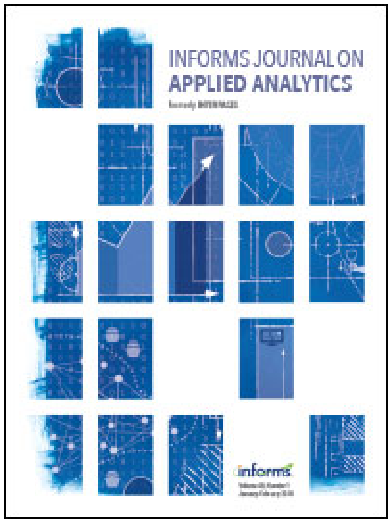

INFORMS Journal on Applied Analytics

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Barrick's Turquoise Ridge Gold Mine Optimizes Underground Production Scheduling Operations

Andrea Brickey, Akshay Chowdu, Alexandra Newman, Marcos Goycoolea, Raphael Godard

informs
http://pubsonline.informs.org/journal/inte

INFORMS JOURNAL ON APPLIED ANALYTICS
Articles in Advance, pp. 1–13
ISSN 0092-2102 (print), ISSN 1526-551X (online)

Barrick's Turquoise Ridge Gold Mine Optimizes Underground Production Scheduling Operations

Andrea Brickey,[a] Akshay Chowdu,[a] Alexandra Newman,[b] Marcos Goycoolea,[c] Raphael Godard[d]

[a] Department of Mining Engineering and Management, South Dakota School of Mines and Technology, Rapid City, South Dakota 57701; [b] Department of Mechanical Engineering, Colorado School of Mines, Golden, Colorado 80401; [c] School of Business, Universidad Adolfo Ibáñez, Peñalolén, Santiago, Chile 7941169; [d] Barrick Gold Corporation, Toronto, Ontario M5J 2S1, Canada

Contact: andrea.brickey@sdsmt.edu, https://orcid.org/0000-0001-6178-5233 (AB); akshay.chowdu@mines.sdsmt.edu, https://orcid.org/0000-0001-5567-5222 (AC); anewman@mines.edu, https://orcid.org/0000-0001-9886-9721 (AN); marcos.goycoolea@uai.cl (MG); rgodard@barrick.com (RG)

Received: November 22, 2018
Revised: March 30, 2019
Accepted: September 11, 2019
Published Online in Articles in Advance: July 24, 2020

https://doi.org/10.1287/inte.2020.1027

Abstract. Mining operations determine a long-term production schedule, often to maximize net present value. For a time horizon of between years and decades, optimization models seek the extraction times—with monthly or yearly fidelity—of three-dimensional, notional blocks of ore and waste within a deposit to satisfy spatial precedence constraints, as well as resource constraints on the amount of material extracted and sent to the mill. With algorithmic advances, as well as those in mine planning software and in hardware, we are able to solve instances with a decade-long horizon at daily fidelity. The resulting objective, repeatable, and defensible schedules inform production and maintenance supervisory decisions based on resource availability, that is, loaders, shovels, haul trucks, and mineral processors. We implement our solutions at the Turquoise Ridge underground gold mine in Nevada, United States. These solutions indicate more than a 2% increase in total ounces extracted over a decade while decreasing development footage by as much as 11% over the same time horizon. Furthermore, we are able to incorporate rules governing a shared resource and to evaluate binding versus nonbinding capacity constraints.

History: This paper was refereed.
Funding: This work was supported by Barrick Gold Corporation [Contract #18-12 with South Dakota School of Mines (A. Brickey)].

Keywords: production scheduling • mine planning • underground mining • integer programming applications

Barrick Gold Corporation is based in Toronto, Canada, and has been publicly traded since 1983. Its first mines were located in Ontario and Quebec, Canada, and in Nevada and Utah, United States. At the time of this writing, Barrick's mines and projects are located in Argentina, Australia, Canada, Chile, the Dominican Republic, Papua New Guinea, Peru, Saudi Arabia, the United States, and Zambia, with more than three-quarters of its production originating in the Americas. As measured by annual production tonnage in 2017, Barrick is the largest international gold mining company.

Turquoise Ridge (TR) underground mine is a division of Barrick Gold Corporation, which, in aggregate, employs about 3,800 workers. The mine is located approximately 44 miles northeast of Winnemucca, Nevada. Barrick and Newmont Mining Corporation co-own the mine in proportions of 75% and 25%, respectively. The average reserve grade of the ore is 0.45 troy ounces per ton (15.56 grams per tonne), the highest in Barrick's operating portfolio, and among the highest in the industry.

Problem Summary

Operational decisions in underground mine production scheduling consist of determining when various activities should occur over the course of a horizon of one to many years. These activities prepare and, subsequently, retire an area from which ore is extracted, and consist of (i) drilling the area and inserting explosives, (ii) blasting the rock (or fragmenting it via mechanical means) to permit it to be extracted, (iii) extracting (or mucking) the broken rock, and loading it into a haul truck, and (iv) transporting it to the surface for further processing and eventual sale. If necessary, for stability, the void left as a result of the extraction is backfilled. A precious metal mining company generally possesses an objective of maximizing net present value (NPV), and constraints consist of rules governing the order in which activities can occur, and the amount of resources over a time span that can be used to conduct these activities, for example, extraction or processing capacity. (For technical terms such as NPV, we refer the reader to Table A.1 in Appendix A.)

Literature Review

There is extensive literature on mathematical programming approaches to production scheduling in both open pit and underground mining applications (Newman et al. 2010), with some of the earlier work applied to the former (Johnson 1968), whereas the latter began with models such as that proposed by Trout (1995). Research on underground mine planning methods has henceforth been applied to several types of mines, including hard-rock (metal) mines containing platinum and palladium (Carlyle and Eaves 2001), iron ore (Kuchta et al. 2004, Martinez and Newman 2011), and lead and zinc (O'Sullivan and Newman 2014), and to soft rock (i.e., coal) mines (Sarin and West-Hansen 2005), all of which have different requirements than the mine we consider here. In particular, the geology specific to each mine lends itself to a given mining method. This difference results primarily in unique precedence constraints, mathematical structures, and, correspondingly, solution techniques. In some cases, the authors simply solve the monolith, whereas in others, they use tailored decompositions.

Recently, Brickey (2015) and Muñoz et al. (2018) have proposed exploiting the so-called resource-constrained project scheduling problem (RCPSP) structure to solve large-scale underground mine planning problems. Their approach is based on extending a linear programming decomposition algorithm (Bienstock and Zuckerberg 2010), and an integer rounding heuristic (Chicoisne et al. 2012), in the context of open pit mining. King et al. (2017) discuss the use of this approach for strategic underground planning, and King et al. (2016) employs it to solve instances of open-pit-to-underground transition problems.

Here, we adopt the same approach to solve a challenging problem faced by Barrick. Although we tailor our model to the company's mine and underground mining method, this type of modeling can be applied, regardless of ore type, to any underground mine as long as the data (as given by the sets and parameters listed in Appendix B) are known. Most open pit operations include the option to stockpile, which, typical of underground mines, we do not. (We refer the interested reader to O'Sullivan et al. 2015 for the differences between open pit and underground mine planning optimization models.) We determine production schedules for long (for example, 10-year) horizons at fine (for example, daily) fidelity. That is, we provide our industry partners with schedules that dictate the activities that start each day for the next 10 years, and that can be used for both mid- and long-range planning.

Operations at Turquoise Ridge

Because Turquoise Ridge (TR) is an underground mine, many types of activities must be scheduled to prepare an area for ore extraction and, subsequently, to backfill the void to mitigate the instability of the host rock. For example, to access a mining zone, we must first excavate long-term, or primary, development areas such as haulage ramps and declines. Within each mining zone, we next build secondary development needed for access to the individual stopes. Primary development is generally a capitalized expenditure, whereas secondary development is considered an operational expense. In TR's case, this primary and secondary development is normally excavated in waste, or lower-grade ore, so as not to inhibit the extraction of higher-grade ore while still satisfying geotechnical and safety requirements.

These activities all require different equipment, and are therefore associated with different resource requirements. For example, development drilling activities use a jumbo, that is, a piece of equipment that drills horizontally. A load-haul-dump unit (Figure 1) is needed for mucking and to place the rock into trucks, which then haul the excavated rock to a muck pass or loading station. In 2017, the mine began using Sandvik MH620 continuous miners (Figure 2) that

Figure 1. (Color online) Photograph of a Load-Haul-Dump Unit (Caterpillar for Hardrock) Representing the Type Used at Turquoise Ridge

Figure 2. (Color online) Photograph of a Mechanical Miner (MH620 Roadheader for Hardrock) Representing the Type Used at Turquoise Ridge

possess rotating drums on a hydraulic arm that cut away at the rock, thus eliminating the need for drilling and blasting (Moore 2017). The cut material is gathered at the front of the continuous miner and then loaded into trucks via a conveyor belt.

Durations of the individual activities vary widely, from a few hours to multiple months. This characteristic is typical for an underground mine, in which exploration, development, extraction, and backfilling are often associated with areas sized to match the nonuniform geology. (See, as another example, O'Sullivan and Newman 2014.) Each activity is assigned a profit or cost based on the activity. For example, the extraction of stopes, or areas of ore, ultimately yields salable gold, that is, revenue, whereas development and backfill activities incur costs. Activities associated with a cost are often the activities whose execution is needed to provide access to a given area such that ore can subsequently be extracted. We treat these data as deterministic for the purposes of our long-term plan and owing to the culture of and data availability at the mine site. Short-term models that might capture the stochastic nature of the operations are beyond the scope of our work.

The Turquoise Ridge mine extracts ore utilizing the underhand cut-and-fill mining method (Figure 3), a top-down method, meaning that extraction progresses downward by mining small cuts and then backfilling them before moving to a cut either adjacent to, or beneath, the current cut. The backfill is a cemented aggregate limestone, which solidifies to provide roof support, necessary for a safe working environment. The precedence structure is dictated by the nature of the extraction above a given area. An intralevel ramp developed in the waste rock provides access to the ore in a topcut, which, once extracted, is backfilled. The ore contained in the ramp is then extracted to reach the undercut directly below, in which the ore is mined out and backfilled. This process continues until all regions associated with this ramp have been mined and backfilled. Lastly, the ramp is backfilled as well. This method possesses a lower production rate than other underground bulk mining methods, but is employed specifically because of the difficult ground conditions (that is, weak host rock) found at Turquoise Ridge.

To determine where and when to mine, engineers continually evaluate mining conditions (such as the strength of the host rock), economic factors (such as the price of gold and the cost of operating or upgrading equipment), and geological factors (such as updated information about the amount of ore contained in the rock) to determine if the mine's production schedule needs to be revised. Associated with these activities is a mining sequence, or a set of rules dictating, to some extent, the precedence that must be followed during activity execution. It is with this fixed design (Sipeki et al. 2020) and set of precedences that we generate a production schedule for the requisite mining zones, that is, designated areas within the larger operation (Figure 4).

Specifically, we seek to maximize discounted NPV subject to the following constraints: (i) precedence between activities, (ii) daily mining capacity for the ore tonnage, (iii) daily capacity of total material movement (waste, ore, and backfill tonnage), (iv) daily number of mechanical mining activities executed, (v) daily mechanical mining advance, (vi) monthly ore tonnage mined, and (vii) other common-sense operational restrictions, for example, an activity can be executed at most once during the planning horizon. In addition to these operational constraints, a small set of activities must start on specific dates or be completed by a specific date, sometimes as a result of our initial conditions.

Previous and New Scheduling Methods

Prior to our work, the scheduling process consisted of the following three steps: mine design, task creation, and scheduling. Mine design incorporates numerous factors such as ground conditions, ore grades, safe mining sequences, and equipment characteristics, and yields an economic extraction envelope. A design produces activities based on characteristics such as location, size, resource quantity, and extraction method (conventional or mechanical), to which are assigned precedence requirements to ensure a safe and logical mining sequence, for example, that access to an extraction activity is created before said extraction, or that ventilation infrastructure is built prior to operating in an area. At the time of this writing, the mine engineers at TR used a commercially available heuristic-based scheduling program. The software employs a genetic algorithm to

Figure 3. (Color online) Mining Sequence in the Underhand Cut-and-Fill Method

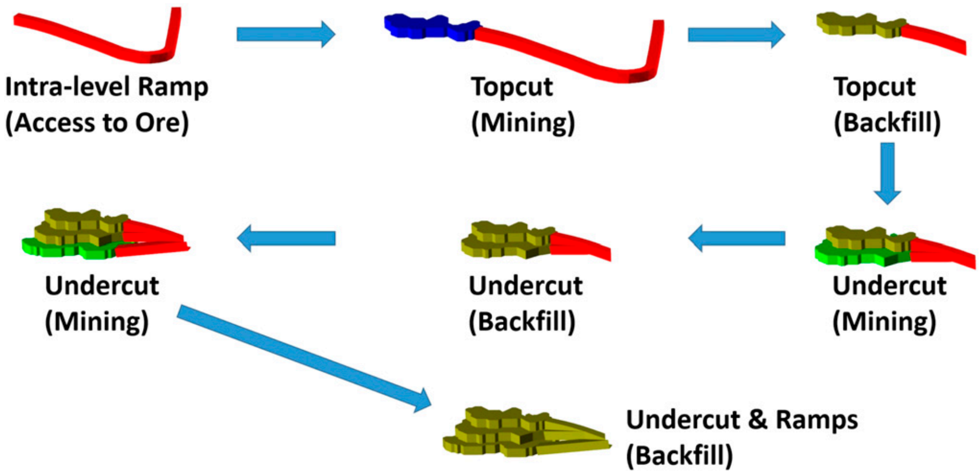

Note. This sequence consists of an intralevel ramp that provides access to the ore, followed by mining and backfilling, where mining always proceeds top-down, that is, the topcut is excavated first, followed by the undercuts.

schedule the activities with a view to improving NPV while adhering to the precedence constraints, and to various resource constraints. The disadvantage of the heuristic-based method is that the quality of the solution is not guaranteed to be within a known percentage of optimality. To compensate for the unknown solution quality, mine planners generate multiple life-of-mine solutions and then evaluate them, choosing the one associated with the highest NPV or with a specific production profile.

The gold mining industry is highly competitive. Price fluctuations and depleting reserves have highlighted the need to continually improve efficiency and reduce costs. Barrick aims "to be the leading mining company focused on gold" (Barrick Gold Corporation 2017, p. 3), and the company is therefore motivated to implement new optimization techniques into their business strategies. With this backdrop, our goal is to improve the production scheduling process by generating provably (near-)optimal solutions more quickly, enabling mine planners to run multiple scenarios depending on the conditions. In fact, due to the mathematical structure of the underlying scheduling models, and the efficiency of the corresponding algorithms, we are able to exceed the planners' expectations in terms of the size of instance solved within an acceptable timeframe.

To create a representative optimization model compatible with the mathematical optimization software, we employ (the Open Mine Planner (OMP) solver; see Appendix B), it is necessary to determine how information such as the mining method and sequence, the equipment used, and any operational limitations are represented in the mine plan. TR utilizes the Deswik suite (Deswik Mining Consultants (Australia) Pty Ltd 2018) of mine planning tools, a commercially available platform, to create a three-dimensional computer-aided design representation of the mine and subsequently to define operational activities.

Figure 4. (Color online) Mining Zones of Interest at the Turquoise Ridge Mine

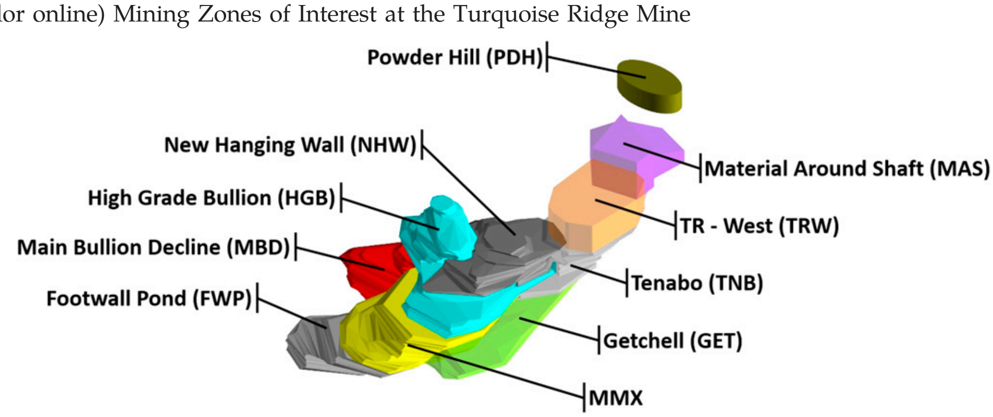

Figure 5. (Color online) Distribution of Activity Durations Before (Left) and After (Right) Splitting

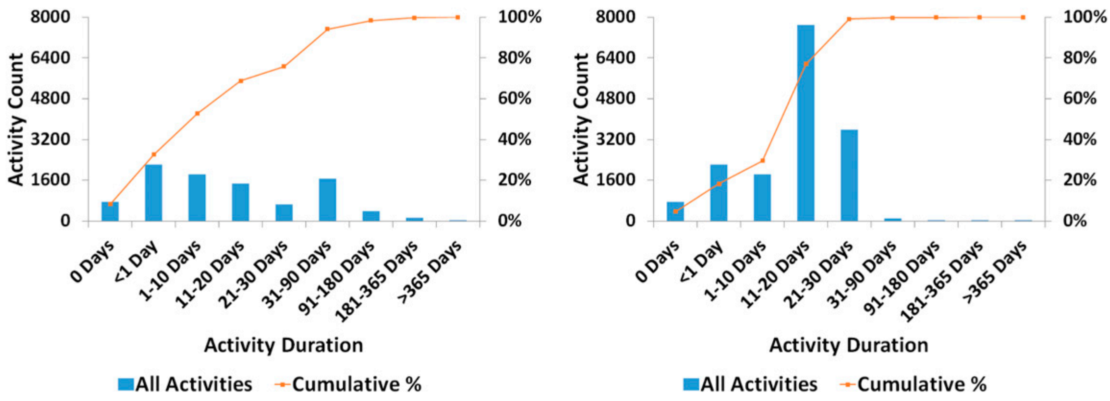

Note. As the graphs show, many more activities require fewer than 31 days to execute after splitting.

We extract the necessary data from the Deswik files to create our optimization model. At a high level, the resulting model can be thought of as consisting of the following components:

1. Parameter data: The required information from Deswik pertaining to each activity, in addition to the time horizon and fidelity.

2. Objective function: The value assigned to each activity or task, which appears as a coefficient in our objective of maximizing NPV.

3. Precedence links: Activity dependence that creates a sequence.

4. Resource constraints: The operational limitations such as the amount of material moved in a day, or the number of activities that can simultaneously be executed on a level (see Table 1).

Each activity is associated with a net profit, which can be positive, zero, or negative, and with an activity type, which dictates the resource(s) that it consumes. Table 1 lists the various resource constraints and the related activity types. Table 2 provides descriptions for the types of activities scheduled. For example, ore-producing activities such as topcut mining, undercut mining, and stope mining are associated with ore tonnage. Therefore, the resource consumptions resulting from executing these activities are summed in the daily ore tonnage constraint.

To enhance model tractability, we excise activities based on the following: (i) whether they were already scheduled in the short term, that is, before January 2019 and therefore fixed, (ii) the length of the critical paths of predecessors based on the durations and resources these predecessors would require to be completed, and (iii) the length of the planning horizon. The optimization model considers a total of 16,432 activities, but then excludes those activities scheduled outside of the time horizon. Additionally, we include 25 activities that produce ore tons scheduled in 2018 but that are to be completed in 2019. We concede that these legacy activities constituting initial conditions compromise some optimality, but to the benefit of creating a seamless plan with which the mine can dovetail its current operations.

To rectify the heterogeneity in the duration of the underground activities, it is necessary to schedule at a fidelity commensurate with the duration of shorter activities while limiting the duration of longer-lasting ones. For example, scheduling at an hourly level of detail for the planning horizon of interest would yield an intractable model. We mitigate this by rounding

Table 1. Resource Constraints in Our Production Scheduling Model

Constraint name	Description	Activity types	Fidelity
Total tonnage	Total amount of material (i.e., waste, ore, and backfill) moved	Backfilling, capitalized projects, primary development, secondary development, stope mining, topcut mining, undercut mining, vertical development	Daily
Daily ore tonnage	Amount of ore tons moved daily	Topcut mining, undercut mining, stope mining	Daily
Mechanical mining activity	Number of concurrent mechanical mining activities	Topcut mining, undercut mining (only activities using mechanical miners)	Daily
Mechanical mining footage	Mechanical mining advance	Topcut mining, undercut mining (only activities mechanical miners)	Daily
Monthly ore tonnage	Amount of ore tons moved monthly	Topcut mining, undercut mining, stope mining	Monthly

Note. Each of the constraints in our model is enforced using a specific time fidelity.

Table 2. Activity Types in the Production Schedule at Turquoise Ridge

Activity type	Description
Backfilling	Filling of voids created by ore extraction with cemented aggregates to maintain geotechnical stability
Delineation drilling	Drilling performed to delineate ore and waste boundaries for smaller sections of the deposit, typically performed a few weeks before actual extraction
Capitalized projects	Large-scale undertakings that would be considered an asset for the mine, for example, sinking shafts, installing or upgrading main ventilation fans
Primary development	Horizontal excavations that provide initial access to the ore body, for example, haulage ramps, declines (Figure 4)
Stope mining	Extraction of ore in a defined area, creating a void
Topcut mining	Extraction of ore in a defined topcut (Figure 3)
Secondary development	Horizontal excavations needed to provide more granular access to the ore body, for example, access ramps, conducted prior to extraction activities
Undercut mining	Extraction of ore in a defined undercut (Figure 3)
Vertical development	Vertical excavations, for example, ventilation raises, service shafts

activity durations up to a day, and, correspondingly, by splitting longer-lasting activities into those requiring no more than 31 days to complete. In this way, we can use a daily time fidelity to schedule, while also forcing a greater degree of homogeneity into activity duration (see Figure 5). Although this may seem overly detailed for a 10-year horizon, many activities require a day or less. These cannot be aggregated without loss of desired detail. Our schedules are therefore intended to provide a longer-term view of the operation, as opposed to a short-term schedule, and are then used as the foundation for near-term planning. From an implementation perspective, mining engineers typically conduct long- to mid-term scheduling events once or twice per year. This may occur more frequently at larger operations depending on the variability of the mining conditions.

For this project, the components necessary to run the OMP solver are represented by six files (Table 3) that together contain the necessary information for the software to build an instance and correspondingly solve it. Specifically, the Deswik model data are exported to a Comma Separated Values file and handed to a Python script in which any required transformations occur, for example, the removal of any tasks that are scheduled prior to the start of our time horizon.

The associated output files are then passed to the solver and a run is initiated. After completion, the solver writes out the integer programming solution. In OMP, time periods are represented as integer values. These are converted into calendar dates, using another script, before the schedule is imported into Deswik, whereupon it can be viewed as a Gantt chart and the results can be visually verified, for example, whether any precedence constraints are being violated and whether daily production limits are being respected. Reports on production metrics and economics can then be generated on a daily, monthly, or yearly basis. Figure 6 provides a schematic of this process.

Fortunately, despite the long horizon (15 years) and fine time fidelity (daily), we are able to obtain reasonable solution quality and times because of the problem structure we can exploit, which yields a 1% optimality gap in approximately eight hours. We use the first 10 years to construct the schedule, resulting in 9,150 scheduled activities out of the approximately 14,000 available over the 15-year horizon. With a mine life at the time of this writing estimated at 22 years, mine planners deem the 10-year horizon to be sufficiently long for their intermediate-term goals.

We provide the formulation of the integer program, the hardware and software on which we run

Table 3. Input File Types and Descriptions

File type	Description
Blocks (*.blocks)	Lists all activities and their associated resource consumption values on a daily basis
Problem (*.prob)	Controls how the overall model is constructed, and contains the time horizon for which to solve the model, the number of constraints, and associated parameter values
Precedence (*.prec)	Lists the predecessors for each activity
Delay (*.delay)	Lists the delay corresponding to each precedence list for each activity
Mapp (*.mapp)	Contains the mapping from the Deswik identifier to the corresponding identifier used by the mathematical optimization software
Time (*.time)	Implements multi-time period constraints

Figure 6. (Color online) OMP Scheduling Process

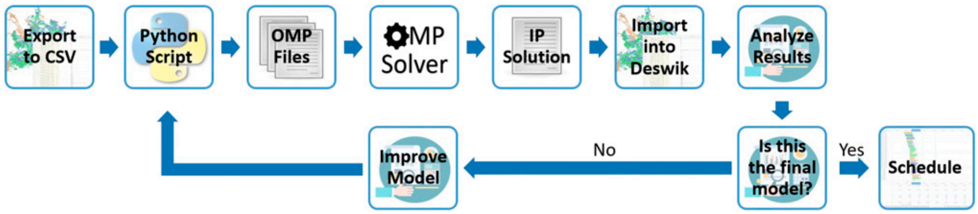

Note. This process consists of data processing, using the solver to produce a schedule, employing mining software to analyze the solution, and outputting the schedule after no adjustments need to be made.

the model instances, and specific information about solution times and optimality gaps in Appendix B.

Results

Barrick's primary goal in enlisting our help to produce production schedules is to mitigate the myopia seen in their current, heuristically generated ones. We are able to incorporate rules whose benefits were not initially apparent. Our model produces an objective, repeatable, and defensible schedule in less time than currently required. Finally, the company is interested in evaluating operational resource utilization, that is, binding versus nonbinding constraints.

A challenge to this operation is the presence of an annual ore production target, rather than a limit, signifying that the mine must meet this value to the extent possible over the course of a year (and cannot exceed it). The origin of this so-called target is the fact that TR shares a processing facility with neighboring mining operations, for which each mine is expected to produce a given amount of material to meet processing facility capacity while not overtaxing it and precluding its use by the neighbors. We could place a constraint on daily ore production such that the total for a year equals the annual ore target. However, it is unclear how to best translate an annual target into a daily one. Daily targets equal to the quotient of the annual level and the number of operational mill days could be overly restrictive during some time periods of the year and therefore lead to lost opportunity for TR with respect to ore production. On the other hand, some relaxation of the daily ore production capacity constraint could cause front-loading, in which the economic driver (in our case, ore production) is accelerated in the early part of the time horizon to increase NPV. Such a schedule would be infeasible from an operational standpoint and would create discord with the neighboring mines. To meet the annual ore production target, to preclude lost opportunity, and to prevent front-loading, we introduce a multi-time period knapsack constraint (Chowdu et al. 2019) that constrains each month's production to be approximately equal to the quotient of the annual production target and 12 (i.e., the number of months in a year). Monthly production targets are slightly higher in aggregate than the annual target to allow for the lumpiness associated with the tonnage contained in different stopes, and to account for the differing number of days per month. Figure 7 provides a depiction of the impacts of front-loading and demonstrates that the use of multi-time period knapsacks curb that behavior.

Our schedule, to which we refer as the underground RCPSP (UG-RCPSP), results in a production profile similar to the original, manually generated one, but tends to forego unprofitable development activities in the near term to a greater extent than the original schedule. The (UG-RCPSP) schedule improves NPV, not only directly by bringing ounces forward from areas with higher grade, but also, secondarily, by reducing the development (scheduled just-in-time relative to the corresponding extraction). Production is moved from areas with higher development to those that require less to mine the associated ore. Table 4 provides quantitative comparisons of total ore tonnage and ounces of gold mined, backfill used, waste extracted, tonnage associated with primary and secondary development, and advancement achieved with the original schedule and for (UG-RCPSP).

Figure 8 contrasts the scaled cumulative NPV over the planning horizon of interest, on an annual basis, for Barrick's original schedule (original) and for the one we generate (UG-RCPSP). The latter yields a higher NPV from the beginning of the horizon, and the difference in NPV grows for about the first six years, after which it remains virtually constant based on an annual discount rate set by the company.

The NPV is inextricably intertwined with the total ounces of gold produced, and with their corresponding average grade (Figure 9(a)). The (UG-RCPSP) schedule consistently produces more ounces for the first 8 years of the 10-year schedule with the exception of year 6, in which (UG-RCPSP) yields a slightly lower value. Nonetheless, owing both to the overall number of ounces produced and to the discount rate, the generally higher level of gold production matches the higher NPV we

Figure 7. Impacts of Front-Loading and Use of Multi-time Period Knapsacks to Curb the Behavior

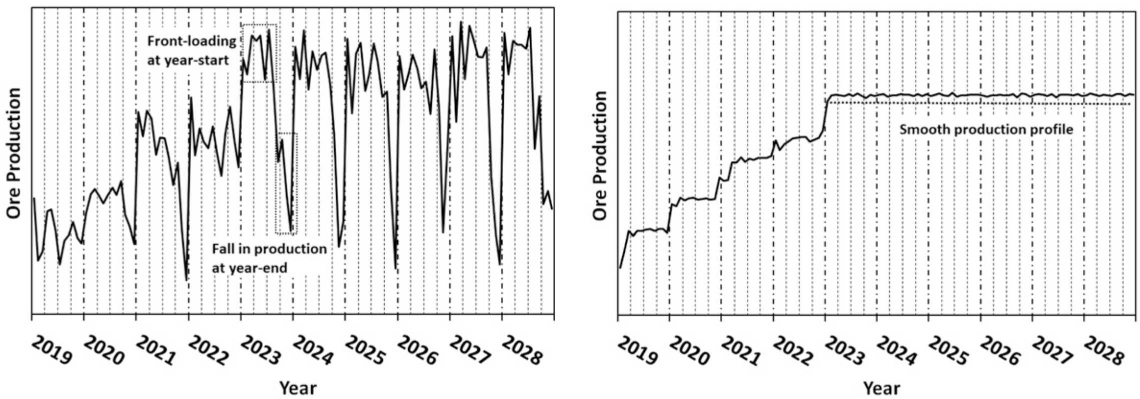

Front-loaded schedule   Schedule with multi-time period knapsacks

Note. The left-hand graph shows an operationally infeasible front-loaded schedule, whereas the right-hand graph demonstrates how the appropriate constraints smooth out ore production (and, hence, mill utilization).

obtain with our schedule. Commensurately, although less monotonic and more difficult to determine dominance between schedules, the overall much higher grade of the mined material in 3 of the first 4 years (and in 6 years of the 10-year horizon) is also consistent with the higher NPV (UG-RCPSP) produces. Barrick's original schedule generates higher-grade ore in the final 2 years of the horizon (at the expense of the same or lower grades throughout all but one of the first 8 years of the horizon) when compared with the optimized schedule. This myopia in the former schedule fails to capture the importance of the discount factor, and, therefore, results in lower NPV.

Activities, including primary and secondary development, backfill, and those associated with other necessary infrastructure, are expensive to conduct and do not directly generate revenue, and (UG-RCPSP) schedules these just-in-time (subject, naturally, to precedence constraints) to minimize the detrimental impact on the objective function value. Figure 9(b) shows a substantial reduction in development footage in the early years of the schedule. Specifically, primary development is reduced by 3.3% and secondary development is reduced by 10.2%, contributing to an overall reduction of 8.8%. More tons of backfill are scheduled in the early years due to the increased ounces extracted, yet the overall amount of backfill required for both schedules is the same.

Finally, although a critical path in the traditional sense is generally based on due dates, a characteristic absent from our production schedules in which activities need not even be executed at all, our model does provide an evaluation of two types of phenomena that restrict the objective function value: (i) bottleneck time periods in which one or more

Table 4. Comparison of Total Ore Tons and Ounces Produced and Total Waste Tons and Development for the OMP-Generated Schedule Compared to the Original

Metric	Difference (%)
Total ore tons	1.0
Total ounces	2.4
Total backfill tons	0.0
Total waste tons	–11.0
Primary development tons	–3.3
Primary development footage	–3.3
Secondary development tons	–10.2
Secondary development footage	–10.2

Note. The OMP-generated schedule improves upon the original schedule in the total ore tons and ounces produced (both increase), as well as total waste tons and development (both decrease).

Figure 8. (Color online) Comparison of the Scaled, Cumulative NPV and the Mine's Original Schedule

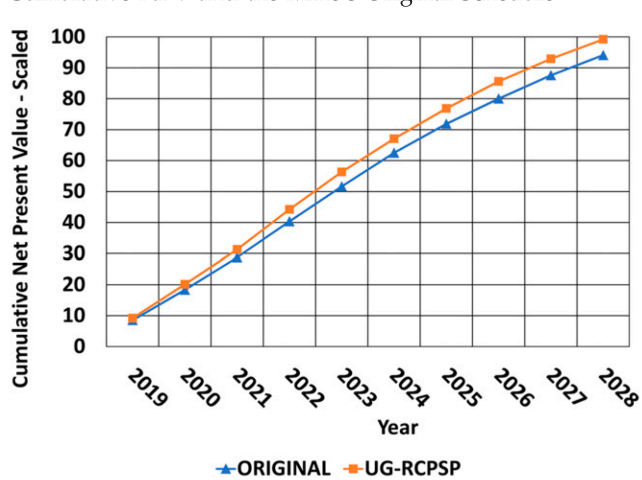

Note. The scaled, cumulative NPV is higher for the model we solve than that corresponding to the mine's original, manual schedule.

Figure 9. (Color online) Comparison of Select Production and Operational Metrics from the (UG-RCPSP) Schedule and the Original Schedule

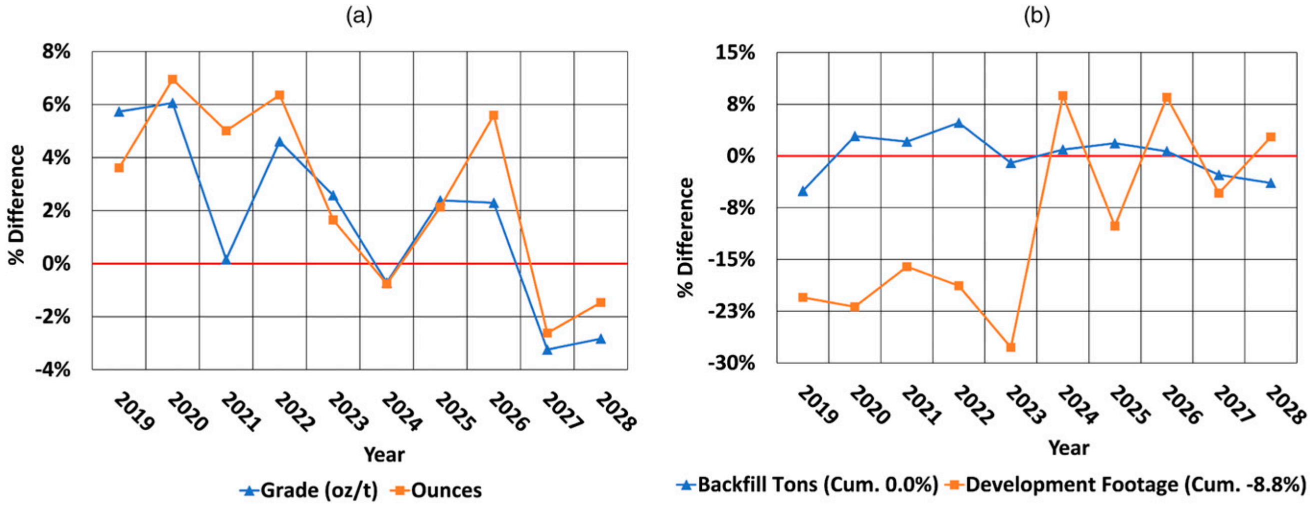

Percentage Difference in Ounces and Average Annual Grade — Percentage Difference in Backfill Tons and Development Footage

Note. The comparison shows that more, higher-grade ounces are produced in the (UG-RCPSP) schedule with about the same number of backfill tons and less development (i.e., with fewer resources overall).

resources are used at capacity, and (ii) activities whose start dates are delayed owing to predecessor activities and their associated durations. The use of Deswik or other mine schedule visualization tools can discern these restrictions and provide engineers and management with feedback regarding appropriate operational responses to changing conditions. Engineers can then determine if additional resources are warranted or if modifications to the design or other fundamental inputs to the scheduling model should be made.

Conclusion

Barrick Gold Corporation approached us with a production scheduling problem that they were handling well, but upon which they wanted to improve. Specifically, they were interested in (i) using a formal optimization approach to increase NPV or demonstrating that the company had a (near-)optimal solution, and (ii) creating production schedules faster, more objectively, and in a reproducible manner. We were able to show that the company was, indeed, generating good schedules, but that we could improve on them both in terms of quality, and, to a much greater extent, in the time required to produce them. Our deterministic-model solutions with daily-fidelity schedules are capable of meeting the annual ore production target, in addition to other resource constraints. The model accommodates heterogeneous activity durations and produces multiyear schedules with reasonable computation times.

Important aspects of our work arose with each iteration of the model, providing us with information about the requirements for producing an operationally feasible schedule. For example, due to the highly precedence-constrained nature of many activities, we are able to run instances containing longer time horizons than initially thought, thereby yielding solutions with a long-term outlook and ensuring that the schedule accounts for a perspective of the operation that includes access development, equipment allocation (that is, mechanical miners), and backfill activities needed for safe execution. Additionally, the schedules are able to provide engineers with a tool to identify deficiencies in resources.

A valued outcome of the model is the significant reduction in development activities early in the time horizon, yielding just-in-time scheduling of expensive activities. The resulting cost reductions, coupled with an increase in the grade of the scheduled ore production, significantly impacts the NPV of the operation. Although the mine design is fixed for any given schedule, the results provide valuable information on critical activities and priority development for higher grade areas. Another important lesson relates to the production profile, initially constrained on an annual basis, but which resulted in front-loaded, operationally infeasible schedules. Incorporating monthly, rather than annual, production targets increases NPV, reduces development costs, and produces more objective and repeatable schedules, with associated parametric analyses. Therefore, the company has begun to think about implementing models such as the one we describe at its other underground mine sites.

Finally, our modeling paradigm of the resource-constrained production schedule problem, and its associated solution procedure, can be applied in a variety of other settings such as nurse staffing and scheduling, supplier selection, pharmaceutical research and development, plant engineering and construction, information systems, and automotive research and development projects (Schwindt and Zimmermann 2015).

Appendix A. Mining-Specific Acronyms

Table A.1 provides a list of mining-specific acronyms and terms to facilitate parsing the paper for those unfamiliar with underground mining.

Table A.1. Mining Terminology Used in This Paper

Term	Description	Source
Backfill	Material used to fill a void created by ore extraction, for example, pastefill (processing tailings mixed with water to form a slurry), rock backfill (waste rock from development sometimes combined with cement)	Hambley (2011)
Decline	A system of ramps and horizontal drives that connects the orebody to a surface portal or to a breakout from existing mine infrastructure	Brazil et al. (2008)
Deswik	Mine planning and scheduling software	Deswik Mining Consultants (Australia) Pty Ltd (2018)
Development	An excavation that provides access to the orebody, categorized as primary, secondary, or vertical	Tuck (2011)
Gold ounces	The final product derived through mineral processing of ore, taking into account processing recovery and dilution	Fuerstenau and Han (2011)
Net present value (NPV)	The difference between the present value of cash inflows and cash outflows over a period of time, used to analyze the profitability of an investment or project	Stermole and Stermole (2014)
OMP solver	Open Mine Planner, a mathematical solver used for large scheduling problems with a particular structure	Moreno et al. (2017)
Ore	An occurrence of rock that contains sufficient minerals to be economically extracted from the deposit	Darling (2011)
RCPSP	Resource-constrained production scheduling problem	Brucker and Knust (2012)
Stope	An excavation in the orebody created to extract ore, resulting in a stable void	Carter (2011)
Topcut	An excavation in the underhand cut-and-fill mining method, in which the excavation's roof is composed of the host rock	Hustrulid and Bullock (2001)
TopoSort	Heuristic employed by OMP to convert a solution to the linear programming relaxation of the scheduling problem into an integer one	Chicoisne et al. (2012)
(UG-RCPSP)	Underground resource-constrained production scheduling problem	King et al. (2017)
Undercut	An excavation in the underhand cut-and-fill mining method, in which the excavation's roof is composed of cemented backfill, formed by backfilling a preceding topcut or undercut	Hustrulid and Bullock (2001)

Appendix B. Formulation

Our underground mine scheduling problem possesses a RCPSP mathematical structure, in which the two main categories of constraints are (i) precedence (the majority), and (ii) knapsacks (the minority). Problems with this structure lend themselves to the use of the OMP solver (Rivera et al. 2015), which implements a number of algorithms published in the academic literature, and executes primarily in two stages: (i) it solves the linear programming relaxation of the problem using a decomposition method (see Muñoz et al. 2018 for details), and then (ii) it applies a TopoSort heuristic to obtain an integer-feasible solution (Chicoisne et al. 2012, Brickey 2015). The latter, although a list-ordering heuristic based on the "expected" completion time of a block or activity in the linear programming solution, tends to provide integer-feasible solutions provably within a few percentage points of optimality. We employ OMP's variant of an early start algorithm (Lambert et al. 2014) to reduce the number of activity-time period pairs under consideration.

We formulate this model by converting Deswik data, via a Python script, into a format compatible with OMP. Barrick's data set contains activities with time-varying durations. Because OMP lacks the flexibility to accommodate this type of data, we approximate the solutions by using the longest (that is, most conservative) duration when optimizing and then post-process into the final schedule the correct duration based on the time at which the activity is executed. We run the corresponding instances using a Lenovo ThinkServer RD350 computer with 16 processors (2.6 GHz) and 32 gigabytes of RAM. All runs are executed with OMP Version 2663. Instances spanning a 15-year horizon and possessing daily fidelity result in 47,410,701 variables, 102,548,859 precedence constraints and 22,102 resource constraints, and require 29,107 seconds (or 8.09 hours) to solve to 1% optimality.

The model formulation that OMP solves is as follows:

Indices and sets:

$a \in \mathcal{A}$ An activity within the set of all activities

$\tilde{a} \in \tilde{\mathcal{A}} \subset \mathcal{A}$ An activity within the set of activities whose start dates have been predetermined

$\bar{a} \in \bar{\mathcal{A}}_a$ An activity $\bar{a}$ within the set of predecessor activities to activity a

$r \in \mathcal{R}$ A resource within the set of resources, such as production and development capacity, whose limits are enforced on a daily basis

$r \in \widehat{\mathcal{R}} \subset \mathcal{R}$ A resource within the set of resources, such as production and development capacity, whose limits are enforced on a monthly basis

$t \in \mathcal{T}$ A day within the set of daily time periods

$m \in \mathcal{M}$ A month within the set of monthly time periods

$t \in \hat{\mathcal{T}}_m$ A day within the set of days contained in month m

Parameters:

c_a Monetary value associated with completing activity a [$]

q_{ra} Daily consumption of resource r associated with completing activity a [tonnes, meters]

$\bar{r}_{rt}$ Maximum amount of resource r available on day t [tonnes, meters]

$\hat{r}_{rm}$ Maximum amount of resource r available in month m [tonnes, meters]

d_a Duration of activity a [days]

$d_{\bar{a}}$ Duration (including mandatory delay) of activity $\bar{a}$ [days]

δ_t Discount factor for period t [fraction]

Decision variables:

X_{at} 1 if activity a is completed by the end of time t, 0 otherwise

$$\text{(UG-RCPSP)} \max \sum_{a\in\mathcal{A}} \sum_{t\in\mathcal{T}} \delta_t c_a (X_{at} - X_{a,t-1}) \tag{B.1a}$$

$$\text{s.t. } X_{a,t-1} \le X_{at} \quad \forall\, a \in \mathcal{A}, t \in \mathcal{T} \tag{B.1b}$$

$$X_{at} \le X_{\bar{a},t-d_{\bar{a}}} \quad \forall\, a \in \mathcal{A}, \bar{a} \in \bar{\mathcal{A}}_a, t \in \mathcal{T} \tag{B.1c}$$

$$\sum_{a\in\mathcal{A}} \frac{q_{ra}}{d_a}\left(X_{at} - X_{a,t-d_a}\right) \le \bar{r}_{rt} \quad \forall\, r \in \mathcal{R}, t \in \mathcal{T} \tag{B.1d}$$

$$\sum_{t\in\hat{\mathcal{T}}_m} \sum_{a\in\mathcal{A}} \frac{q_{ra}}{d_a}\left(X_{at} - X_{a,t-d_a}\right) \le \hat{r}_{rm} \quad \forall\, r \in \widehat{\mathcal{R}}, m \in \mathcal{M} \tag{B.1e}$$

$$X_{\tilde{a}1} = 1 \quad \forall \tilde{a} \in \tilde{\mathcal{A}} \tag{B.1f}$$

$$X_{at} \text{ binary} \quad \forall a \in \mathcal{A}, t \in \mathcal{T} \tag{B.1g}$$

The objective (B.1a) maximizes NPV, which is a discounted function of the monetary value associated with the (on-time) completion of activity a and the time at which said activity is completed. We express the latter as the difference of two variables corresponding to the time by which an activity is completed (Lambert et al. 2014). Constraints (B.1b) ensure that once an activity is completed at time $t-1$, it remains completed for all future time periods $t, \dots, |\mathcal{T}|$. Constraints (B.1c) enforce precedence between an activity a and its predecessors $\bar{a}$, such that a cannot start unless $\bar{a}$ starts sufficiently early that, when accounting for its duration, it is finished by the time a starts. Constraints (B.1d) constitute knapsacks and ensure that the amount of a resource of a particular type consumed by all activities on any given day cannot exceed the availability of said resource. Constraints (B.1e) do the same for a subset of the resources whose consumption must be restricted on a monthly basis. Activities whose start dates have been previously determined to coincide with the beginning of our time horizon must be inserted into the schedule per constraints (B.1f). All variables are required to be binary by constraints (B.1g).

Table B.1. OMP Algorithm Performance as a Function of Time Horizons

Time horizon	Runtime (hours)	Optimality gap (%)
3 years	<0.1	6.2
4 years	0.2	4.5
5 years	0.4	4.2
7 years	1.0	2.9
10 years	3.5	2.1
15 years	8.1	1.7
20 years	14.6	4.8

Note. The OMP algorithm requires more time to execute instances with longer time horizons, though optimality gaps are not commensurately monotonically increasing.

Table B.1 provides OMP solver performance as a function of different time horizon lengths. For a relatively constant gap, runtimes increase exponentially with the number of time periods we include in an instance. The case in the final row in the table corresponds to a life-of-mine instance, and one in which almost all of the approximately 16,000 activities are scheduled. The optimality gaps, which range (generally) between 1% or 5%, are considered acceptable for our type of planning, and result from the termination of the algorithm after executing the TopoSort heuristic to obtain a lower bound and after solving the linear programming relaxation to obtain an upper bound. With such a style of execution, longer runtimes would not close the gap.

References

Barrick Gold Corporation (2017) Operations and technical update. Accessed June 23, 2018, https://www.barrick.com/investors/presentations-reports/event-details/2017/Operations-and-Technical-Update/default.aspx.

Bienstock D, Zuckerberg M (2010) Solving LP relaxations of large-scale precedence constrained problems. Eisenbrand F, Shepherd FB, eds. *Integer Programming and Combinatorial Optimization*, Lecture Notes in Computer Science, vol. 6080 (Springer-Verlag, Berlin, Heidelberg), 1–14.

Brazil M, Grossman PA, Lee DH, Rubinstein JH, Thomas DA, Wormald NC (2008) Decline design in underground mines using constrained path optimisation. *Mining Tech.* 117(2):93–99.

Brickey A (2015) Underground production scheduling optimization with ventilation constraints. PhD thesis, Colorado School of Mines, Golden.

Brucker P, Knust S (2012) *Complex Scheduling* (Springer-Verlag, Berlin, Heidelberg).

Carlyle M, Eaves B (2001) Underground planning at Stillwater Mining Company. *Interfaces* 31(4):50–60.

Carter P (2011) Selection process for hard-rock mining. Darling P, ed. *SME Mining Engineering Handbook*. 3rd ed. (Society for Mining, Metallurgy, and Exploration, Englewood, CO), 357–376.

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

Chowdu AA, Goycoolea M, Brickey AJ (2019) Mine schedule optimization and mine operational realities: Bridging the gap. Mueller C, Assibey-Bonsu W, Baafi E, Dauber C, Doran C, Jasczuk MJ, Nagovitsyn O, eds. *Mining Goes Digital: Proc. 39th Internat. Sympos. Appl. Comput. Oper. Res. Mineral Indust.* (APCOM) (Taylor & Francis, London), 412–418.

Darling P (2011) Mining: Ancient, modern and beyond. Darling P, ed. *SME Mining Engineering Handbook*. 3rd ed. (Society for Mining, Metallurgy, and Exploration Englewood, CO), 3–9.

Deswik Mining Consultants (Australia) Pty Ltd (2018) Deswik mine planning software. Accessed July 12, 2018, https://www.deswik.com/.

Fuerstenau M, Han K (2011) Introduction to mineral processing. Darling P, ed. *SME Mining Engineering Handbook*. 3rd ed. (Society for Mining, Metallurgy, and Exploration, Englewood, CO), 1455–1459.

Hambley D (2011) Backfill mining. Darling P, ed. *SME Mining Engineering Handbook*. 3rd ed. (Society for Mining, Metallurgy, and Exploration, Englewood, CO), 1375–1384.

Hustrulid WA, Bullock RC (2001) *Underground Mining Methods: Engineering Fundamentals and International Case Studies* (Society for Mining, Metallurgy, and Exploration, Englewood, CO).

Johnson T (1968) Optimum open pit mine production scheduling. PhD thesis, University of California, Berkeley.

King B, Goycoolea M, Newman A (2016) Optimizing the open pit-to-underground mining transition. *Eur. J. Oper. Res.* 257(1):297–309.

King B, Goycoolea M, Newman A (2017) New integer programming models for tactical and strategic underground production scheduling. *Mining Engrg.* 69(3):37–42.

Kuchta M, Newman A, Topal E (2004) Implementing a production schedule at LKAB's Kiruna mine. *Interfaces* 34(2):124–134.

Lambert B, Brickey A, Newman A, Eurek K (2014) Open-pit block-sequencing formulations: A tutorial. *Interfaces* 44(2):127–142.

Martinez M, Newman A (2011) A solution approach for optimizing long-and short-term production scheduling at LKAB's Kiruna mine. *Eur. J. Oper. Res.* 211(1):184–197.

Moore P (2017) Turquoise Ridge transitioning to mechanical cutting with Sandvik Roadheader. *Internat. Mining* (November 22), https://im-mining.com/2017/11/22/turquoise-ridge-transitioning-mechanical-cutting-sandvik-roadheader/.

Moreno E, Rezakhah M, Newman A, Ferreira F (2017) Linear models for stockpiling in open-pit mine production scheduling problems. *Eur. J. Oper. Res.* 260(1):212–221.

Muñoz G, Espinoza D, Goycoolea M, Moreno E, Queyranne M, Rivera O (2018) A study of the Bienstock–Zuckerberg algorithm: Applications in mining and resource constrained project scheduling. *Comput. Optim. Appl.* 69(2):501–534.

Newman A, Rubio E, Caro R, Weintraub A, Eurek K (2010) A review of operations research in mine planning. *Interfaces* 40(3):222–245.

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

O'Sullivan D, Brickey A, Newman A (2015) Is openpit production scheduling "easier" than its underground counterpart? *Mining Engrg.* 67(4):68–73.

Rivera O, Goycoolea M, Moreno E, Espinoza D (2015) The OMP Guide. Unpublished software documentation, Universidad Adolfo Ibáñez, Santiago, Chile.

Sarin S, West-Hansen J (2005) The long-term mine production scheduling problem. *IIE Trans.* 37(2):109–121.

Schwindt C, Zimmermann J, eds. (2015) *Handbook on Project Management and Scheduling*, vol. 2 (Springer International Publishing, Cham, Switzerland).

Sipeki L, Newman AM, Yano CA (2020) Selecting support pillars in underground mines with ore veins. *IISE Trans.*, ePub ahead of print January 29, https://doi.org/10.1080/24725854.2019.1699978.

Stermole FJ, Stermole JM (2014) *Economic Evaluation and Investment Decision Methods*. 14th ed. (Investment Evaluations Corp., Lakewood, CO).

Trout LP (1995) Underground mine production scheduling using mixed integer programming. *Proc. 25th Internat. APCOM Sympos.* (Australasian Institute of Mining and Metallurgy, Brisbane, Australia), 395–400.

Tuck M (2011) Underground horizontal and inclined development methods. Darling P, ed. *SME Mining Engineering Handbook*. 3rd ed. (Society for Mining, Metallurgy, and Exploration, Englewood, CO), 1179–1202.

Verification Letter

Trey Williams, Mining Engineer, Turquoise Ridge Joint Venture, MC66 Box 220, 28 Miles NE Golconda, Golconda, Nevada 89414, writes:

"This letter is regarding the submitted paper entitled 'Barrick's Turquoise Ridge Gold Mine Optimizes Underground Production Scheduling Operations.' The long-term planning process at Turquoise Ridge typically consists of using heuristics for scheduling. While these heuristics are good at managing resources and creating feasible schedules, they are often time consuming to produce. This, in addition to the overwhelming complexity of the Turquoise Ridge orebody, means that optimization occurs in very limited ways and there often isn't enough time to do the full optimization that the mathematical method presented in this paper is capable of. Follow-up work using this optimizer has informed the mine's capital development schedule and plans to include it in Turquoise Ridge's long-term planning process in more intimate ways are being considered. The optimized schedules produced have been invaluable in identifying key development time lines and targeting clear value in a highly complex mine."

Andrea Brickey, an associate professor and assistant department head for the mining engineering and management department at the South Dakota School of Mines and Technology, has a PhD in mining and earth systems engineering with a minor in operations research from the Colorado School of Mines. Upon receiving her bachelor of science in mining engineering, she worked in the mining industry for more than 10 years before pursuing a career in academia.

Akshay Chowdu is a doctoral candidate in mining engineering and management at the South Dakota School of Mines and Technology, with a focus on underground production scheduling optimization. He received his master of science and bachelor of science in mining engineering from South Dakota School of Mines and Technology and Indian Institute of Technology (BHU), Varanasi, respectively.

Alexandra Newman is a professor in the mechanical engineering department at the Colorado School of Mines (CSM). She obtained her PhD in industrial engineering and options research at the University of California, Berkeley.

She received a Fulbright Fellowship to work on mining problems at the University of Chile in 2010 and was awarded the INFORMS Prize for the Teaching of OR/MS Practice in 2013.

Marcos Goycoolea is a professor of operations research in the School of Business at Universidad Adolfo Ibáñez, in Santiago, Chile. His primary research interests are in integer programming and natural research management. He is the director of AlicantoLabs, an applied research lab that translates academic research into optimization technologies for strategic mine planning and advises major mining companies in making projects more profitable and sustainable. He holds a PhD in industrial engineering from Georgia Tech.

Raphael Godard is a manager in long-range planning at Barrick Gold Corporation. He is a professional engineer and holds a bachelor of science in mining engineering from McGill University.