A study of the Bienstock-Zuckerberg algorithm[*]

Applications in Mining and Resource Constrained Project Scheduling

Gonzalo Muñoz[1], Daniel Espinoza[2], Marcos Goycoolea[3], Eduardo Moreno[4], Maurice Queyranne[5], and Orlando Rivera[6]

[1]*Industrial Engineering and Operations Research, Columbia University*
[2]*Gurobi Optimization*
[3]*School of Business, Universidad Adolfo Ibañez*
[4]*Faculty of Engineering, Universidad Adolfo Ibañez*
[5]*School of Business, University of British Columbia*
[6]*School of Business and Faculty of Engineering, Universidad Adolfo Ibañez*

April 11, 2017

We study a Lagrangian decomposition algorithm recently proposed by Dan Bienstock and Mark Zuckerberg for solving the LP relaxation of a class of open pit mine project scheduling problems. In this study we show that the Bienstock-Zuckerberg (BZ) algorithm can be used to solve LP relaxations corresponding to a much broader class of scheduling problems, including the well-known Resource Constrained Project Scheduling Problem (RCPSP), and multi-modal variants of the RCPSP that consider batch processing of jobs. We present a new, intuitive proof of correctness for the BZ algorithm that works by casting the BZ algorithm as a column generation algorithm. This analysis allows us to draw parallels with the well-known Dantzig-Wolfe decomposition (DW) algorithm. We discuss practical computational techniques for speeding up the performance of the BZ and DW algorithms on project scheduling problems. Finally, we present computational experiments independently testing the effectiveness of the BZ and DW algorithms on different sets of publicly available test instances. Our computational experiments confirm that the BZ algorithm significantly outperforms the DW algorithm for the problems considered. Our computational experiments also show that the proposed speed-up techniques can have a significant impact on the solve time. We provide some insights on what might be explaining this significant difference in performance.

Keywords: Column generation, Dantzig-Wolfe, Optimization, RCPSP

1. Introduction

Resource constrained project scheduling problems (RCPSPs) seek to optimally schedule activities over time in such a way as to comply with precedence and resource usage constraints. These problems can be notoriously difficult. Despite great progress in optimization methodologies during the last fifty years, there are instances of RCPSP involving as few as sixty activities that cannot be solved with today's most effective algorithms [24]. In multi-modal extensions of RCPSP, jobs can be processed in different ways (or modes). Changing the mode of a job affects its resource utilization, and its processing costs. In some applications, changing the mode of a job changes the duration of its processing time. In batch-processing extensions of RCPSP jobs are grouped together in clusters that must be processed together. For simplicity of notation, we will henceforth refer to multi-modal batch resource constrained project scheduling problems, simply as General Production Scheduling Problems, or GPSPs.

To date, the most effective methods for solving GPSPs, especially modal and batch variants, are based on integer programming methods that use the so-called time index formulations [39, 3, 7, 9]. These formulations define a binary variable for each job-execution time combination. An important limitation of these formulations is that the linear programming (LP) relaxations are very difficult to solve. This is because they tend to be large and degenerate, even for small problem instances.

While classical decomposition techniques and column generation approaches can be used for addressing some of these issues (specially complications related to the large number of variables), they often suffer from slow convergence rate. In this context, an important contribution was made by Bienstock and Zuckerberg [5, 6], where the authors presented an alternative to tackle these limitations effectively. They developed a new algorithm that can considerably outperform classical methods in a broad class of problems, thus providing a novel set of tools with a high practical impact and a wide range of potential extensions.

In this paper we study the Bienstock-Zuckerberg (BZ) algorithm [5, 6] in depth, and find that it can be used to overcome the stated limitations on a wide range of scheduling problems. We provide additional evidence to that in [5, 6], advocating for the efficacy of the algorithm in practice, and we provide new insights on it that allow further extensions. Specifically, we study the BZ algorithm as an approach to batch, multi-modal production scheduling problems where jobs can have arbitrary durations, but where these durations are not affected by changes of mode. The application of the BZ algorithm to this class of problems requires us to extend the algorithmic template as it was originally proposed. We are specifically concerned about this class of problems because, in addition to generalizing the well-known RCPSP, it generalizes three very important problems that arise in the context of mine planning: Underground Production Scheduling Problems (UPSPs), Open Pit Phase Design Problems (OPPDPs), and Open Pit Production Scheduling Problems (OPPSPs).

We present a new proof of algorithm correctness by casting the BZ algorithm as a column generation method, discuss algorithmic speed-ups, computationally test it on a variety of problem instances, and compare the methodology to the more traditional Dantzig-Wolfe decomposition (DW) method. As part of our analysis we prove that the BZ algorithm is closely related to a decomposition scheme that produces a bound somewhere in between that of the DW method, and that of the LP relaxation. This study allows us to conclude that the BZ algorithm is an effective

way of solving the LP relaxation of large time index formulations, significantly outperforming the DW method on all considered problem classes, and provides insights on the effectiveness of the methodology.

This article is organized as follows. In Section 2 we present a literature review of related scheduling work in mine planning applications and mathematical programming methodologies. In Section 3 we present an integer programming model that generalizes the classes of problems we are interested in studying. In Section 4 we present a reformulation of this model that fits the algorithmic framework we will be analyzing. In Section 5 we present a general column generation framework that can be used to solve the problem. We use this column generation approach to motivate the BZ algorithm, and facilitate a comparison to the DW method. We also introduce important speed ups for the BZ algorithm. Computational results analyzing the performance of BZ and comparing this performance to that of DW are presented in Section 6.

2. Background

Scheduling applications in Mine Planning

In this article we are mainly interested in addressing scheduling problems that are of relevance to planning applications in the mining industry.

The class of mining problems we are interested in include open pit and underground mine planning problems. In these problems, deposits are discretized into three-dimensional arrays known as block models. The problem to solve consists of deciding which blocks should be extracted, when they should be extracted, and what to do with the blocks once they are extracted. In this context, jobs correspond to extraction activities, modes correspond to different processing options and batches correspond to groups of blocks that must be extracted concurrently in order to comply with equipment extraction requirements. Resource constraints would be used to impose extracting, milling and refining capacity constraints. In an open-pit mine, precedence constraints would be used to impose that a mine must be extracted from the surface on downwards. In an under-ground mine, specifically in a stopping operation, precedence constraints would be used to impose that selected blocks branch out from a network of tunnels descending from the surface.

In all of these mining problems the objective is to maximize net-present-value of the extracted minerals. This is different from traditional scheduling problems, in which the objective is typically to minimize completion time metrics such as makespan, maximum tardiness, or total throughput time. Another difference between mine planning problems and traditional scheduling problem is that in mining it is not a strict requirement to execute all jobs. In all other ways, the Underground Production Scheduling Problems (UPSPs) is identical to the RCPSP. The Open Pit Phase Design Problems (OPPDPs) is a multi-modal RCPSP in which all jobs take a single-time period to complete, regardless of the selected mode. What the mode affects is the objective function value and the resource consumption of each activity. The Open Pit Production Scheduling Problems (OPPSPs) is just like the OPPDP, with the addition that there are batch constraints that force groups of blocks to be extracted simultaneously. These batch constraints are used to enforce equipment operability constraints, with each batch corresponding to contiguous set of

blocks typically called a *bench-phase* or *increment* in the mining industry.

For an introduction to optimization in underground mining see Alford et al. [1]. For related work in underground mining see Martinez and Newman [27], Newman and Kuchta [32] and O'Sullivan et al. [35, 36]. The user manuals of Deswik Scheduler [15] and MineMax iGantt [28] illustrate how UPSP is solved in practical mining applications. For an introduction to Open Pit Mining see Hustrulid and Kuchta [22], and the user manuals of Dassault Whittle [13], Mine-Max Scheduler [30], and MineMax Planner [29]. For a discussion on OPPSP see Goycoolea et al. [18]. For a general survey on OR applications in mine planning see Newman et al. [33].

Mathematical programming methodologies

To our knowledge, the first mathematical programming formulations of GPSPs dates back to Johnson [23] and Pritsker et al. [38], in the late 1960s. Each of these articles spawned its own track of academic articles on production scheduling problems. The first track, following the work of Johnson, mostly focused on strategic open pit mine planning problems. The second track, following the work of Pritsker et al., took a more general approach. Surprisingly, though many of the results found in these two tracks coincide, there are few, if any, citations connecting the two literatures.

The academic literature on exact optimization methods for GPSPs is immensely rich. Well-known surveys from the scheduling track include those of Graham et al. [19], Brucker et al. [8] and Hartmann and Briskorn [20]. In a recent book by Artigues et al. [2] a very thorough survey of GPSP is presented, including a computational study that compares existing methodologies on benchmark instances. Surveys from the mining track include those of Newman et al. [33], Espinoza et al. [16] and Osanloo et al. [34].

A brief summary of advances on solving the LP relaxation of time-index formulations is as follows. Since as early as the 1960s, most methods have been based on some form of decomposition that reduces the problem to a sequence of maximum closure or minimum cut problems. Johnson [23], in the context of mining (single mode OPPDPs), was the first to attempt such an approach with a Dantzig-Wolfe decomposition. Shortly after, and independently, Fisher [17], in the context of scheduling, proposed a Lagrangian Relaxation decomposition approach. Since then, a number of decomposition algorithms, primarily Lagrangian Relaxation decomposition algorithms, have been proposed for the problem in both literatures. Important examples include Dagdelen and Johnson [11] and Lambert and Newman [26] in mining (single mode OPPDPs), and Christofides et al. [10] and Mohring et al. [31] in scheduling.

Some recent approaches are as follows. Chicoisne et al. [9] consider single-mode OPPDPs with a single renewable resource, and propose a decomposition algorithm that solves the continuous relaxation in polynomial time. Boland et al. [7] consider a multi-modal variant of the same problem with two renewable resources, and propose a decomposition algorithm for the continuous relaxation that, while not provably polynomial, is very effective in practice. The BZ algorithm [5], developed shortly after, can be considered a generalization of this last algorithm that extends to multi-modal OPPMPs with an arbitrary number of renewable or non-renewable resources. This algorithm proved to be very efficient, even for extremely large instance sizes; it reduced the solving times drastically and it was able to tackle instances that could not be solved before. Berthold et al. [3] developed a branch-and-bound algorithm that combines mathematical

programming and constraint programming methods to close a large number of open benchmark instances of RCPSP. Zhu et al. [39] developed a mathematical programming based algorithm for solving multi-modal variants of RCPSP without batch constraints, but where mode changes affect duration times. In his work he closes a large percentage of open benchmark instances.

3. Integer programming formulation

We now describe an integer programming formulation for a class of production scheduling problems that generalizes the RCPSP, UPSP, OPPDP and OPPSP. As mentioned before, we simply refer to this class of problems as the General Production Scheduling Problem (GPSP).

Sets:

- $\mathscr{A}$: activities that must be scheduled in the problem.

- $\mathscr{C}$: clusters of activities that define a partition of $\mathscr{A}$.

- $prec(c)$: clusters that must be initiated no later than cluster $c \in \mathscr{C}$.

- $\mathscr{R} = \{1, \ldots, R\}$: resources that are consumed when carrying out activities.

- $\mathscr{T} = \{1, \ldots, T\}$: time periods in which it is possible to initiate activities.

- $\mathscr{M}_a = \{1, \ldots, M_a\}$: possible modes for activity $a \in \mathscr{A}$.

Parameters:

- t_a^- : release date for activity $a \in \mathscr{A}$.

- t_a^+ : due date for activity $a \in \mathscr{A}$.

- $d_{a,m}$: duration (number of time periods) of activity $a \in \mathscr{A}$ when executed in mode $m \in \mathscr{M}_a$.

- $p_{a,m,t}$: profit obtained if activity $a \in \mathscr{A}$ is initiated in period $t \in \mathscr{T}$ using mode $m \in \mathscr{M}_a$.

- $l(c_1, c_2)$: lag (number of time periods) that must elapse before activities in cluster $c_2 \in \mathscr{C}$ can be initiated, after activities in cluster $c_1 \in prec(c_2)$ have been initiated.

- $Q_{r,t}$: amount of resource $r \in \mathscr{R}$ available in time period $t \in \mathscr{T}$.

- $q_{r,a,m}$: amount of resource $r \in \mathscr{R}$ consumed by activity $a \in \mathscr{A}$ in each time period it is being executed, when executed in mode $m \in \mathscr{M}_a$.

Variables:

$$x_{c,t} = \begin{cases} 1 & \text{if the activities in cluster } c \text{ all start in time period } t \\ 0 & \text{otherwise} \end{cases}$$

$$y_{a,m,t} = \begin{cases} 1 & \text{if activity } a \text{ starts in time period } t \text{ using mode } m \\ 0 & \text{otherwise} \end{cases}$$

Objective function:

$$\text{maximize} \quad \sum_{a \in \mathscr{A}} \sum_{m \in \mathscr{M}_a} \sum_{t \in \mathscr{T}} p_{a,m,t} y_{a,m,t} \qquad (1)$$

Note that we express the objective function only in terms of the y variables. There is no loss of generality in this due to constraints (3), below.

Constraints:

Clusters can only be initiated once over the time horizon [$\forall c \in \mathscr{C}$]:

$$\sum_{t \in \mathscr{T}} x_{c,t} \leq 1. \qquad (2)$$

Activities in a cluster must start simultaneously, and must be carried out using a unique mode [$\forall c \in \mathscr{C}, \forall a \in c, \forall t \in \mathscr{T}$]:

$$x_{c,t} = \sum_{m \in \mathscr{M}_a} y_{a,m,t}. \qquad (3)$$

In order to initiate the activities in a cluster, all activities in preceding clusters must be initiated early enough so as to satisfy the lag requirement   [$\forall c_2 \in \mathscr{C}, \forall c_1 \in prec(c_2), \forall t \in \mathscr{T}$]:

$$\sum_{s \leq t} x_{c_2,s} \leq \sum_{s \leq t - l(c_1,c_2)} x_{c_1,s}. \qquad (4)$$

The amount of resources consumed cannot exceed the amount of resources available each period [$\forall r \in \mathscr{R}, t \in \mathscr{T}$]:

$$\sum_{a \in \mathscr{A}} \sum_{m \in \mathscr{M}_a} q_{r,a,m} \sum_{s = \max\{1, t - d_{a,m} + 1\}}^{t} y_{a,m,s} \leq Q_{r,t}. \qquad (5)$$

Activities can only be scheduled after their release dates [$\forall a \in \mathscr{A}$]:

$$\sum_{m \in \mathscr{M}_a} \sum_{t=1}^{t_a^- - 1} y_{a,m,t} = 0. \qquad (6)$$

Activities must be terminated no later than due dates and exactly one mode must be chosen for each executed activity [$\forall a \in \mathscr{A} : t_a^+ < \infty$]:

$$\sum_{m \in \mathscr{M}_a} \sum_{t=1}^{t_a^+ - d_{a,m} + 1} y_{a,m,t} = 1. \qquad (7)$$

Whenever $t_a^+ = \infty$, there is no due date for activity a and we do not include constraint (7). Note, however, that due to (2) and (3) we always have:

$$\sum_{m \in \mathscr{M}_a} \sum_{t \in \mathscr{T}} y_{a,m,t} \leq 1.$$

Thus, even when $t_a^+ = \infty$, there is an implicit constraint enforcing the choice of one mode and one time period at most.

It should be noted that the GPSP described by Formulation (1)-(7) generalizes the scheduling formulations found in Bienstock and Zuckerberg [6] in that it allows activities to have durations that span multiple time periods. This allows us to consider instances of RCPSP and UPSP which fell outside the scope of the Bienstock and Zuckerberg studies [5, 6].

Though this formulation is intuitive, and also the most commonly used formulation in the academic literature, it must be reformulated so as to fit the algorithmic schemes that will be presented.

4. Reformulation

In this section we present a reformulation of the GPSP formulation described in Section 3. As we will see in Section 5, this reformulation is key for the decomposition algorithms that we will discuss in this article.

For each $c \in \mathscr{C}, a \in \mathscr{A}, m \in \mathscr{M}_a$ and $t \in \mathscr{T}$ define,

$$w_{c,t} = \sum_{s=1}^{t} x_{c,s} \qquad \text{and} \qquad z_{a,m,t} = \sum_{i \in \mathscr{M}_a} \sum_{s=1}^{t-1} y_{a,i,s} + \sum_{i=1}^{m} y_{a,i,t}.$$

In this way, $w_{c,t}$ is a binary variable that takes value one if and only if cluster c is initiated "by" time t (i.e., no later than t). Likewise, $z_{a,m,t}$ is a binary variable that takes value one if and only if activity a is initiated by time $t-1$, or, if it is initiated in time period t, and its mode i is such that $i \leq m$.

To see that it is possible to formulate the production scheduling problem given by (1)-(7), using the (z, w) variables, note that we can map between the two variable spaces by means of the following linear sets of equations:

$$y_{a,m,t} = z_{a,m,t} - z_{a,m-1,t} \qquad \forall a \in \mathscr{A}, \, m = 2, \dots, M_a, \, t \in \mathscr{T} \qquad (8a)$$

$$y_{a,1,t} = z_{a,1,t} - z_{a,M_a,t-1} \qquad \forall a \in \mathscr{A}, \, t = 2, \dots, T \qquad (8b)$$

$$y_{a,1,1} = z_{a,1,1} \qquad \forall a \in \mathscr{A} \qquad (8c)$$

$$x_{c,t} = w_{c,t} - w_{c,t-1} \qquad \forall c \in \mathscr{C}, \, \forall t = 2, \dots, T \qquad (8d)$$

$$x_{c,1} = w_{c,1} \qquad \forall c \in \mathscr{C} \qquad (8e)$$

Using this mapping we can substitute out the variables in (1)-(7), to trivially obtain the following equivalent formulation:

Objective function:

$$\max \sum_{a \in \mathscr{A}} \sum_{m \in \mathscr{M}_a} \sum_{t \in \mathscr{T}} \tilde{p}_{a,m,t} z_{a,m,t} \qquad (9)$$

where,

$$\tilde{p}_{a,m,t} = \begin{cases} p_{a,m,t} - p_{a,m+1,t} & \text{if } m < M_a \\ p_{a,m,t} - p_{a,1,t+1} & \text{if } m = M_a, \, t < T \\ p_{a,m,t} & \text{if } m = M_a, \, t = T \end{cases}$$

7

Constraints:

For $a \in \mathscr{A}$, $m \in \{1, \ldots, M_a - 1\}$, $t \in \mathscr{T}$,:

$$z_{a,m,t} \leq z_{a,m+1,t}. \tag{10}$$

For $a \in \mathscr{A}$, $m = M_a$, $t \in \{1, \ldots, T-1\}$,:

$$z_{a,m,t} \leq z_{a,1,t+1}. \tag{11}$$

For $c \in \mathscr{C}$, $a \in c$, $t \in \mathscr{T}$,

$$w_{c,t} = z_{a,M_a,t}. \tag{12}$$

For $c_2 \in \mathscr{C}$, $c_1 \in prec(c_2)$, $t \in \{1 + l(c_1, c_2), \ldots, T\}$,

$$w_{c_2,t} \leq w_{c_1,t-l(c_1,c_2)}. \tag{13}$$

For $r \in R$, $t \in \mathscr{T}$,

$$\sum_{a \in \mathscr{A}} \sum_{m \in \mathscr{M}_a} \sum_{s \in \mathscr{T}} \tilde{q}_{a,m,s}^{r,t} z_{a,m,s} \leq Q_{r,t}. \tag{14}$$

where

$$\tilde{q}_{a,m,s}^{r,t} = \begin{cases} q_{r,a,m} \mathbf{1}_{[t-d_{a,m}+1,t]}(s) - q_{r,a,m+1} \mathbf{1}_{[t-d_{a,m+1}+1,t]}(s) & \text{if } m < M_a \\ q_{r,a,M_a} \mathbf{1}_{[t-d_{a,M_a}+1,t]}(s) - q_{r,a,1} \mathbf{1}_{[t-d_{a,1},t-1]}(s) & \text{if } m = M_a \end{cases}$$

and,

$$\mathbf{1}_{[x,y]}(s) = \begin{cases} 1 & \text{if } x \leq s \leq y \\ 0 & \text{otherwise.} \end{cases}$$

(The derivation of this last constraint from (5) can be found in Appendix A).

For $a \in \mathscr{A}$, $m \in \mathscr{M}_a$, $t < t_a^-$:

$$z_{a,m,t} = 0. \tag{15}$$

For $a \in \mathscr{A}$, $m \in \mathscr{M}_a$, $t \geq t_a^+$:

$$z_{a,m,t} = 1. \tag{16}$$

After the reformulation, if we substitute out the w variables by using linear equalities (12), we obtain what Bienstock and Zuckerberg [6] call a General Precedence Constrained Problem (GPCP) :

$$Z^* = \quad \max \quad c'z \tag{17}$$
$$\text{s.t.} \quad z_i \leq z_j \quad \forall (i,j) \in I, \tag{18}$$
$$Hz \leq h, \tag{19}$$
$$z \in \{0,1\}^n \tag{20}$$

In this problem constraints (18) correspond to constraints (10), (11), and (13), and constraints (19) correspond to constraints (14), (15) and (16).

It is interesting that despite the fact that the General Production Scheduling Problem (GPSP) formulated in Section 3 is more general than the scheduling problem formulations presented by Bienstock and Zuckerberg [5], both problems can be reformulated as an instance of GPCP.

5. Methodology.

In this section we describe a generalized version of the Bienstock-Zuckerberg (BZ) algorithm, originally introduced in [5, 6]. More specifically, we describe a decomposition algorithm well suited for solving the LP relaxation of large mixed integer programming problems having form,

$$
\begin{aligned}
Z^{IP} = \max \quad & c'z + d'u \\
\text{s.t.} \quad & z_i \le z_j, \quad \forall (i,j) \in I, \\
& Hz + Gu \le h, \\
& z \in \{0,1\}^n.
\end{aligned} \tag{21}
$$

Like Bienstock and Zuckerberg [6] before us, we will call this problem the General Precedence Constrained Problem (GPCP). It should be noted, however, that in our definition of GPCP we consider the presence of the extra u variables. These variables can be used for other modeling purposes. For example, they can be used to model variable capacities, or, in the context of mining problems, they can be used to model the use of stockpiles or other requirements.

Our presentation of the BZ algorithm is different than the original presentation in two respects: first, we cast it as a column generation method, and second, as stated before, we consider the presence of extra variables (the u variables in (21)). These differences require a new proof of algorithm correctness that we present below. The advantage of presenting the algorithm this way is that it is easier to compare to existing decomposition algorithms, and thus, in our opinion, becomes easier to understand and extend it. It also opens up the possibility of developing new classes of algorithms that might be useful in other contexts.

Before we present the BZ algorithm, we present a generic column generation (GCG) algorithm that can be used as a framework for understanding the BZ algorithm. This column generation scheme can be used to solve a relaxation of mixed integer problems having form:

$$
\begin{aligned}
Z^{IP} = \max \quad & c'z + d'u \\
\text{s.t.} \quad & Az \le b, \\
& Hz + Gu \le h, \\
& z \in \{0,1\}^n.
\end{aligned} \tag{22}
$$

Much like DW algorithm, this relaxation can yield bounds that are tighter than those obtained by solving the LP relaxation of the same problem. Throughout this section we will assume, without loss of generality, that $Az \le b$ includes $0 \le z \le 1$. Other than this, we make no additional assumption on problem structure.

5.1. A general column generation algorithm

In this section we introduce a General Column Generation (GCG) algorithm that will later motivate the BZ algorithm. This column generation algorithm is presented as a decomposition scheme for computing upper bounds of (22) that are no weaker than those provided by the LP relaxation.

Given a set $S \subseteq R^n$, let `lin.hull`(S) denote the linear space spanned by S. That is, let `lin.hull`(S) represent the smallest linear space containing S. Let $P = \{z \in \{0,1\}^n : Az \le b\}$. Define problem,

$$
\begin{aligned}
Z^{LIN} = \max \quad & c'z + d'u \\
\text{s.t.} \quad & Az \le b, \\
& Hz + Gu \le h, \\
& z \in \mathtt{lin.hull}(P),
\end{aligned}
\tag{23}
$$

The GCG algorithm that we present computes a value Z^{UB} such that $Z^{IP} \le Z^{UB} \le Z^{LIN}$. Observe that the optimal value Z^{LIN} of problem (23) is such that $Z^{IP} \le Z^{LIN} \le Z^{LP}$, where Z^{LP} is the value of the linear relaxation of problem (22).

The key to solving this problem is the observation that

$$
\mathtt{lin.hull}(P) = \{z : z = \sum_{i=1}^{d} \lambda_i v^i, \text{ for some } \lambda \in \mathbb{R}^n\},
\tag{24}
$$

where $\{v^1, \ldots, v^d\}$ is a basis for $\mathtt{lin.hull}(P)$. If V is a matrix whose columns $\{v^1, \ldots, v^d\}$ define a basis of $\mathtt{lin.hull}(P)$, it follows that problem (23) is equivalent to solving,

$$
\begin{aligned}
Z(V) = \quad \max \quad & c'V\lambda + d'u \\
\text{s.t.} \quad & AV\lambda \le b \quad (\alpha) \\
& HV\lambda + Gu \le h \quad (\pi).
\end{aligned}
\tag{25}
$$

In order for the optimal solution of (25) to be optimal for (23), it is not necessary for the columns of V to to define a basis of $\mathtt{lin.hull}(P)$. It suffices for the columns of V to span a subset of $\mathtt{lin.hull}(P)$ containing an optimal solution of (23). This weaker condition is what the column generation seeks to achieve. That is, starting with a matrix V with columns spanning at least one feasible solution of (23), the algorithm iteratively computes new columns until it computes the optimal value Z^{LIN}. As we will see, in some cases it is possible for the algorithm to terminate before having computed Z^{LIN}. In these cases the algorithm will terminate having computed a value Z^{UB} such that $Z^{UB} \le Z^{LIN}$.

Henceforth, let α, π denote the dual variables corresponding to the constraints of problem (25). To simplify notation, we will henceforth assume that α and π are always row vectors. Note that $\alpha, \pi \ge 0$.

To solve problem (23) with column generation we begin each iteration with a set of points $\{v^1, \ldots, v^k\}$ such that $\mathtt{lin.hull}(\{v^1, \ldots, v^k\}) \subseteq \mathtt{lin.hull}(P)$ contains at least one feasible solution of problem (23). At each iteration we add a linearly independent point v^{k+1} to this set, until we can prove that we have generated enough points so as to generate the optimal solution of (23).

The first step of each iteration consists in solving problem (25), with a matrix V^k having columns $\{v^1, \ldots, v^k\}$. Let $Z_k^L = Z(V^k)$. Let (λ^k, u^k) and (α^k, π^k) be the corresponding optimal primal and dual solutions. Note that $Z_k^L = \alpha^k b + \pi^k h$.

Define $z^k = V^k \lambda^k$. It is clear that (z^k, u^k) is feasible for (23); hence, $Z_k^L \le Z^{LIN}$. However, is (z^k, u^k) optimal for (23)? To answer this question, observe that (α, π) is dual-feasible for problem (25) if and only if $\pi G = d'$ and,

$$
\bar{c}(v, \alpha, \pi) \equiv c'v - \alpha Av - \pi Hv = 0, \qquad \forall v \in V.
$$

Let V^P represent a matrix whose columns $\{v^1,\ldots,v^{|P|}\}$ coincide with set P. Since V^P is clearly a generator of `lin.hull`(P), we know that the optimal solution of (25) (for V^P) corresponds to an optimal solution of (23). Thus, if $\bar{c}(v,\alpha^k,\pi^k)=0$ for all $v \in V^P$ (or equivalently, $v \in P$), we conclude that (z^k,u^k) is optimal for (23), since we already know that $\pi^k G = d'$. On the other hand, if (z^k,u^k) is not optional for (23), we conclude that there must exist $v \in P$ such that $\bar{c}(v,\alpha^k,\pi^k) \neq 0$.

This suggests how the column generation scheme for computing an optimal solution of (23) should work. Solve (25) with V^k to obtain primal and dual solutions (z^k,u^k) and (α^k,π^k). Keeping with traditional column generation schemes, we refer to problem (25) as the *restricted master problem*. If $\bar{c}(v,\alpha^k,\pi^k)=0$ for all $v \in P$, then (z^k,u^k) is an optimal solution of (23). If this condition does not hold, let $v^{k+1} \in P$ be such that $\bar{c}(v^{k+1},\alpha^k,\pi^k) \neq 0$. Note that v^{k+1} must be linearly independent of the columns in V^k. In fact, every column v of V^k must satisfy $\bar{c}(v,\alpha^k,\pi^k)=0$, hence so must any vector that can be generated with this set. We refer to the problem of computing a vector v^{k+1} with $\bar{c}(v^{k+1},\alpha^k,\pi^k) \neq 0$ as the *pricing problem*. Obtain a new matrix V^{k+1} by adding column v^{k+1} to V^k, and repeat the process. Given that the rank of V^k increases strictly with each iteration, the algorithm must finitely terminate.

The pricing problem can be tackled by solving:

$$
\begin{aligned}
L(\pi) = \quad \max \quad & c'v - \pi(Hv - h) \\
\text{s.t.} \quad & Av \leq b, \\
& v \in \{0,1\}^n
\end{aligned}
\tag{26}
$$

For this, at each iteration compute $L(\pi^k)$, and let $\hat{v}$ be the corresponding optimal solution. Observe that $L(\pi^k) \geq Z^{IP}$ for all $k \geq 1$. In fact, since the u variables are free in problem (25), and since (α^k,π^k) is dual-feasible in problem (25), we have that $d' - \pi^k G = 0$. This implies that problem (26), for π^k, is equivalent to

$$
\begin{aligned}
\max \quad & c'v + d'u - \pi^k(Hv + Gu - h) \\
\text{s.t.} \quad & Av \leq b, \\
& v \in \{0,1\}^n.
\end{aligned}
\tag{27}
$$

Given that $\pi^k \geq 0$, problem (27) is a relaxation of problem (22), obtained by penalizing constraints $Hz + Gu \leq h$. Hence $L(\pi^k) \geq Z^{IP} \; \forall k \geq 1$.

Define $Z_k^U = \min\{L(\pi^i) : i = 1,\ldots,k\}$ and note that $Z_k^U \geq Z^{IP}$, by the argument above. If $Z_k^U \leq Z_k^L$, we can define $Z^{UB} = Z_k^U$, and terminate the algorithm, as we will have that $Z^{IP} \leq Z^{UB} \leq Z^{LIN}$. In fact, since $Z_k^L \leq Z^{LIN}$, we have $Z^{IP} \leq Z^{UB} = Z_k^U \leq Z_k^L \leq Z^{LIN}$.

On the other hand, if $Z_k^U > Z_k^L$, then the optimal solution $\hat{v}$ of (26) is such that $\bar{c}(\hat{v},\alpha^k,\pi^k) \neq 0$, so we obtain an entering column by defining $v^{k+1} = \hat{v}$. To see this, note that

$$
\begin{aligned}
Z_k^U - Z_k^L > 0 &\Leftrightarrow c'\hat{v} - \pi^k(H\hat{v} - h) - \alpha^k b - \pi^k h > 0 \\
&\Leftrightarrow c'\hat{v} - \pi^k H\hat{v} - \alpha^k b > 0 \\
&\Rightarrow c'\hat{v} - \pi^k H\hat{v} - \alpha^k A\hat{v} > 0 \quad (\text{since } \alpha^k b \geq \alpha^k A\hat{v}) \\
&\Rightarrow \bar{c}(\hat{v},\alpha^k,\pi^k) > 0.
\end{aligned}
$$

Observe that, to be precise, we do not need to know a starting feasible solution in order to solve problem (23). In fact, if we do not have a feasible solution, we can add a variable to the right-hand-side of each constraint $Hz + Gu \leq h$, and heavily penalize it to proceed as in a Phase-I simplex algorithm (see [4]). The first iteration only needs a solution v^1 such that $Av^1 \leq b$.

5.2. Comparison to the Dantzig-Wolfe decomposition algorithm.

When solving problems with form (22) a common technique is to use the Dantzig-Wolfe decomposition [12] to compute relaxation values. The Dantzig-Wolfe decomposition (DW) algorithm is very similar to the General Column Generation (GCG) algorithm presented in the previous section. In fact, the DW algorithm computes the optimal value of problem,

$$
\begin{aligned}
Z^{DW} = \max \quad & c'z + d'u \\
\text{s.t.} \quad & Az \leq b, \\
& Hz + Gu \leq h, \\
& z \in \texttt{conv.hull}(P).
\end{aligned}
\tag{28}
$$

Given that $\texttt{conv.hull}(P) \subseteq \{z : Az \leq b\}$, this problem is typically written as,

$$
\begin{aligned}
Z^{DW} = \max \quad & c'z + d'u \\
\text{s.t.} \quad & z \in \texttt{conv.hull}(P), \\
& Hz + Gu \leq h,
\end{aligned}
\tag{29}
$$

or equivalently, as

$$
\begin{aligned}
Z^{DW} = \max \quad & c'V\lambda + d'u \\
\text{s.t.} \quad & \lambda \cdot 1 = 1, \\
& HV\lambda + Gu \leq h, \\
& \lambda \geq 0.
\end{aligned}
\tag{30}
$$

In this formulation, the columns $\{v^1, \ldots, v^k\}$ of V correspond to the extreme points of $\texttt{conv.hull}(P)$.

Given that $\texttt{conv.hull}(P) \subseteq \texttt{lin.hull}(P)$, it follows that $Z^{IP} \leq Z^{DW} \leq Z^{LIN} \leq Z^{LP}$. When $\texttt{conv.hull}(P) = \{z : Az \leq b\}$ it is easy to see that $Z^{DW} = Z^{LIN} = Z^{LP}$. However, the following example shows that all of these inequalities can also be strict:

$$
\begin{aligned}
Z^{IP} = \max \quad & -x_1 + 2x_2 + x_3 \\
\text{s.t.} \quad & -x_1 + x_2 \leq 0.5, \\
& x_1 + x_2 \leq 1.5, \\
& x_3 \leq 0.5, \\
& x \in \{0,1\}^3.
\end{aligned}
\tag{31}
$$

By decomposing such that $Hz + Gu \leq h$ corresponds to $-x_1 + x_2 \leq 0.5$ we get,

$$
\{z : Az \leq b\} = \{z \in [0,1]^3 : x_1 + x_2 \leq 1.5, \, x_3 \leq 0.5\},
$$

and,

$$
P = \{z \in \{0,1\}^n : Az \leq b\} = \{(0,0,0), (0,1,0), (1,0,0)\},
$$

and so,

$$\texttt{lin.hull}(P) = \{(x_1, x_2, x_3) : x_3 = 0\},$$
$$\texttt{conv.hull}(P) = \{(x_1, x_2, x_3) : 0 \leq x_1, \ 0 \leq x_2, \ x_1 + x_2 \leq 1, \ x_3 = 0\}.$$

From this it can be verified that $z^{IP} = 0.0$, $z^{DW} = 1.25$, $z^{LIN} = 1.5$, and $z^{LP} = 2.0$.

The DW algorithm is formally presented in Algorithm 1. The proof of correctness is strictly analogous to the proof presented in Section 5.1 for the generic column generation algorithm.

Observe that the DW pricing problem (see (35)) is exactly the same as the GCG pricing problem. However, since optimizing over $\{v : Av \leq b, \ v \in \{0, 1\}\}$ is exactly the same as optimizing over $\texttt{conv.hull}(P)$, we will have at every iteration $j \geq 1$ that $L(\pi^j) \geq Z^{DW}$. Thus, the bounds will never cross as they might in GCG, and there will be no early termination condition.

The main difference between the algorithms is in the master problem, where both solve very different linear models. One would expect that solving the master problem for the DW decomposition method would be much faster. In fact, let r_1 and r_2 represent the number of rows in the A and H matrices, respectively. The DW master problem has $r_2 + 1$ rows, compared to the $r_1 + r_2$ rows in the GCG master problem. However, this is only the case because of the way in which problem (22) is presented. If the first system of constraints, $Az \leq b$, instead has form $Az = b$, the algorithm can be modified to preserve this property. This equality form version of GCG, which we call GCG-EQ, is presented in Appendix B.

The discussion above suggests that for the GCG algorithm to outperform the DW algorithm, it would have to do less iterations. It is natural to expect that this would happen. After all, given a common set of columns, the feasible region of the GCG master problem is considerably larger than the corresponding feasible region of the DW master problem. That is, the dual solutions obtained by solving the GCG master problem should be "better" than those produced by the DW master problem, somehow driving the pricing problem to find the right columns quickly throughout the iterative process. This improved performance, of course, would only compensate on problems in which the DW algorithm has a tough time converging, and could come at the cost of a worse upper bound to the underlying integer programming problem. In fact, this slow convergence rate is exactly what happens in General Production Scheduling Problems (GPSPs), and is what motivates the BZ algorithm in the first place.

As column generation algorithms, the size of the master problem will be increasing by one in each iteration, both for the GCG and DW. At some point the number of columns could grow so large that solving the master problem could become unmanageable. An interesting feature of the DW algorithm is that this can be practically managed by removing columns that are not strictly necessary in some iterations. The key is the following Lemma, which is a direct result of well-known linear programming theory (Bertsimas and Tsisiklis [4]).

Lemma 1. *Let (λ^*, u^*) represent an optimal basic solution of problem*

$$
\begin{aligned}
Z(V) = \quad \max \quad & c'V\lambda + d'u \\
\text{s.t.} \quad & 1 \cdot \lambda = 1 \\
& HV\lambda + Gu \leq h \\
& \lambda \geq 0.
\end{aligned}
$$

Then, $|\{i : \lambda_i^* \neq 0\}| \leq r_2 + 1$, where r_2 is the number of rows in matrix H.

Algorithm 1: The DW Algorithm.

Input: A feasible mixed integer programming problem of form

$$Z^* = \quad \max \quad c'z + d'u$$
$$\text{s.t.} \quad Az \leq b$$
$$Hz + Gu \leq h \tag{32}$$
$$z \in \{0,1\}.$$

A matrix V whose columns are feasible 0-1 solutions of $Az \leq b$.

Output: An optimal solution (z^*, u^*) to problem

$$Z^{DW} = \max \quad c'z + d'u$$
$$\text{s.t.} \quad z \in \texttt{conv.hull}(P), \tag{33}$$
$$Hz + Gu \leq h.$$

1 $j \leftarrow 1$ and $Z_0^U \leftarrow \infty$.

2 $V^1 \leftarrow V$.

3 Solve the *restricted DW master problem*,

$$Z_j^L = \quad \max \quad c'V^j\lambda + d'u$$
$$\text{s.t.} \quad 1 \cdot \lambda = 1$$
$$HV^j\lambda + Gu \leq h \tag{34}$$
$$\lambda \geq 0$$

Let (λ^j, u^j) be an optimal solution to problem (34), and let (z^j, u^j), be the corresponding feasible solution of (33), where $z^j = V^j\lambda^j$. Let π^j be an optimal dual vector corresponding to constraints $HV^j\lambda + Gu \leq h$.

4 Solve the *DW pricing problem*. For this, consider the problem,

$$L(\pi) = \quad \max \quad c'v - \pi'(Hv - h)$$
$$\text{s.t.} \quad Av \leq b, \tag{35}$$
$$v \in \{0,1\}^n$$

and let v^j represent an optimal solution of $L(\pi^j)$. Let $Z_j^U \leftarrow \min\{L(\pi^j), Z_{j-1}^U\}$.

5 if $Z_j^L = Z_j^U$ **then**

6 $\quad$ $z^* \leftarrow z^j$. Stop.

7 else

8 $\quad$ $V^{j+1} \leftarrow [V^j, v^j]$.

9 $\quad$ $j \leftarrow j + 1$. Go to step 3.

In fact, what this Lemma says is that after m iterations, no matter how large m, we can always choose to remove all but $r_2 + 1$ columns, thus making the problem smaller. As a means of ensuring that the resulting column generation algorithm does not cycle, a reasonable technique

would be to remove columns only after the primal bound in the DW algorithm changes. That is, only in those iterations k such that $Z_k^L > Z_{k-1}^L$.

It should be noted that a variant of this Lemma is also true for the GCG-EQ algorithm (see Appendix B). However, there is no similar result for the GCG algorithm. That is, while unused columns can be discarded in GCG, there is no guarantee that the columns remaining in the master problem will be few in number.

5.3. The BZ algorithm

The BZ algorithm, originally proposed by Bienstock and Zuckerberg [5, 6] is a variant of the GCG algorithm that is specialized for General Precedence Constrained Problems (GPCPs). That is, the BZ algorithm assumes that constraints $Az \le b$ correspond to precedence constraints having form $z_i - z_j \le 0$, for $(i, j) \in I$, and bound constraints, having form $0 \le z \le 1$. This assumption allows the BZ to exploit certain characteristics of the optimal master solutions in order to speed up convergence. As we will see, the BZ algorithm is very effective when the number of rows in H and the number of u variables is small relative to the number of z variables. It should be noted, however, that the optimal value of the problem solved by the BZ algorithm will have value $Z^{BZ} = Z^{LP}$. That is, the bound will be no tighter than that of the LP relaxation. This follows directly from the fact that $\{z : z_i \le z_j \quad \forall (i, j) \in I\}$ defines a totally unimodular system.

The BZ algorithm is exactly as the GCG algorithm, with the exception of two changes:

First, the BZ algorithm uses a very specific class of generator matrices V^k. These matrices have two main properties. The first is that the linear space spanned by V^{k+1} will always contain z^k, the optimal solution of the master problem in the previous iteration. The second is that the columns of V^k are orthogonal 0-1 vectors. That is, for every column v^q of matrix V^k define $I^q = \{i \in \{1, \dots, n\} : v_i^q \ne 0\}$ (note that I^q describes the *support* of v^q, for $q = 1, \dots, k$). Then, 0-1 vectors v^q and v^r are orthogonal if and only if their supports are disjoint, i.e., $I^q \cap I^r = \emptyset$.

An intuitive explanation for why this would be desirable is that, when V^k has orthogonal 0-1 columns, problem (25) is equivalent to

$$
\begin{aligned}
v^* = \quad \max \quad & c'z + d'u \\
\text{s.t.} \quad & Az \le b \\
& Hz + Gu \le h \\
& z_i = z_j \qquad \forall i, j \in I^q, \ \forall q = 1, \dots, k.
\end{aligned}
\tag{36}
$$

That is, restricting the original feasible region to the linear space spanned by V^k is equivalent to equating the variables corresponding to the non-zero entries of each column in V^k. In a combinatorial optimization problem, this resembles a contraction operation, which is desirable because it can lead to a significant reduction of rows and variables, while yet preserving the original problem structure.

The matrix V^k used in the BZ algorithm is obtained by the following recursive procedure. Let $\hat{v}$ be the 0-1 vector obtained from solving the pricing problem. Let $[V^k, \hat{v}]$ be the matrix obtained by appending column $\hat{v}$ to V^k. We would like to compute a matrix V^{k+1} comprised of orthogonal 0-1 columns such that, first, $Span(V^{k+1}) \supseteq Span([V^k, \hat{v}])$; and second, such that V^{k+1} does not have too many columns. This can be achieved as follows. For two vectors $x, y \in \{0, 1\}^n$, define $x \wedge y \in \{0, 1\}^n$ such that $(x \wedge y)_i = 1$ if and only if $x_i = y_i = 1$. Likewise, define $x \setminus y \in \{0, 1\}^n$

15

such that $(x \setminus y)_i = 1$ if and only if $x_i = 1$ and $y_i = 0$. Assume that V^k is comprised of columns $\{v^1, \ldots, v^r\}$. Let V^{k+1} be the matrix made of the non-zero vectors from the collection:

$$\{v^j \wedge \hat{v} : 1 \leq j \leq r\} \cup \{v^j \setminus \hat{v} : 1 \leq j \leq r\} \cup \left\{\hat{v} \setminus \left(\sum_{j=1}^{k} v^j\right)\right\}.$$

We call this procedure of obtaining the matrix V^{k+1} from V^k and $\hat{v}$ the *refining* procedure. Note that in each iteration k, the refining procedure will produce a matrix V^{k+1} with at most $2r + 1$ columns.

The second difference between BZ and GCG is that it introduces a mechanism for preventing an unmanageable growth in the number of columns used in the master problem.

Consider any nonzero vector $z \in \mathbb{R}^n$. Let $\lambda_1, \ldots, \lambda_d$ denote the distinct non-zero values of the components of z, and for $i = 1, \ldots, d$ let $v^i \in \mathbb{R}^n$ denote the indicator vector of the components of z equal to λ_i, that is, v^i has components $v^i_j = 1$ if $z_j = \lambda_i$, and 0 otherwise. Note that $1 \leq d \leq n$ and $v^1, \ldots, v^d$ are orthogonal 0-1 vectors. Thus we can write $z = \sum_{i=1}^{d} \lambda_i v^i$ and we say that $v^1, \ldots, v^d$ define the *elementary basis* of z.

If at some iteration the number of columns in matrix V^k is too large, the BZ algorithm will replace matrix V^k with a matrix whose columns make up an elementary basis of the incumbent master solution z^k.

Again, there are two possible problems with this. On the one-hand, it is possible that such a coarsification procedure leads to cycling. On the other-hand, it might still be the case that after applying coarsification procedure the number of columns remains prohibitively large.

To alleviate the first concern, the BZ algorithm applies the coarsification procedure in exactly those iterations in which the generation of a new column leads to a strict improvement of the objective function. That is, when $Z^L_k > Z^L_{k-1}$. The second concern is alleviated by the fact that under the BZ algorithm assumptions, the elementary basis associated to a basic feasible solution of problem (36) will always have at most r_2 elements, where r_2 is the number of rows in matrix H (see Bienstock and Zuckerberg [5] for the original proof, or Appendix C for the proof adapted to our notation and extra variables).

A formal description of the column generation algorithm that results from the previous discussion is summarized in Algorithm 2.

5.4. BZ algorithm speedups

In this section we describe a number of computational techniques for improving the performance of the BZ algorithm. The computational techniques that we present for speeding up the master problem are generic, and can be used more generally in the GCG algorithm. However, the speedups that we present for the pricing problem are specific for solving generalized production scheduling problems (GPSPs), and thus, are specific for the BZ algorithm.

Algorithm 2: The BZ Algorithm.

Input: A feasible linear programming problem of form

$$Z^* = \begin{array}{ll} \max & c'z + d'u \\ \text{s.t.} & Az \leq b \\ & Hz + Gu \leq h \end{array} \tag{37}$$

where constraints $Az \leq b$ correspond to precedence constraints having form $z_i - z_j \leq 0$, for $(i,j) \in I$, and bound constraints, having form $0 \leq z \leq 1$. A matrix V whose columns are orthogonal 0-1 vectors spanning at least one feasible solution of (37).

Output: An optimal solution (z^*, u^*) to problem (37).

1 $j \leftarrow 1$ and $Z_0^U \leftarrow \infty$.
2 $V^1 \leftarrow V$.
3 Solve the *restricted BZ master problem*,

$$Z_j^L = \begin{array}{ll} \max & c'V^j\lambda + d'u \\ \text{s.t.} & AV^j\lambda \leq b \\ & HV^j\lambda + Gu \leq h \end{array} \tag{38}$$

Let (λ^j, u^j) be an optimal solution to problem (38), and let (z^j, u^j), be the corresponding feasible solution of (37), where $z^j = V^j\lambda^j$. Let π^j be an optimal dual vector corresponding to constraints $HV^j\lambda + Gu \leq h$.

4 Solve the *BZ pricing problem*. For this, consider the problem,

$$L(\pi) = \begin{array}{ll} \max & c'v - \pi'(Hv - h) \\ \text{s.t.} & Av \leq b, \end{array} \tag{39}$$

and let v^j represent an optimal solution of $L(\pi^j)$. Let $Z_j^U \leftarrow \min\{L(\pi^j), Z_{j-1}^U\}$.

5 **if** $Z_j^L = Z_j^U$ **then**
6 | $z^* \leftarrow z^j$. Stop.
7 **else**
8 | Obtain V^{j+1}, a matrix of orthogonal 0-1 columns, by refining V^j with v^j.
9 | $j \leftarrow j + 1$. Go to step 3.

5.4.1. Speeding up the BZ pricing algorithm

The pricing problem consists of solving a problem of the form

$$\begin{array}{ll} \max & \bar{c}'z \\ \text{s.t.} & z_i \leq z_j \qquad \forall (i,j) \in I \\ & 0 \leq z \leq 1 \end{array} \tag{40}$$

for some objective function $\bar{c}$, and a set of arcs I. This class of problems are known as *maximum closure* problems. In order to solve (40), we use the Pseudoflow algorithm of [21]. We use the following two features to speed-up this algorithm in the present context.

Pricing hot-starts [PHS]. The pricing problem is solved once per iteration of the BZ algorithm, each time with a different objective function. An important feature of the Pseudoflow algorithm is that it can be hot-started. More precisely, the Pseudoflow algorithm uses an internal data structure called a *normalized tree* that can be used to re-solve an instance of a maximum closure problem after changing the objective function vector. Rather than solve each pricing problem from scratch, we use the normalized tree obtained from the previous iteration to hot-start the algorithm. Because the changes in the objective function are not componentwise monotonous from iteration to iteration, we need to re-normalize the tree each time. For more information on the detailed working of the Pseudoflow algorithm, with indications on how to incorporate hot-starts, see [21].

Path contractions [PC]. Consider problem (40), and define an associated directed acyclic graph $G = (\mathcal{V}, \mathcal{E})$ as follows. For each variable z_i define a vertex v_i. For each precedence constraint $z_i \leq z_j$ with $(i, j) \in I$, define a directed arc (v_i, v_j) in $\mathcal{E}$. We say that a directed path $P = (v(1), v(2), \ldots, v(k))$ in G is *contractible* if it is a maximal path in G such that (i) $k \geq 3$, and (ii) every internal vertex $v(2), \ldots, v(k-1)$ has both in-degree and out-degree equal to one.

Observe that the variables associated to this path can only take $k+1$ different combinations of 0-1 values: either (i) $z_{v(i)} = 0$ for $i = 1, \ldots, k$ and the total contribution of the nodes in P is zero; or else, (ii) there exists an index $j \in \{1, \ldots, k\}$ such that $z_{v(i)} = 0$ for all $i = 1, \ldots, j-1$ and $z_{v(i)} = 1$ for all $i = j, \ldots, k$.

Since (40) is a maximization problem, in an optimal integer solution to (40) the contribution of the path will either be 0 (this will happen when $z_{v(k)} = 0$), or the contribution will be $\max_{1 \leq j \leq k} \sum_{i=j}^{k} \bar{c}_{v(i)}$ (which will happen when $z_{v(k)} = 1$).

This suggests the following arc-contracting procedure.

Let $j(P, \bar{c}) = \operatorname{argmax}_{1 \leq j \leq k} \sum_{i=j}^{k} \bar{c}_{v(i)}$ and define a new graph $\hat{G} = (\hat{\mathcal{V}}, \hat{\mathcal{E}})$ from G by eliminating the vertices $v(2), \ldots, v(k-1)$ from $\mathcal{V}$ and the arcs incident to them, and then adding an arc that connects $v(1)$ and $v(k)$. Define $\hat{c}$ with $\hat{c}_v = \bar{c}_v$ for all vertices $v \notin \{v_1, \ldots, v_k\}$, $\hat{c}_{v(k)} = \sum_{i=j(P,\bar{c})}^{k} \bar{c}_{v(i)}$ and $\hat{c}_{v(1)} = \sum_{i < j(P,\bar{c})} \bar{c}_{v(i)}$.

Solving the pricing problem in this smaller graph $\hat{G}$ with the objective $\hat{c}$ gives an optimal solution to the original problem on graph G with objective $\bar{c}$, with the same optimal objective value.

This procedure can be used to contract all contractible paths in G in order to obtain, in some cases, a significantly smaller graph. In fact, problem instances with multiple destinations and batch constraints induce many contractible paths in the pricing problem graph. This is illustrated with an example in Figure 1.

It should be noted that identifying and contracting paths only needs to be done in the first iteration. In all subsequent runs of the pricing problem, it is possible to use the same contracted graph, after updating the indices $j(P, \bar{c})$ and the objective function $\hat{c}$.
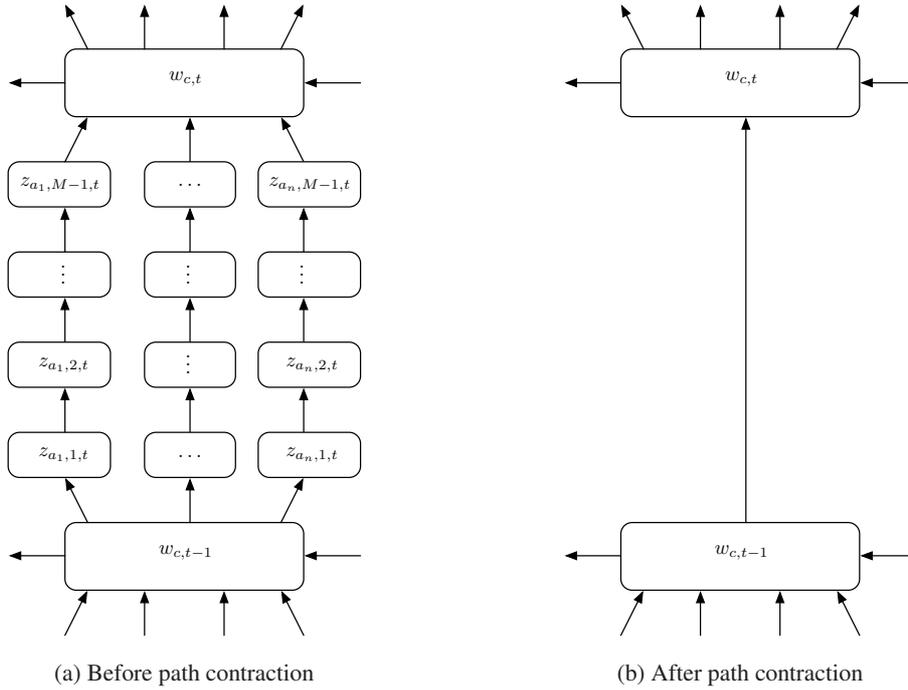
(a) Before path contraction	(b) After path contraction

Figure 1: Example illustrating the possible effect of the path contraction speed-up. This example assumes a cluster c comprised of activities $\{a_1, \ldots, a_n\}$. Each activity is assumed to have M possible modes. The graph includes an arc between all pairs of variables for which there is a precedence relationship.

5.4.2. Speeding up the BZ master algorithm

The following ideas can be used to speed up solving the master problem (38) in the BZ algorithm:

Starting columns [STCOL]. As input, the BZ algorithm requires an initial set of columns inducing a (possibly infeasible) solution of the problem. Such columns can be obtained by computing an elementary basis associated to a solution obtained with heuristics.

If, when starting the algorithm, we do not have an initial set of columns describing a feasible the problem, we can proceed as in the Phase-I simplex algorithm. For this, start with some arbitrary set of columns, and add "artificial variables" for each row.

Master hot-starts [MHS]. Because of the refining procedure, the restricted BZ master problem (38) in an iteration may be quite different from that in the previous iteration. To reduce its solve time, we can feed any simplex-based solver the solution from the previous iteration so that it can be used to attempt and identify a good feasible starting basis.

k-Step column management [k-Step]. In the original BZ algorithm, the coarsification procedure is applied in every iteration where there is a strict increase in the primal objective function value. Sometimes, however, it is better to keep the existing columns so as to improve the convergence rate. To avoid having too many columns in each iteration, we use what we refer to as

a k-step column management rule. This rule, which requires as input an integer parameter k, works as follows: assume that we are at the m-th iteration of the BZ algorithm, where $m > k$. Further assume that in the previous iteration (iteration $m − 1$) there was a strict increase in the primal objective function value. The column matrix V^m used in the m-th iteration is built from scratch, by first obtaining an elementary basis associated to solution $z^{m−k}$, and then successively refining this basis with the vectors $v^{m−k+1}, v^{m−k+2}, \ldots, v^{m−1}$ (see Section 5.3).

6. Computational results

Our computational tests consider solving the linear relaxation of three different classes of problems. Specifically, we consider instances of OPPSP, OPPDP and RCPSP. The OPPSP and OP-PDP instances are based on both real and hypothetical mine-planning instances obtained from industry partners, and from the MineLib website [16]. In order to protect the confidentiality of the data sets obtained from our industry partners, we have renamed all of the instances using the names of Chilean volcanoes. In Table 1 we present some characteristics of these instances. Note that for each instance we have a specification regarding the maximum number of time periods to consider. For each of the OPPSP problems we count with two versions: one with with less, and one with more clusters. We refer to these as the fine and course versions of each problem. These clusters correspond to what mining engineers refer to as a bench-phase, or increment (see [22] for formal definitions). We do not have cluster definition data for all of the problem instances, thus, the set of instances considered for OPPSP is a subset of the instances considered for OPPDP. The RCPSP instances that we consider are obtained from PSPlib [25] (datasets *j30*, *j60*, *j90* and *j120*) and from [14] (dataset *RG300*). It should be noted that these problems have a different objective function than the OPPDP and OPPSP problems. That is, these problems seek to minimize Makespan, rather than maximize net present value (NPV).

We evaluate the performance of the BZ algorithm, comparing it to the DW algorithm, and to a version of DW that incorporates the dual-stabilization techniques (DW-S) proposed by Pessoa et al. [37]. For each of the algorithms we attempt to compute the linear relaxation solution of the test problems, stopping early if the relative difference between the master problem bound and the pricing problem bound is less than $10^{−6}$ (the default reduced cost tolerance used by the CPLEX simplex algorithm). We do not report the bounds obtained by each algorithm since they will all coincide with the value of the LP relaxation for each problem.

We also evaluate the performance of the speed-up techniques described in this paper. Specifically, we analyze the performance of Pricing hot-starts (PHS), Path contractions (PC), Starting columns (STCOL), Master hot-starts (MHS), and k-step column management (k-Step), with $k = 10$. All of these speed-up techniques, with the exception of k-step (which is not applicable) are also implemented for the DW and DW-S algorithms. In fact, our code uses the exact same implementations of these speed-up techniques for the BZ, DW and DW-S algorithms.

To assess the value of our proposed speed-up techniques we defined a *default* set of speed-up-techniques to activate for each class of problems considered. For the OPPDP and OPPSP we defined the default speed-up-techniques to be PHS, MHS and PC. We found these to be the

Table 1: Description of the different instances of the test set used in our computational study.

Instance	Blocks (UPIT)	Modes	Periods	Resources	Clusters	
					Course	Fine
antuco	3 525 317	2	45	2	—	—
calbuco	198 248	3	21	1	324	548
chaiten	287 473	2	20	2	273	475
dospuntas	897 609	2	40	4	—	—
guallatari	57 958	3	21	3	272	460
kd	12 154	2	12	1	53	93
lomasblancas	1 492 024	3	45	3	—	—
marvin	8 516	2	20	2	56	98
mclaughlin_limit	110 768	2	15	1	166	290
palomo	97 183	2	40	2	44	93
ranokau	304 977	2	81	2	186	296
sm2	18 388	2	30	2	—	—
zuck_large	96 821	2	30	2	—	—
zuck_medium	27 387	2	15	2	—	—
zuck_small	9 399	2	20	2	—	—

features that improved performance of the algorithm, without requiring a heuristic to generate a solution. For the RCPSP class of problems we use a different set of default speed-up options. Since we have to compute a feasible solution to each problem in order to define the number of time periods, we change the default settings so as to use this as a starting solution (STSOL). In addition, we observe that these problem instances have signficantly more time periods per activities when compared to the OPPSP and OPPDP instances. This resulted in a very different algorithm behavior that prompted us to include k-Step column management as a default option, as it improved problem performance.

In order to assess the contribution of each individual speed-up technique in the algorithm, we turned each of these off (or on) to measure the impact from the change relative to that of the default settings.

All algorithms were implemented using C programming language, using CPLEX® 12.6 as optimization solver. The machines were running Linux 2.6.32 under x86_64 architecture, with four eight-core Intel® Xeon® E5-2670 processors and with 128 Gb of RAM. For all results, we present normalized geometric means comparing the performance versus BZ algorithm under default configuration.

6.1. Results for OPPDP

Table 2 shows the time required to solve each problem using the different proposed algorithms. The table also shows the number of iterations required by each algorithm, as well as the final number of columns generated by BZ. We do not report the bound generated by each relaxation, since they all generate the same bound for each instance.

Table 2: Comparison between the different algorithms for OPPDP instances

Instance	Time (sec)				Iterations			BZ cols
	BZ	DW	DW+S	CPLEX	BZ	DW	DW+S	
antuco	208 000	1 542 712	353 313	-	97	2 539	559	31 640
calbuco	1 480	6 715	3 602	-	34	243	132	1 876
chaiten	2 620	35 339	9 625	-	34	884	259	5 800
dospuntas	36 100	451 515	108 105	-	52	1 542	385	112 943
guallatari	143	1 668	494	-	30	444	158	2 421
kd	7	22	16	15 560	23	115	90	365
lomasblancas	86 600	521 279	208 256	-	75	773	360	5 635
marvin	9	28	17	34 601	27	154	99	206
mclaughlin_limit	167	1 096	539	-	27	208	102	1 273
palomo	1 290	9 550	3 637	-	45	587	264	2 089
ranokau	192 000	1 526 274	210 506	-	108	7 261	1 025	124 406
sm2	44	256	109	254	52	874	281	3 594
zuck_large	634	3 564	1 761	-	46	416	194	3 195
zuck_medium	40	112	93	954 405	28	107	82	229
zuck_small	11	37	24	56 165	29	150	113	294
Norm. G. Mean	1	5.98	2.42	-	1	11.8	4.9	

As can be seen, DW is nearly six times slower than BZ. Using stabilization techniques, the DW-S is able to signficantly reduced the time required to reach the optimality condition. In fact, stabilization techniques reduce the time of DW in a 60%, but it is still 2.42 times slower than BZ. This can be explained by the number of iterations required to converge by each algorithm. Even though a BZ iteration is, in average, 2.1 times slower than a DW iteration, DW and DW+S require 11.8 and 4.9 times the number of iterations of BZ to converge. This is possible because the BZ algorithm can produce a large number of useful columns in few iterations. In fact, the number of columns generated by BZ is considerably larger than the number of columns produced by DW and DW-S (the number of columns for these algorithms is equal to the number of iterations). Finally, note that CPLEX is only able to solve the five smallest instances.

In Table 3 we show what happens when we change the tolerance used to define optimality in all of the algorithms. It can be seen that BZ algorithm still outperforms the DW and DW-S algorithms after reducing the target optimality gap to 10^{-4} and 10^{-2}. However, the performance difference narrows as the target gap becomes smaller. The resulting times, normalized to the time of BZ under default setting, are presented in Table 3.

Finally, we study the impact of the different speed-up techniques on BZ. From the default settings, we disable the speed-ups PHS, MHS and PC, one-by-one, and also, we disable all three of them together. We also run BZ under our default setting after adding a starting column

Table 3: Normalized time required for different optimality gaps for OPPDP instances

Instance	10^{-6}			10^{-4}			10^{-2}		
	BZ	DW	DW+S	BZ	DW	DW+S	BZ	DW	DW+S
antuco	1	7.42	1.70	0.96	6.28	1.49	0.93	4.63	1.20
calbuco	1	4.54	2.43	0.92	3.49	1.94	0.74	1.55	1.08
chaiten	1	13.49	3.67	0.95	10.53	2.92	0.79	4.30	1.57
dospuntas	1	12.51	2.99	0.83	7.77	2.16	0.67	2.40	1.15
guallatari	1	11.66	3.45	0.86	7.35	2.44	0.70	2.60	1.37
kd	1	3.05	2.25	0.90	2.40	1.78	0.76	1.13	1.18
lomasblancas	1	6.02	2.40	0.95	5.04	2.05	0.87	3.26	1.38
marvin	1	3.19	1.99	0.91	2.26	1.47	0.81	0.94	0.81
mclaughlin_limit	1	6.56	3.23	0.86	4.75	2.47	0.68	1.89	1.47
palomo	1	7.40	2.82	0.98	6.10	2.33	0.88	2.69	1.50
ranokau	1	7.95	1.10	0.82	4.79	0.86	0.63	1.87	0.50
sm2	1	5.88	2.50	0.95	3.77	1.89	0.88	2.03	1.35
zuck_large	1	5.62	2.78	0.96	4.42	2.28	0.84	2.21	1.45
zuck_medium	1	2.80	2.31	0.87	2.15	1.78	0.66	1.23	0.97
zuck_small	1	3.22	2.13	0.86	2.20	1.52	0.77	1.24	0.85
Norm. G. Mean	1	5.98	2.42	0.90	4.36	1.89	0.77	2.04	1.14

(STCOL) generated using the Critical Multiplier algorithm (See [9]), and after enabling the k-Step feature. The resulting times, normalized to the time of BZ under default setting, are presented in Table 4.

We can see that all speeding features provide, in average, an improvement on the time required by BZ to converge. However, the individual performance of each feature differs instance by instance. The most important speed-up technique is path contraction (PC), as disabling this feature increases the time required to converge by 60%. Since the reduction in number of variables of this speed-up in OPPDP is proportional to the number of modes of the problem, this feature is particularly important for the Guallatari instance, which has 3 different modes. Disabling these three features, the total time required to converge more than doubles. On the other hand, if we start with a preliminary set of starting columns, the time required is decreased by 32%. Finally, we see that k-Step does not improve the convergence time, making the algorithm run 35% slower. In fact, it makes every problem converge slower, with the exception of the ranokau instance.

6.2. Results for OPPSP

For these instances we use the same default settings as those used for OPPDP. We note that these problems are considerably smaller than the OPPSP problems. This is both in the number of variables and precedence constraints. All algorithms are able to solve these problems in a few

Table 4: Normalized time required for BZ algorithm with/without features for OPPDP instances

Instance	Default	No PHS	No PC	No MHS	No PHS/PC MHS	Default + STCOL	Default + k-Step
antuco	1	1.57	1.60	0.98	2.50	0.57	2.56
calbuco	1	1.48	1.77	1.10	2.78	0.85	1.09
chaiten	1	1.38	1.49	1.01	1.83	0.62	1.20
dospuntas	1	1.36	1.57	1.03	1.75	0.73	1.52
guallatari	1	1.47	2.18	1.21	2.88	1.00	0.99
kd	1	1.30	1.63	1.00	2.22	0.69	1.51
lomasblancas	1	1.84	1.68	1.01	2.94	0.25	1.98
marvin	1	1.32	1.27	0.97	2.47	0.91	1.49
mclaughlin_limit	1	1.27	1.78	0.94	2.06	0.77	1.20
palomo	1	1.60	1.47	0.98	2.21	0.94	1.53
ranokau	1	0.93	1.04	0.97	1.15	0.22	0.57
sm2	1	1.33	1.93	1.16	2.35	1.11	2.06
zuck_large	1	1.15	1.60	1.00	2.08	0.62	1.31
zuck_medium	1	1.75	1.59	1.20	2.74	0.84	1.13
zuck_small	1	1.32	1.72	0.95	2.41	0.92	1.27
Norm. G. Mean.	1	1.39	1.60	1.03	2.24	0.68	1.35

hours, with the exception of the *ranokau* instance, which CPLEX fails to due to the memory limit (128Gb). We present the results in Table 5.

We can see that in this class of problems the performance of DW and DW-S is more similar to that of BZ. This is probably explained by the fact that the number of iterations required by DW and DW-S is much smaller. Note also that stabilization techniques for DW only marginally reduce the number of iterations required by DW to converge, resulting in that DW-S is 18% slower than DW.

Comparing the impact of the different features on BZ (see Table 6), we see that the most important feature is, again, path contraction (PC). Disabling this feauture makes the algorithm run almost 10 times slower. Similarly to the OPPDP problems, providing starting columns to BZ reduces the time by 23%, and introducing k-Step column management makes the problem run slower (83%).

6.3. Results for RCPSP instances

In order to formulate the RCPSP instances we need a maximum number of time periods to consider. Since the objective of these problems is to minimize Makespan, the number of time

Table 5: Comparison between the different algorithms for OPPSP instances

Instance	Time (sec)				Iterations			BZ cols
	BZ	DW	DW+S	CPLEX	BZ	DW	DW+S	
calbuco	88	90	122	56 102	32	96	95	167
chaiten	86	174	184	4 595	34	189	149	392
guallatari	29	39	50	11 757	30	117	102	188
kd	1	1	1	5	18	44	40	70
marvin	2	2	2	17	27	70	73	82
mclaughlin_limit	14	19	22	2 310	25	78	73	100
palomo	49	59	95	429	44	138	137	160
ranokau	2 034	5 102	4084	-	186	1 461	781	2 108
Norm. G. Mean	1	1.38	1.63	54.57	1	3.64	3.17	

Table 6: Normalized time required for BZ algorithm with/without features for OPPSP instances

Instance	Default	No PHS	No PC	No MHS	Default + STCOL	Default + k-Step
calbuco	1	0.92	9.96	1.03	0.82	2.10
chaiten	1	0.95	14.63	1.04	0.81	2.06
guallatari	1	1.01	4.91	1.05	0.73	1.77
kd	1	1.04	5.63	1.05	0.65	1.59
marvin	1	0.90	3.51	0.92	0.68	1.26
mclaughlin_limit	1	0.99	14.65	0.93	0.98	1.91
palomo	1	0.87	13.33	0.97	0.77	1.83
ranokau	1	1.03	19.48	1.03	0.78	2.37
Norm. G. Mean	1	0.96	9.25	1.00	0.77	1.83

periods should be an upper bound on the number of periods required to schedule all of the activities. Such a bound can be computed by running any heuristic to compute any feasible integer solution. For this purpose, we use a greedy TOPOSORT heuristic [9], which takes fractions of a second to solve for all of our instances.

Table 7 describes the performance of the different algorithms on our RCPSP test instances. Note that we only consider instances that are solved by CPLEX in more than 1 second, obtaining a total of 1575 instances. The running times presented in the table are geometric means over instances in the same dataset.

Table 7: Comparison between the different algorithms for RCPSP instances

Dataset †		Time (sec)				Iterations		
		BZ	DW	DW+S	CPLEX	BZ	DW	DW+S
j30	(51 inst.)	1.23	4.71	1.91	1.61	141.4	650.4	312.9
j60	(152 inst.)	0.98	8.19	2.55	3.77	68.5	477.0	232.1
j90	(293 inst.)	0.54	2.27	1.07	5.33	23.3	115.4	73.4
j120	(599 inst.)	3.95	33.13	11.94	31.78	58.2	440.5	230.5
RG300	(480 inst.)	22.86	43.76	24.87	240.51*	91.1	393.9	260.5
Norm. G. Mean		1	4.76	2.00	7.25	1	5.8	3.3

†: We only consider instances which took CPLEX more than a second to solve.

∗: We only consider the 438 instances which were solved within 48 hours.

Table 7 shows that BZ is again faster than the other algorithms when solving RCPSP instances. In fact, it is 2 times faster than DW+S and 4.7 faster than DW. Note that this difference is particularly large for the *j120* instances from PSPLIB repository, where DW and DW+S run 8.4 and 3 times slower, respectively. It would seem that the performance of these algorithms is greatly dependent on problem structure. This can be seen when considering the *RG300* instances, which are generated in a different way. In these instances, DW with stabilization is only marginally slower than BZ.

Table 8 shows how much the performance of the BZ algorithm is affected by turning off each of the default speed up features. It is interesting to note that on RCPSP instances the *k*-Step column management rule is very important. Turning it off makes the problem run, in average, 2.51 times slower. This is in stark contrast to what happens with the OPPSP and OPPDP instances, where activating the *k*-Step feature actually makes the problem run slower. We speculate that this is due to the fact that the RCPSP problems have a significantly greater number of time periods per activity than the OPPSP and OPPDP instances. This results in a problem with significantly more resource consumption constraints per variable. It is also interesting to note that disabling the PC and MHS features actually makes BZ run faster in the RCPSP instances. We speculate that the MHS feature does not improve performance because, having enabled the *k*-Step feature as well, succesive master problem formulations significantly differ from each other. We speculate that the PC feature does not improve performance because

in RCPSP instances there are not as many paths to contract due to the structure of the precedence graphs. This is due to the fact that there are no clusters and that there is just a single mode per activity.

Table 8: Normalized time required for BZ algorithm with/without features for RCPSP instances

Dataset †	Default	No PHS	No PC	No MHS	No k-Step	No STCOL
j30 (51 inst.)	1	1.30	0.78	0.80	0.74	1.44
j60 (152 inst.)	1	1.53	0.99	1.03	1.44	1.91
j90 (293 inst.)	1	1.15	0.86	0.91	1.51	1.80
j120 (599 inst.)	1	1.40	0.92	0.95	2.39	1.23
RG300 (453 inst.)	1	1.22	0.87	0.89	5.78	1.04
Norm. G. Mean	1	1.30	0.90	0.93	2.51	1.33

†: We only consider instances for which the BZ algorithm finished in 48 hours in all settings.

6.4. Concluding remarks

We summarize with three important conclusions. First, the BZ algorithm significantly outperforms DW and DW+S on all GPCP problem classes. Second, the speed-up features proposed in this article significantly improve the performance of the BZ algorithm. Third, the algorithm and speed-up performances greatly depend on the problem classes that are being considered. These conclusions suggest that the BZ algorithm with the proposed speed-ups is a technique that can be used to effectively be used to compute the LP relaxation solution of different classes of scheduling problems. Such an algorithm might be useful as a means to provide upper bounds for scheduling problems, or as a part of the many rounding heuristics that have recently been proposed, or eventually, as part of a branch-and-bound solver. In addition, they suggest that the BZ algorithm's potential use for other classes of problems should also be further studied.

References

[1] Alford, C., Brazil, M., Lee, D.: Optimisation in underground mining. In: Handbook of operations research in natural resources, pp. 561–577. Springer (2007)

[2] Artigues, C., Demassey, S., Neron, E.: Resource-constrained project scheduling: models, algorithms, extensions and applications, vol. 37. John Wiley & Sons (2010)

[3] Berthold, T., Heinz, S., Lübbecke, M., Möhring, R., Schulz, J.: A constraint integer programming approach for resource-constrained project scheduling. In: Integration of AI and OR techniques in constraint programming for combinatorial optimization problems, pp. 313–317. Springer (2010)

[4] Bertsimas, D., Tsitsiklis, J.: Introduction to linear optimization, vol. 6. Athena Scientific Belmont, MA (1997)

[5] Bienstock, D., Zuckerberg, M.: A new LP algorithm for precedence constrained production scheduling. Optimization Online (2009)

[6] Bienstock, D., Zuckerberg, M.: Solving LP relaxations of large-scale precedence constrained problems. Proceedings from the 14th conference on Integer Programming and Combinatorial Optimization (IPCO). Lecture Notes in Computer Science 6080 pp. 1–14 (2010)

[7] Boland, N., Dumitrescu, I., Froyland, G., Gleixner, A.: LP-based disaggregation approaches to solving the open pit mining production scheduling problem with block processing selectivity. Computers & Operations Research **36**, 1064–1089 (2009)

[8] Brucker, P., Drexl, A., Möhring, R., Neumann, K., Pesch, E.: Resource-constrained project scheduling: Notation, classification, models, and methods. European Journal of Operational Research **112**(1), 3–41 (1999)

[9] Chicoisne, R., Espinoza, D., Goycoolea, M., Moreno, E., Rubio, E.: A new algorithm for the open-pit mine production scheduling problem. Operations Research **60**(3), 517–528 (2012)

[10] Christofides, N., Alvarez-Valdés, R., Tamarit, J.: Project scheduling with resource constraints: A branch and bound approach. European Journal of Operational Research **29**(3), 262–273 (1987)

[11] Dagdelen, K., Johnson, T.: Optimum open pit mine production scheduling by lagrangian parameterization. In: Proceedings of the 19th International Symposium on the Application of Computers and Operations Research in the Mineral Industry (APCOM) (1986)

[12] Dantzig, G., Wolfe, P.: Decomposition principle for linear programs. Operations research **8**(1), 101–111 (1960)

[13] Dassault Systèmes: GEOVIA Whittle (2015). URL http://www.gemcomsoftware.com/products/whittle

[14] Debels, D., Vanhoucke, M.: A decomposition-based genetic algorithm for the resource-constrained project-scheduling problem. Operations Research **55**(3), 457–469 (2007)

[15] Deswik: Deswik.sched (2015). URL https://www.deswik.com/product-detail/deswik-scheduler/

[16] Espinoza, D., Goycoolea, M., Moreno, E., Newman, A.: Minelib: A library of open pit mining problems. Annals of Operations Research **206**, 93–114 (2013)

[17] Fisher, M.: Optimal solution of scheduling problems using lagrange multipliers: Part i. Operations Research **21**(5), 1114–1127 (1973)

[18] Goycoolea, M., Espinoza, D., Moreno, E., Rivera, O.: Comparing new and traditional methodologies for production scheduling in open pit mining. In: Proceedings of the 37th International Symposium on the Application of Computers and Operations Research in the Mineral Industry (APCOM), pp. 352–359 (2015)

[19] Graham, R., Lawler, E., Lenstra, J., Kan, A.: Optimization and approximation in deterministic sequencing and scheduling: a survey. Annals of Discrete Mathematics **5**, 287–326 (1979)

[20] Hartmann, S., Briskorn, D.: A survey of variants and extensions of the resource-constrained project scheduling problem. European Journal of Operational Research **207**(1), 1–14 (2010)

[21] Hochbaum, D.: The pseudoflow algorithm: A new algorithm for the maximum-flow problem. Operations Research **56**, 992–1009 (2008)

[22] Hustrulid, W., Kuchta, K. (eds.): Open Pit Mine Planning and Design. Taylor and Francis, London, UK (2006)

[23] Johnson, T.: Optimum open pit mine production scheduling. Ph.D. thesis, Operations Research Department, Universiry of California, Berkeley (1968)

[24] Kolisch, R., Sprecher, A.: PSP-Library. Online at `http://www.om-db.wi.tum.de/psplib/datamm.html` (1997). [Online; accessed March-2015]

[25] Kolisch, R., Sprecher, A.: PSPLIB - a project scheduling problem library. European Journal of Operational Research **96**(1), 205 – 216 (1997)

[26] Lambert, W.B., Newman, A.M.: Tailored lagrangian relaxation for the open pit block sequencing problem. Annals of Operations Research **222**(1), 419–438 (2014)

[27] Martinez, M., Newman, A.: A solution approach for optimizing long-and short-term production scheduling at LKAB's Kiruna mine. European Journal of Operational Research **211**(1), 184–197 (2011)

[28] MineMax: iGantt. `https://www.minemax.com/solutions/products/igantt` (2015)

[29] MineMax: Planner. `http://www.minemax.com/solutions/requirements/strategic-planning` (2015)

[30] MineMax: Scheduler (2015). URL `http://www.minemax.com/solutions/requirements/strategic-p`

[31] Möhring, R., Schulz, A., Stork, F., Uetz, M.: Solving project scheduling problems by minimum cut computations. Management Science **49**(3), 330–350 (2003)

[32] Newman, A., Kuchta, M.: Using aggregation to optimize long-term production planning at an underground mine. European Journal of Operational Research **176**(2), 1205–1218 (2007)

[33] Newman, A., Rubio, E., Caro, R., Weintraub, A., Eurek, K.: A review of operations research in mine planning. Interfaces **40**, 222–245 (2010)

[34] Osanloo, M., Gholamnejad, J., Karimi, B.: Long-term open pit mine production planning: a review of models and algorithms. International Journal of Mining, Reclamation and Environment **22**(1), 3–35 (2008)

[35] O'Sullivan, D., Newman, A.: Extraction and backfill scheduling in a complex underground mine. Interfaces **44**(2), 204–221 (2014)

[36] O'Sullivan, D., Newman, A., Brickey, A.: Is open pit production scheduling 'easier' than its underground counterpart? Mining Engineering **67**(4), 68–73 (2015)

[37] Pessoa, A., Sadyvok, R., Uchoa, E., Vanderbeck, F.: In-out separation and column generation stabilization by dual price smoothing. In: 12th International Symposium on Experimental Algorithms(SEA), Rome, Lecture Notes in Computer Science, vol. 7933, pp. 354–365 (2013)

[38] Pritsker, A., Waiters, L., Wolfe, P.: Multiproject scheduling with limited resources: A zero-one programming approach. Management science **16**(1), 93–108 (1969)

[39] Zhu, G., Bard, J., Yu, G.: A branch-and-cut procedure for the multimode resource-constrained project-scheduling problem. INFORMS Journal on Computing **18**(3), 377–390 (2006)

Appendix

A. Resource constraints in GPSP reformulation

In this section we provide a detailed proof on how to derive (14) from (5). This is needed in order to reformulate an instance of a General Production Scheduling Problem (GPSP), described in Section 3, as an instance of a General Precedence Constrained Problem (GPCP), presented in Section 4. For the benefit of the reader, we recall (14) and (5).

Proposition 2. *Consider inequality* (5)

$$\sum_{a \in \mathscr{A}} \sum_{m \in \mathscr{M}_a} q_{r,a,m} \sum_{s=\max\{1, t-d_{a,m}+1\}}^{t} y_{a,m,s} \leq Q_{r,t},$$

as defined in Section 3. If we apply the variable substitution scheme introduced in Section 4:

$$y_{a,m,t} = z_{a,m,t} - z_{a,m-1,t} \qquad \forall a \in \mathscr{A}, \, m = 2, \ldots, M_a, \, t \in \mathscr{T},$$
$$y_{a,1,t} = z_{a,1,t} - z_{a,M_a,t-1} \qquad \forall a \in \mathscr{A}, \, t = 2, \ldots, T,$$
$$y_{a,1,1} = z_{a,1,1} \qquad \forall a \in \mathscr{A},$$

we obtain inequality (14)*:*

$$\sum_{a \in \mathscr{A}} \sum_{m \in \mathscr{M}_a} \sum_{s \in \mathscr{T}} \tilde{q}_{a,m,s}^{r,t} z_{a,m,s} \leq Q_{r,t},$$

where

$$\tilde{q}_{a,m,s}^{r,t} = \begin{cases} q_{r,a,m} \mathbf{1}_{[t-d_{a,m}+1,t]}(s) - q_{r,a,m+1} \mathbf{1}_{[t-d_{a,m+1}+1,t]}(s) & \text{if } m < M_a \\ q_{r,a,M_a} \mathbf{1}_{[t-d_{a,M_a}+1,t]}(s) - q_{r,a,1} \mathbf{1}_{[t-d_{a,1},t-1]}(s) & \text{if } m = M_a \end{cases}$$

Proof.

$$
\begin{aligned}
Q_{r,t} \;\geq\; & \sum_{a\in\mathscr{A}}\sum_{m\in\mathscr{M}_a} q_{r,a,m} \sum_{s=\max\{1,t-d_{a,1}+1\}}^{t} y_{a,m,s} \\[2mm]
=\; & \sum_{a\in\mathscr{A}}\left(\sum_{m=2}^{M_a}\sum_{s=\max\{1,t-d_{a,m}+1\}}^{t} q_{r,a,m}y_{a,m,s} + \sum_{s=\max\{1,t-d_{a,1}+1\}}^{t} q_{r,a,1}y_{a,1,s}\right) \\[2mm]
=\; & \sum_{a\in\mathscr{A}}\left(\sum_{m=2}^{M_a}\sum_{s=\max\{1,t-d_{a,m}+1\}}^{t} q_{r,a,m}\left(z_{a,m,s}-z_{a,m-1,s}\right)\right. \\[2mm]
& \left. + \sum_{s=\max\{1,t-d_{a,1}+1\}}^{t} q_{r,a,1}z_{a,1,s} - \sum_{s=\max\{2,t-d_{a,1}+1\}}^{t} q_{r,a,1}z_{a,M_a,s-1}\right) \\[2mm]
=\; & \sum_{a\in\mathscr{A}}\left(\sum_{m=2}^{M_a}\sum_{s=\max\{1,t-d_{a,m}+1\}}^{t} q_{r,a,m}z_{a,m,s} - \sum_{m=1}^{M_a-1}\sum_{s=\max\{1,t-d_{a,m+1}+1\}}^{t} q_{r,a,m+1}z_{a,m,s}\right. \\[2mm]
& \left. + \sum_{s=\max\{1,t-d_{a,1}+1\}}^{t} q_{r,a,1}z_{a,1,s} - \sum_{s=\max\{1,t-d_{a,1}\}}^{t-1} q_{r,a,1}z_{a,M_a,s}\right) \\[2mm]
=\; & \sum_{a\in\mathscr{A}}\left(\sum_{m=1}^{M_a}\sum_{s=\max\{1,t-d_{a,m}+1\}}^{t} q_{r,a,m}z_{a,m,s} - \sum_{m=1}^{M_a-1}\sum_{s=\max\{1,t-d_{a,m+1}+1\}}^{t} q_{r,a,m+1}z_{a,m,s}\right. \\[2mm]
& \left. - \sum_{s=\max\{1,t-d_{a,1}\}}^{t-1} q_{r,a,1}z_{a,M_a,s}\right).
\end{aligned}
$$

$\square$

B. Equality form version of the GCG algorithm

In this section we present the GCG algorithm when, instad of the form (22), the mixed integer problem to be solved has form,

$$
\begin{aligned}
Z^{IP} = \max\quad & c'z + d'u \\
\text{s.t.}\quad & Az = b, \\
& Hz + Gu \leq h, \\
& z \in \{0,1\}^n.
\end{aligned}
\tag{41}
$$

That is, when the first set of inequalities have been replaced with equalities. As before, define

$$
P = \{z : Az = b,\ z \in \{0,1\}\}.
$$

We present a column generation algorithm for computing the optimal solution of

$$
\begin{aligned}
Z^{LIN} = \max\quad & c'z + d'u \\
\text{s.t.}\quad & Az = b, \\
& Hz + Gu \leq h, \\
& z \in \texttt{lin.hull}(P).
\end{aligned}
\tag{42}
$$

Let v^o be such that $Av^o = b$. For each $k \geq 1$ let V^k be a matrix comprised of linearly independent points $\{v^1, \ldots, v^k\}$ such that $Av^i = 0$ for $i = 1, \ldots, k$. Define,

$$Z(V) = \begin{array}{ll} \max & c'v^o + c'V\lambda + d'u \\ \text{s.t.} & HV\lambda + Gu \leq h - Hv^o \quad (\pi) \end{array} \tag{43}$$

Begin each iteration by solving problem (43) with $V = V^k$. Define $Z_k^L = Z(V^k)$ and let (λ^k, u^k) and π^k be the corresponding optimal primal and dual solutions respectively. Define $z^k = v^o + V^k\lambda^k$ and let,

$$L(\pi) = \begin{array}{ll} \max & c'v - \pi(Hv - h) \\ \text{s.t.} & Av = b, \\ & v \in \{0,1\}^n. \end{array} \tag{44}$$

Observe that since $\pi^k G = d'$, we have:

$$L(\pi^k) = \begin{array}{ll} \max & c'v + d'u - \pi^k(Hv + Gu - h) \\ \text{s.t.} & Av = b, \\ & v \in \{0,1\}^n. \end{array} \tag{45}$$

Thus, $L(\pi^k)$ is a relaxation of (41) obtained by penalizing constraints $Hz + Gu \leq h$ with $\pi^k \geq 0$. This implies $L(\pi^k) \geq Z^{IP}$ for all iterations k.

Solve problem (44) with $\pi = \pi^k$ and let $\hat{v}$ be the corresponding optimal solution. Define $Z_k^U = \min\{L(\pi^i) : i = 1, \ldots, k\}$. As before, we will have that $Z_k^U \geq Z^{IP}$ and $Z_k^L \leq Z^{LIN}$.

If, at any iteration, we have that $Z_k^L \geq Z_k^U$ we can halt the algorithm. By defining $Z^{UB} = Z_k^U$ we will have that $Z^{IP} \leq Z^{UB} \leq Z^{LIN}$.

On the other hand whenever $Z_k^L < Z_k^U$ we will have that column $\hat{v} - v^o$ has positive reduced cost in the master problem. Thus, letting $v^{k+1} = \hat{v} - v^o$ and $V^{k+1} = [V^k, v^{k+1}]$ we can continue iterating the algorithm.

C. Distinct fractional values Lemma

In this section we show that the BZ master problem will have bounded number of distinct fractional values. This fact is important in Section 5.3, in order to argue why the use of the *elementary basis* can reduce the number of columns in the BZ algorithm.

Consider the setting for the BZ algorithm, i.e, a problem of the form

$$Z^{BZ} = \begin{array}{ll} \max & c'z + d'u \\ \text{s.t.} & Az \leq b \\ & Hz + Gu \leq h \end{array} \tag{46}$$

where $\{z : Az \leq b\} = \{z : z_i \leq z_j \ \forall \ (i,j) \in I, \ 0 \leq z \leq 1\}$. It is well known that A defines a totally unimodular matrix in such case, thus all extreme points of $\{z : Az \leq b\}$ are 0-1 vectors.

Recall that z is an n-dimensional vector, and let m be the dimension of u. In this section we prove the following result:

Lemma 3. *Let $(\bar{z}, \bar{u})$ be an extreme point of* (46)*, and let q be the number of rows of H (and G). Then the number of distinct fractional values of $\bar{z}$ is at most $q - m$.*

Note that by assuming an extreme point $(\bar{z}, \bar{u})$ of (46) exists, we are implicitly assuming $q \geq m$. In fact, if $(\bar{z}, \bar{u})$ is an extreme point of (46), then $\bar{u}$ must be an extreme point of $\{u : Gu \leq h - H\bar{z}\}$. This, however, requires that G have at least m rows, thus implying that $q - m \geq 0$ holds.

An almost identical result is originally proved in [5]. We present such proof here, adapted to our notation and including the u variables, which modifies the final result. Note that this theorem also applies to every *restricted BZ master problem* (38), since the latter is obtained by simply equating sets of variables of (37), thus maintaining the same structure.

To prove Lemma 3 we will make use of the *elementary basis* introduced in Section 5.3, that is, we write

$$\bar{z} = \sum_{i=1}^{k} \lambda_i \bar{v}^i$$

where $\{\lambda_1, \ldots, \lambda_k\}$ are the distinct non-zero values of $\bar{z}$, and each $\bar{v}^i$ is the corresponding indicator vector for $\bar{z} = \lambda_i$. Without loss of generality we assume $\lambda_1 = 1$, so the fractional values of $\bar{z}$ are given by $\{\lambda_2, \ldots, \lambda_k\}$.

Lemma 4. *Let $(\bar{z}, \bar{u})$ be an extreme point of* (46)*, and decompose $\bar{z}$ as above. Additionally, denote $\bar{A}$ the sub-matrix of A corresponding to binding constraints at $(\bar{z}, \bar{u})$. Then $\bar{v}^i \in null(\bar{A})$ for all $2 \leq i \leq k$, and they are linearly independent.*

Proof. The $\bar{v}^i$ vectors are clearly linearly independent since they have disjoint support. As for the first claim, consider a precedence constraint $z_i \leq z_j \ (i, j) \in I$. If such constraint is binding in $(\bar{z}, \bar{u})$ then $\bar{z}_i = \bar{z}_j$, which implies

$$\bar{v}_i^l = 1 \Longleftrightarrow \bar{v}_j^l = 1, \quad l = 1, \ldots, k$$

thus, $\bar{v}_i^l = \bar{v}_j^l \ \forall l$. On the other hand, if a constraint $z_i \leq 1$ or $z_i \geq 0$ is binding at $(\bar{z}, \bar{u})$ then clearly $\bar{v}_j^l = 0$ for $l \geq 2$. This proves $\bar{A}\bar{v}^l = 0$ for $2 \leq l \leq k$. $\qquad\square$

Lemma 5. *Let $(\bar{z}, \bar{u})$ and q be as in Lemma 3. Additionally, let $\bar{A}$ be the sub-matrix of A corresponding to binding constraints at $(\bar{z}, \bar{u})$. Then $dim\left(null(\bar{A})\right) \leq q - m$.*

Proof. Let $\bar{H}, \bar{G}$ be the set of rows of H and G corresponding to active constraints in $(\bar{z}, \bar{u})$. Since this is an extreme point, the matrix

$$\begin{bmatrix} \bar{A} & 0 \\ \bar{H} & \bar{G} \end{bmatrix}$$

must contain at least $n + m$ linearly independent rows. Suppose $[\bar{H} \ \bar{G}]$ has $\bar{q}$ linearly independent rows. This implies $\bar{A}$ must have $n + m - \bar{q}$ linearly independent rows, and in virtue of the rank-nullity theorem,

$$\dim\left(\text{null}(\bar{A})\right) = n - \text{rank}(\bar{A}) = \bar{q} - m \leq q - m$$

$\qquad\square$

Lemma 3 is then obtained as a direct corollary of Lemmas 4 and 5, since these two prove $k-1 \leq q-m$, and $k-1$ is exactly the number of fractional values of $\bar{z}$.