

$$\sum_{t=1}^{t_{\max}} \sum_{b \in B} p_{b,t} (x_{b,t} - x_{b,t-1}) \quad (4)$$

where $p_{b,t}$ represents the net present value of extracting block b in time period t . For consistency purposes of these last two family of constraints we will assume that $x_{b,t} = 0$ for $t = 0$.

Henceforth we will refer to the problem of maximizing (4) subject to (1), (2), (3) and integrality constraints as the Resource-Constraint Pit optimization problem, or RC-PIT.

The RC-PIT problem is very difficult to solve in practice because real block models can be very large, thus leading to problems with an intractable amount of variables. A number of authors have contributed to improved integer programming techniques for the RC-PIT problem. Some important contributions, among many others, include those of Johnson (1968), Dagdelen (1985), Dagdelen and Johnson (1986), Caccetta and Hill (2003), Ramazan and Dimitrakopoulos (2004), Fricke (2006), Boland et al. (2007) and Gaupp (2008). Interested readers should refer to Osanloo et al. (2008) for a more detailed review of exact optimization work on this problem.

While each of these methods has contributed to faster solution times, solving problem instances with more than 100,000 blocks remains elusive. Perhaps the only exception to this can be found in the work of Caccetta and Hill (2003) where it is claimed that instances with up to 250,000 blocks are solved. Unfortunately, due to commercial reasons, the authors do not back this claim with any replicable algorithm or methodology, except for preprocessing scheme. Moving beyond 250,000 blocks is made more difficult by the fact that just representing such large problems in memory can be very difficult. Today commercial packages such as Whittle 4X can handle up to 1.5 million blocks that need to be re-blocked in order to find a pit limit.

We now present a methodology for tackling very large problems (several millions of blocks). This methodology does not require block aggregation in order to work, and builds on previously developed integer programming developments. Thus, any improvements to state-of-the art integer programming methodologies for RC-PIT can be combined with our approach. Finally, this method is scalable and parallelizable. That is, the methodology will work regardless of the problem size and its performance will greatly improve in parallel computation frameworks.

The idea of the method is simple. Starting from a known feasible solution (which we call the incumbent), our approach seeks to find a solution which is “similar” and which yields an improved objective function value. In order to do this, the algorithm uses the incumbent as a guide by only considering alternative solutions that partially coincide with the incumbent. This is accomplished by means of a random search that iteratively fixes parts of the schedule and tries to optimally improve the remaining “unfixed” part using an integer-programming model. Whenever improvements are found the incumbent is updated, and the process is repeated. This type of approach is known in the literature as a Local Search methodology. For an introduction to this type of approach, see Aarts and Lenstra (2003).

In Section 2 we explain the methodology in more detail. In Section 3 we present computational results. In Section 4 we conclude with some final remarks.

Improving Feasible Solutions: A Local Search Heuristic

We will now represent feasible RC-PIT solutions with a vector u indicating the time at which each block is extracted. In this way, if $u_b = t$, we understand that block b is extracted in time t . We convene that $u_b = \infty$ whenever a block b is never extracted in a solution.

Given a solution u and a subset of blocks A we define the A -neighborhood of u as the set of all solutions v which coincide with u everywhere except, possibly, the blocks in A . More formally, v is in the A -neighborhood of u if and only if $u_a = v_a$ for all blocks $a \notin A$.

Given an incumbent solution u , our local-search algorithm works by iteratively defining different sets of blocks A and finding, for each of these, a solution v in the A -neighborhood of u having the best possible objective function value. Given u and A , this can be accomplished using the formulation described in Section 1, adding additional constraints to ensure that blocks $a \notin A$ are scheduled in the exact time that they are scheduled in u . More specifically, for each $a \notin A$ define constraints,

$$x_{a,t} = 1 \quad \text{for all } t \geq u_a, \quad (5)$$

$$x_{a,t} = 0 \quad \text{for all } t < u_a, \quad (6)$$

and maximize (4) subject to (1), (2), (3), (5) and (6). Observe that solving this is equivalent to solving a smaller version of the original RC-PIT problem, since constraints (5) and (6) are imposed in practice by eliminating the corresponding variables and adjusting the right-hand sides. After solving the resulting integer programming problem, we obtain a solution v by setting $v_b = t$ if and only if $x_{b,t} - x_{b,t-1} = 1$. If solution v has a better objective function value than u , we update the incumbent and repeat the process. If not, we attempt again with a different set A .

The key to using this method effectively is in choosing the proper sets A . Ideally, one would like to choose sets A which are small enough to ensure that the reduced RC-PIT instances can be solved quickly, and yet large enough so as to ensure that there exist improving solutions in the neighborhood. We next describe three basic strategies which we found work well in our test data.

a. The ‘‘Cone-Above’’ strategy.

Consider a block b , and define $P(b)$ as the set of all blocks which are predecessors of b . That is, block a will belong to $P(b)$ if and only if block a must be extracted before block b due to wall-slope constraints, or because block a is immediately above block b . In order to find a local improvement to a solution u we randomly select a block b and find the best solution in the $P(b)$ -neighborhood of u as indicated above. This is illustrated in Figure 1.

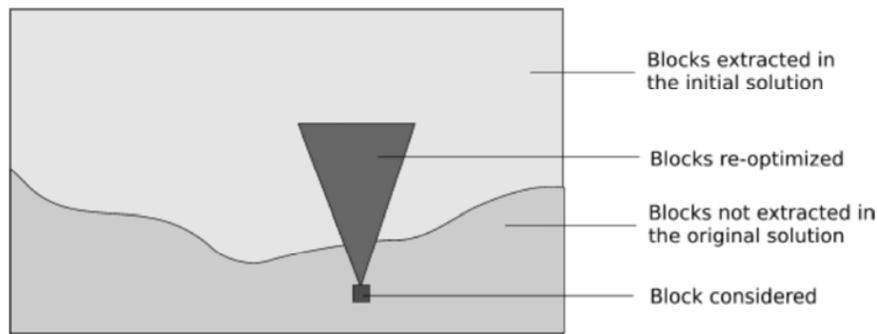


Figure 1: Illustration of the “cone-above” strategy

b. The “Periods” strategy.

Consider time periods t_1 and t_2 , and a solution vector u . Define the set,

$$I(t_1, t_2, u) = \{b \in B : t_1 \leq u_b \leq t_2\}$$

That is, $I(t_1, t_2, u)$ represents the set of all blocks extracted in solution u between time periods t_1 and t_2 . Observe that if $t_2 = \infty$ this could include blocks that are not extracted.

In order to find a local improvement of solution u we randomly select a pair of time periods t_1 and t_2 such that $|t_2 - t_1|$ is not too large, and find the best solution in the $I(t_1, t_2, u)$ -neighborhood of u .

c. The “Transversal” strategy.

Consider a distance d , and a height h . Define $D(d, h)$ as the set of all blocks at a vertical distance no greater than d of the set of blocks with height h . In order to find a local improvement of solution u we randomly select a height h and a distance d that is not too large, and find the best solution in the $D(d, h)$ -neighborhood of u .

We found that combining the use of these strategies helped to avoid getting stuck in local optima.

Case Studies: Analysis of Four Mines

Computational experiments were performed in four different ore bodies that are presented as follows:

Table 1: Description of the ore bodies used for the study

Name	# Blocks	Grade range	Observations
Marvin	61x60x17	0.03-1.46 %Cu 0.1-1.2 ppm Au	fictional copper gold ore body included in the Whittle 4X mine planning software
AmericaMine	61x42x60	% Cu : 0.08-3.68	hard rock polymetallic mine
AsiaMine	112x230x38	0-1.91 % Cu	Polymetallic ore body with a pipe shape
Andina	184x269x121	0.02-3.64 % Cu 0-0.42 % Mo	Copper molybdenum ore body taken from Andina Sur Sur deposit located at 50 Km north of Santiago. Typical porphyry copper ore body

The computations were performed on a Dell Poweredge 1950 Server with two Intel Xeon Quad Core E5345 processors (2.33 Ghz each) and 16GB of RAM. The software was written in the C programming language and compiled using GCC v3.4.6 on a Linux CentOS 4.5 operating system. All optimization runs were performed with the ILOG CPLEX v11.0 (henceforth CPLEX) optimization software callable library (running on the default settings).

Table 2 summarizes information regarding the block models corresponding to each ore body. “N. Blocks” describes the x,y,z dimensions of the raw block model. “Real Blocks” describes the amount of blocks remaining after removing air blocks. “P.P. Blocks” describes the number of blocks remaining after applying the ultimate-pit preprocessing scheme of Caccetta and Hill (2003) and the constrained precedence knapsack preprocessing schemes used by Samphai boon and Yamada (2000) and Gaupp (2008). “N. Periods” describes the number of time periods considered. All subsequent computations were performed on the preprocessed instances. As can be seen in Table 2, preprocessing reduces problem size dramatically.

The first step in our computational experiments consisted of running the heuristic described by Gershon (1987) in each of the problem instances. Second, we ran the proposed local search heuristic (starting from the Gershon solution) for a period of 4 hours. Third, we solved the LP relaxation of the full problem formulation. That is, the value obtained when solving the RC-PIT formulation directly using CPLEX, but replacing the integrality requirement with the condition that variables should take continuous values between 0 and 1. This value is a valid upper bound of the optimal solution, and hence serves as a guide to estimate the quality of our methodology. Finally, we let CPLEX integer programming solver attempt to solve each of the full RC-PIT problems to optimality. In Table 3 we present the results of these experiments. In column “Local Search” we indicate the objective value reached by the local search. In column “LP relaxation” we indicate the objective value of the LP relaxation. Finally, in the column “LP time” we describe the time (in minutes) required to solve each LP relaxation. All objective function values are given relative to the value obtained by the Gershon heuristic.

The first observation is that we were unable to solve the LP relaxation of the Andina mine since it was much too large to even load in memory. Second, we let the CPLEX IP solver run for 12 hours on each of the problem instances (except Andina), attempting to solve these problems and it could not find a feasible solution for any of the problem instances.

As can be seen in Table 3, in Marvin, AmericaMine and AsiaMine we are able to obtain near-optimal solutions in under 4 hours. While we can’t estimate how near optimality we are in the Andina mine (since

we could not solve the LP relaxation, and hence have no upper bound to compare against), it can be seen that the heuristic was able to afford a sizeable improvement in 4 hours despite the size of the problem.

The performance of our algorithm can be observed in more detail in Figure 2, where we track the objective function value over time, from 0 to 3 hours. In order to compare the different instances we normalize objective function values so that the LP relaxation upper bound has value 1.0. As can be seen in the Figure 1 in the Marvin, AmericaMine and AsiaMine mines we are able to obtain solutions near 1% optimality in just minutes.

As a final test, we let the local search algorithm run on the Andina mine for a full day. The objective function values at 8, 16 and 24 hours of running were 1.15, 1.16 and 1.17 respectively. This shows that though there are still important improvements to be found, they are modest relative to what is achieved in 4 hours.

Table 2: Description of the test set instances used for the study

	N.Blocks	Real Blocks	P.P. Blocks	N. Periods
Marvin	61x60x17	53668	8553	13
AmericaMine	61x42x60	19320	6445	18
AsiaMine	112x230x38	772800	97900	15
Andina	184x269x121	4320480	3340898	15

Table 3: Summary of local search performance after running 4 hours

	Gershon	Local Search (4 hrs)	LP relaxation	LP time
Marvin	1.0	1.08	1.09	26 min
AmericaMine	1.0	1.15	1.15	19 min
AsiaMine	1.0	1.23	1.24	4h 13 min
Andina	1.0	1.15	Unknown	Unknown

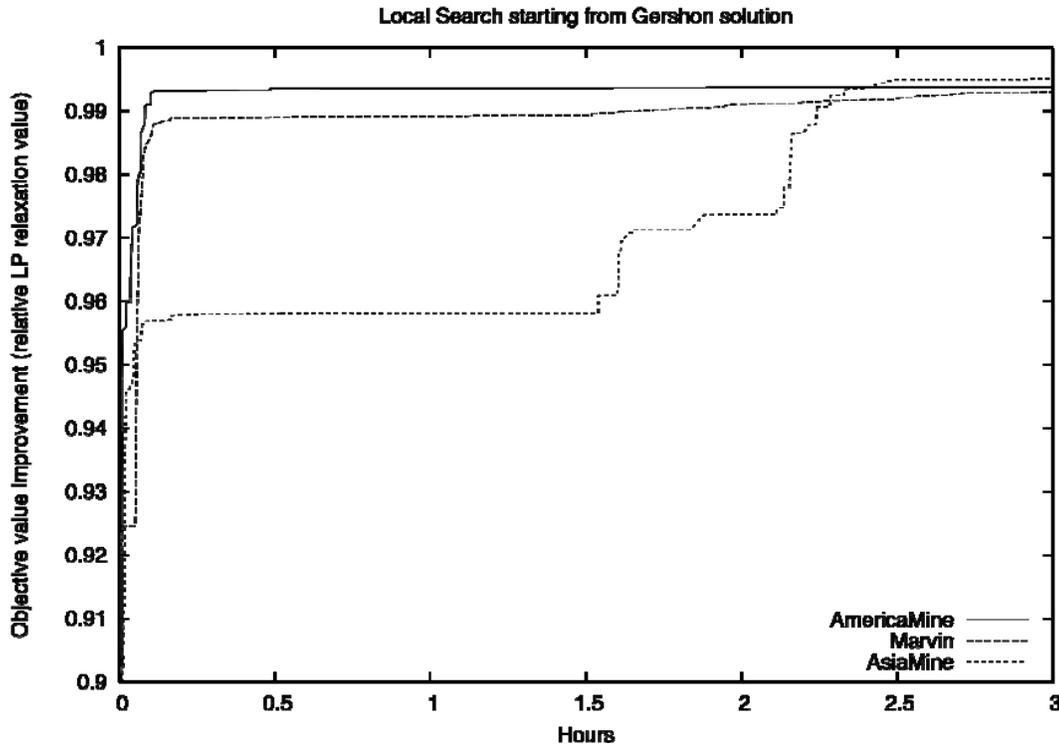


Figure 2: Objective value improvements over time when using local search heuristic

Final Remarks

In this article we have shown how to use a simple local-search based framework in order to tackle large instances of the RC-PIT problem. Our preliminary computational results are very promising, and suggest that our approach should be extended to consider more detailed models that incorporate multiple possible destinations for blocks after extraction, variable cut-off grades, use of stockpiles and other features. The results obtained can very likely be improved with a more thorough study of different neighborhoods and by the use of a distributed computing system in which different processors are independently and simultaneously searching for improvements, and being synchronized when these are found.

Acknowledgements

This interdisciplinary research was funded by the Chilean Government through FONDEF grant D06I1031 and through Basal Center Programs CMM (Center of Mathematical Modeling) and CTM (Center of Mining Technology) of Universidad de Chile. In addition, Marcos Goycoolea received funding from Fondecyt grant 11075028 and Daniel Espinoza from the Millennium Institute on Complex Systems. An important part of this work was done by Thomas Prevost as part of his master thesis, prior to working at Network Mining. The authors would like to thank IBM CPLEX for their support in terms of software and assistance. Also, the authors would like to thank Alexandra Newman for providing the AmericaMine data.

Bibliography

- AARTS E. and J.K. LENSTRA, 2003. Local search in combinatorial optimization. Princeton University Press, New Jersey. 536p.
- BOLAND N., C. FRICKE, and G.FROYLAND, 2007. A strengthened formulation for the open pit mine production scheduling problem. Available online at Optimization Online.
- CACCETTA L. and S.P. HILL, 2003. An application of branch and cut to open pit mine scheduling. *Journal of global optimization*, 27:349–365.
- DAGDALEN K., 1985. Optimum multi-period open pit mine production scheduling. PhD thesis, Colorado School of Mines, Golden, Colorado.
- DAGDALEN K. and T.B. JOHNSON, 1986. Optimum open pit mine production scheduling by lagrangian parameterization. 19th APCOM Symposium of the society of mining engineers (AIME).
- FRICKE C., 2006. Applications of Integer Programming in Open Pit Mining. PhD thesis, Department of Mathematics and Statistics, The University of Melbourne.
- GAUPP, M. 2008. Methods for improving the tractability of the block sequencing problem for an open pit mine. PhD thesis, Division of Economics and Business, Colorado School of Mines.
- GERSHON, 1987. Heuristic approaches for mine planning and production scheduling. *Int. Journal of Mining and Geological Engineering*, 5:1–13.
- HUSTRULID W. and M. KUCHTA, 2006. Open pit mine planning and design, 2nd Edition. Taylor and Francis. London, England. 971p.
- JOHNSON, T.B., 1968. Optimum open pit mine production scheduling. PhD thesis, Operations Research Department, University of California, Berkeley.
- LERCHS, H. and I.F. GROSSMAN, 1965. Optimum design of open-pit mines. *Transactions*, 68:17–24.
- OSANLOO, M., J. GHOLAMNEJAD, and B. KARIMI, 2008. Long-term open pit mine production planning: a review of models and algorithms. *International Journal of Mining, Reclamation and Environment*, 22:3–35.
- RAMAZAN, S. and R. DIMITRAKOPOLOUS, 2004. Recent applications of operations research and efficient mip formulations in open pit mining. *Society for mining, metallurgy and exploration meeting transactions*, 316:73–78.
- SAMPHAIBOON, N. and T. YAMADA, 2000. Heuristic and exact algorithms for the precedence constrained knapsack problem. *Journal of optimization theory and applications*, 105:659–676.