European Journal of Operational Research 249 (2016) 1169–1177

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

# Aggregation heuristic for the open-pit block scheduling problem

Enrique Jélvez [a,b,c], Nelson Morales [a,b], Pierre Nancel-Penard [a], Juan Peypouquet [d,*], Patricio Reyes [e]

[a] Advanced Mining Technology Center, Universidad de Chile, Santiago, Chile
[b] Delphos Mine Planning Laboratory, Department of Mining Engineering, Universidad de Chile, Santiago, Chile
[c] CSIRO Chile International Centre of Excellence, Santiago, Chile
[d] Department of Mathematics, Universidad Técnica Federico Santa María, Av España 1680, Valparaíso 2390123, Chile
[e] Technological Institute for Industrial Mathematics (ITMATI), University of Santiago de Compostela, Santiago de Compostela, Spain

ARTICLE INFO

*Article history:*
Received 15 October 2014
Accepted 20 October 2015
Available online 11 November 2015

*Keywords:*
Mine planning
Block aggregation
Open-pit block scheduling
Integer programming
Heuristics

ABSTRACT

In order to establish a production plan, an open-pit mine is partitioned into a three-dimensional array of blocks. The order in which blocks are extracted and processed has a dramatic impact on the economic value of the exploitation. Since realistic models have millions of blocks and constraints, the combinatorial optimization problem of finding the extraction sequence that maximizes the profit is computationally intractable. In this work, we present a procedure, based on innovative aggregation and disaggregation heuristics, that allows us to get feasible and nearly optimal solutions. The method was tested on the public reference library `MineLib` and improved the best known results in the literature in 9 of the 11 instances of the library. Moreover, the overall procedure is very scalable, which makes it a promising tool for large size problems.



## 1. Introduction

The mining industry is a very relevant economic sector. In Chile, where this research has been carried out, copper exports account for about 62.5 percent of the total exports and represent a 12 percent of the GDP (Cochilco, 2013).

Mines can be either open-pit or underground, the actual decision depending on different economic and technical considerations. In this paper we focus on *open-pit* mines, in which mineral is extracted by digging from the surface. Open-pit mines are preferred to underground mines because they can reach higher production levels, and have smaller operational costs. However, most of the time, it is necessary to remove material with poor or none ore content (*waste*) in order to have access to economically profitable material.

The actual value of a mine strongly depends on the order in which the material is extracted and processed. In order to define what portions of the terrain must be mined at different moments during the life-time of the mine, the planning horizon is discretized into *time-periods* (or *time-slots*). In turn, the terrain is divided into regular *blocks*, which are arranged in a 3-dimensional array. For each block, estimations on the ore content, density and other relevant attributes are constructed by using geostatistical methods. A *block model*, namely, the set of all blocks and their attributes, is the main input to the mine planning process. Using this information, it is possible to build a *block scheduling*, which specifies an extraction time-period for each block. The final value of a mine is therefore determined by the block model and the block scheduling.

The feasibility of a block scheduling for the open-pit method depends on *accessibility* and *extraction* constraints. First, before extracting one block, all the blocks above it must have been extracted. Moreover, stability of the walls must be ensured. This is expressed in terms of slope angles that must be satisfied at each moment. All these constraints are translated into *precedences* between blocks. On the other hand, there are certain capacity constraints, as well as other limitations, that are inherent to the process. The amount of material to be transported and processed at each period is subject to upper bounds given by transportation and plant capacity, respectively, which are usually expressed either in tons or hours. Further on, processed material must satisfy some *blending* constraints as well. The efficiency (or even feasibility) of the plant process depends on the attributes of the combination of blocks that are processed at a given period. For example, it may not be feasible to process alone a block with a high content of a certain pollutant (say arsenic), even it has a very high ore grade. Mixing it with another block (even a low ore-grade one) and processing them simultaneously may be possible because the blending

* Corresponding author. Tel.: +56 223037335.
*E-mail addresses:* ejelvez@ug.uchile.cl (E. Jélvez), nelson.morales@amtc.cl (N. Morales), pierre.nancel@amtc.uchile.cl (P. Nancel-Penard), juan.peypouquet@usm.cl (J. Peypouquet), patricio.reyes@usc.es (P. Reyes).

http://dx.doi.org/10.1016/j.ejor.2015.10.044
0377-2217/

provides an acceptable amount of the pollutant. Blending constraints can be either upper or lower bounds and apply to certain attributes of the blocks to be processed. Finally, the decision of how to process a block may depend on different parameters. Indeed, it is quite common for a mine to have more than one processing possibility (considering a block as waste and sending it to a waste pile is already one possibility). Depending on the final process or destination of the block, the net profit perceived by the mine is different, as are the blending constraints that apply to the process and the resource required to achieve this processing (different plant capacities, for example).

In this work, we propose and test a new numerical method to determine the block scheduling that maximizes the net present value for the exploitation of an open-pit mine. Our proposal is based on a combination of two approaches, that allow us to reduce the size of the problem and make it computationally tractable. The procedure aggregates blocks, uses integer programming techniques to solve incrementally the aggregated problem and produces solutions for the original instance in an innovative fashion. Using this methodology, we are able to provide nearly optimal solutions for some realistic-size problems that are otherwise numerically inaccessible.

The procedure was proved on the instances of the public reference library named `MineLib` (Espinoza, Goycoolea, Moreno, & Newman, 2013), which has three different types of open-pit mine planning problems for which good feasible solutions have been reported: the ultimate pit limit problem and two variants of open-pit production scheduling problems, for fictional cases, but also for real-life mine (for example, the instances KD, P4HD, W23 and McLaughlin correspond to actual copper and gold mines located in North America). We focus on the Constrained Pit Limit Problem (`CPIT`), which consists of the maximization of the net present value (NPV) of the exploitation over the time horizon, subject to precedence and operational constraints. The results obtained by our procedure improve nine out of the eleven instances available in the `MineLib` library. Moreover, the remaining two cases are within a gap of 0.2 percent of the optimum solution.

The paper is organized as follows: In Section 2 we provide a brief summary of the most relevant (and best-known) approaches found in the literature. Section 3 contains all the details concerning the modeling, notation and problem statement. The description of our methodological proposal, as well as the different heuristics involved, are presented in Section 4. All the implementation details, and the numerical results obtained are given and commented in Section 5. Finally, Section 6 contains some concluding remarks and perspectives.

## 2. Related work

A very general formulation, due to Johnson (1968), presents the block scheduling problem under slope, capacity and blending constraints (the last ones given by ranges of the processed ore grade) within a multi-destination setting, i.e., the optimization model decides what is the best process to apply to a given block. Unfortunately, at the time of its publication, the model was too complex to be solved in realistic case studies.

As an alternative to the work of Lerchs and Grossman (1965) proposed a very simplified version of the problem in which block destinations are fixed in advance, slope constraints are considered, but capacity or blending constraints are not. In this case, the problem reduces to selecting a subset of blocks such that the contained value is maximized while the precedence constraint induced by the slope angles are held. This problem is known as the *ultimate* or *final pit* problem. Lerchs and Grossman presented an efficient (polynomial) algorithm for solving the ultimate pit problem, and showed that reducing the economic value of any given block makes the optimal solution of the ultimate pit problem to *shrink*, in the sense that, if the values of the blocks decrease, the new solution is a subset of the original one. Therefore, it is possible to produce *nested pits* and, by trial and error, construct block schedules that satisfy other constraints like capacity. Present-day commercial software, like Gemcom (2011), is based on these facts.

As it turns out, while the model proposed by Johnson (and others) has always been regarded as superior in terms of the value it can add to a mining plan, it has been only recently that new developments (especially in algorithms) have allowed to solve or approximate this kind of models. Indeed, a main motivation of this work is to contribute to transform the theoretical superiority of these mathematical models into a practical one.

Picard (1976) showed that the ultimate pit problem is equivalent to the *maximum closure problem* in which, given a directed graph $G = (V, A)$ with weight function $w$ defined over the nodes, one looks for a subset of nodes $U \subset V$ such that $\sum_{u \in U} w(u)$ is maximal but $u \in U, (u, v) \in A \Rightarrow v \in U$. The maximum closure problem, in turn, can be reduced to the *min cut* problem (for more details see Nemhauser & Wolsey, 1988). Using this fact, Hochbaum and Chen (2000) proposes to attack the ultimate pit problem by means of existing efficient algorithms for the min cut problem.

Caccetta and Hill (2003) use a customized version of the *branch-and-bound* algorithm to solve problems up to a few hundreds of thousands of blocks under blending and capacity constraints. Their method can be used only for upper bounds. Bley, Boland, Fricke, and Froyland (2010) use a similar model but incorporating additional cuts based on the capacity constraints that strengthen the formulation of the problem, in the sense that the value of the linear relaxation provides a tighter bound. They test this approach on small instances (up to 500 blocks and 10 time-periods) on which they show very interesting improvements in the computational time. Unfortunately, it is not clear how to scale the technique for larger instances, as the number of cuts may explode very quickly. A closely related strategy is used by Fricke (2006), in order to find inequalities that improve various integer formulations of the same model. Gaupp (2008) reduces the size of the problem by deriving minimum and maximum extraction periods for each block, from the capacity constraints, and eliminating some of the variables. The method then applies Lagrangian relaxation to solve the problem.

The next two papers address the problem under consideration, but considering only upper bounds on resources consumption constraints: First, Amaya et al. (2009) starting from an initial feasible solution and then iteratively fix parts of the incumbent solution and re-optimize the complement. At each iteration, this defines an integer programming sub-problem that is solved exactly. They are able to solve instances of up to 4 million blocks and 15 time periods in 4 hours. In turn, Lamghari, Dimitrakopoulos, and Ferland (2014) use a hybrid method based on linear programming and variable neighborhood descent. The authors introduce a two-phase solution method: in the first one, they solve a series of linear programming problems to generate an initial solution. In the second phase, a variable neighborhood descent procedure is applied to improve the solution. The method is tested on some benchmark instances from the literature (some of `MineLib`), showing new best-known solutions for almost all of the instances, when compared to the solutions reported in Lamghari et al. (2014) and Espinoza et al. (2013). Indeed, only in two of these instances the solutions obtained have a larger gap, but this is still at much 0.2 percent.

Following Picard and Hochbaum ideas, Chicoisne, Espinoza, Goycoolea, Moreno, and Rubio (2012) and Bienstock and Zuckerberg (2010) address a problem which is very close to the one considered in this paper. However, they use Lagrangian relaxation on all but the precedence constraints (in this case the problem reduces to the ultimate pit problem). Using this approach, Chicoisne et al. focus on the case where there exists only one destination and one capacity constraint per period, and develop a customized algorithm (`CMA`) for