Interfaces
Vol. 44, No. 2, March–April 2014, pp. 204–221
ISSN 0092-2102 (print) | ISSN 1526-551X (online)
http://dx.doi.org/10.1287/inte.2013.0730

Extraction and Backfill Scheduling in a Complex Underground Mine

Dónal O'Sullivan, Alexandra Newman
Division of Economics and Business, Colorado School of Mines, Golden, Colorado 80401
{donalos@gmail.com, anewman@mines.edu}

We use mathematical optimization to determine a production schedule (i.e., a schedule of extraction and subsequent filling of the voids) for a complex underground mine in Ireland. The goal of the mining operation is to maximize metal production over the life of the mine, subject to constraints on maximum monthly extraction and backfilling quantities, maximum and minimum monthly metal production, and sequencing between extraction and backfilling operations. We solve our integer programming model with a heuristic to produce a schedule that adds value to the mining operation by (1) shifting metal production forward, (2) reducing waste mining and backfilling delays, (3) avoiding expensive mill-halting drops in ore production, and (4) enabling smoother workforce management.

Keywords: underground mine production scheduling; integer programming applications.
History: This paper was refereed.

During the 1980s, surveys demonstrated that the Lisheen orebody, located in south central Ireland, might be profitable; correspondingly, a drilling program commenced. The seventh hole, drilled in the spring of 1990, revealed nearly 6.5 meters of an orebody containing almost 15 percent sphalerite (zinc) and three percent galena (lead). By the mid-1990s, after 550 holes had been drilled, a 22.5 million-ton deposit, which was comprised of two separate horizontal ore bodies of similarly high-grade zinc and lead content, was defined.

In 1999, the first ore was extracted. Since then, the Lisheen mine has become one of Europe's largest zinc producers; at the time of this writing, over three million tonnes of ore have been extracted and shipped from the mine. As many as 6,300 tonnes of ore are transported to the surface daily. The mine operates six days a week and employs nearly 400 people. Approximately two years of mine life remain, although continued mineral exploration at the fringes of the mine may extend that time.

The two ore bodies at Lisheen are partitioned into 11 zones, which are subsequently divided into 88 panels. Within each panel, the ore body is further discretized into (1) stopes—areas outlined solely for the purpose of extracting ore, (2) drifts—areas that are mined to develop access to a stope and (or) for the extraction of ore, and (3) pillars—large ore blocks that support the mine infrastructure. For ease of presentation, we refer to any stope, drift, or pillar as a block. At the time of this writing, the Lisheen deposit contained 1,193 mining blocks (i.e., candidates for extraction).

A critical factor in planning production is the determination of a profit-maximizing cut-off grade, which classifies a block as either ore or waste based on the percentage of mineral content or grade of that block. Because Lisheen is a polymetallic deposit with a zinc-to-lead ratio of 5:1, a combined mineral content or zinc-equivalent grade is calculated for each block. Ore blocks are candidates for extraction and refinement to mineral concentrate, whereas waste blocks are only extracted to facilitate extraction of adjacent ore blocks and are disposed of underground.

To determine the best extraction technique to employ in a given area of the mine, geotechnical engineers consider the strength of the host rock (i.e., the rock that encompasses the ore body), in addition to other factors such as the ore body slope and thickness. With varying host rock strength and blocks that range in thickness from 1 to 30 meters, the Lisheen mine applies three mining methods to extract ore: room-and-pillar, long-hole stoping, and drift-and-fill.

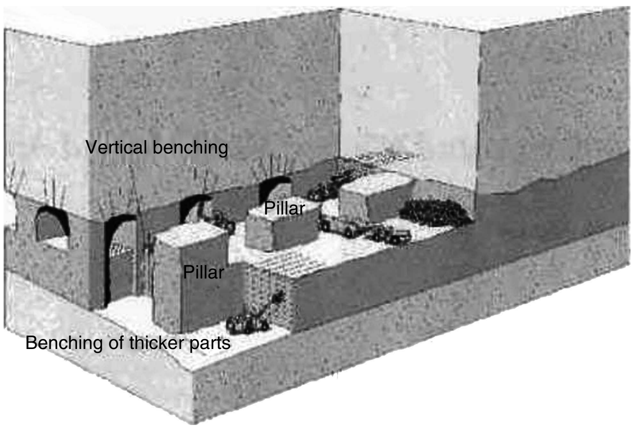

Figure 1: In room-and-pillar mining, a self-supporting method, some pillars of ore are left in place to support the rock that encompasses the ore body, while the remaining ore is extracted. When only pillars remain, pillars are removed in a sequence that starts with the furthest pillar from the mine exit, allowing the host rock to cave in as the extraction proceeds toward the exit (Hamrin 1997).

Where the host rock is strong and the ore body is not steeply angled, room-and-pillar mining is preferred (see Figure 1). Areas of the mine in which the ore is particularly thick and the host rock is strong are suited to the large-scale and economically efficient long-hole stoping method (see Figure 2). Finally, where the mine has poor host rock strength, drift-and-fill mining is practiced (see Figure 3). Figures 1–3 are representative of the operations at Lisheen; however, in reality, the operations are tailored. Long-hole stoping dominates 70 percent of the extraction, drift-and-fill mining accounts for another 20 percent, and room-and-pillar mining constitutes the remainder.

With 12 years of production already completed, the mine has an extensive network of haulage routes, which are used to transport the ore from the panels to the crusher, a machine used to break the pieces of ore into manageable sizes before conveyance to the surface. Once above ground, the ore is transported to the mill to be refined into metal concentrate. The operational plan at Lisheen recommends filling the mill (i.e., continually running the mill at as close to full capacity as possible). This requirement is difficult to satisfy because the quantity and the average quality (i.e., head grade) of the ore that enters the mill during a given period must be blended within a specific range; otherwise, the process stops. Many mining companies stockpile ore to more precisely control the quality and quantity of the ore that feeds into the mill. At Lisheen, a small above-ground surge pile buffers ore production to ensure continual mill operations over holiday weekends or during unplanned

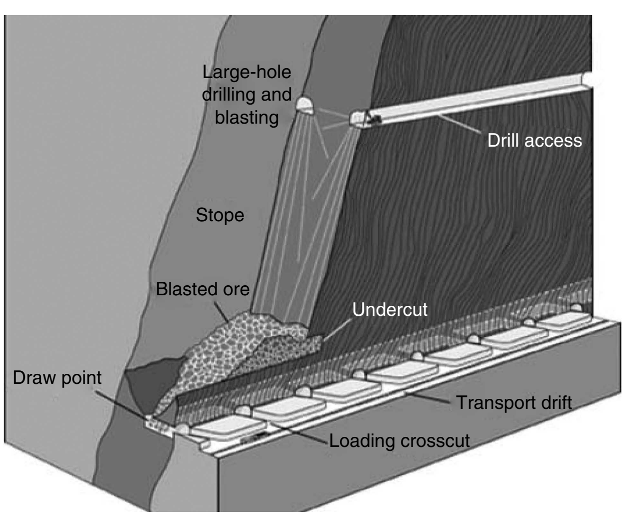

Figure 2: In the long-hole stoping process, a sublevel drift is developed below the ore to be excavated. The ore is then drilled and blasted and falls to the sublevel where load-haul-dump vehicles scoop and load it into trucks for transport to the crusher (Hamrin 1997).

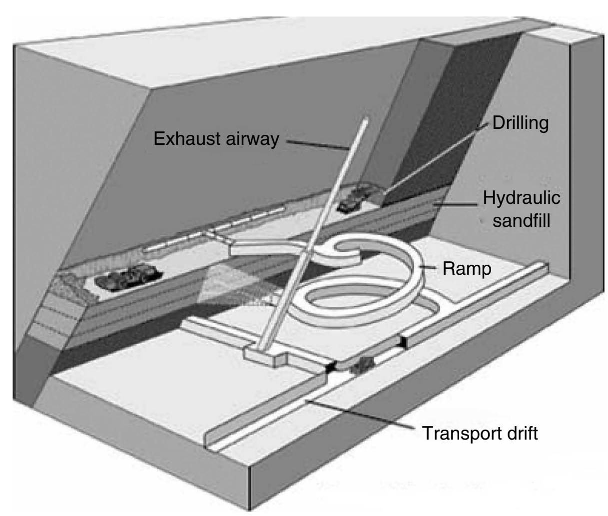

Figure 3: In drift-and-fill mining, a corridor of ore, known as a drift, is removed from the ore body. The remaining void is then filled using a cement-based mixture to provide structural support before cutting an adjoining drift through the ore (Hamrin 1997).

production shutdowns. However, this stockpile is not suited to blending operations. Therefore, with limited stockpiling capability, mine planners must carefully select the blocks that they schedule for extraction at any one time so that the quantity and average head grade of the ore reaching the surface is within the mill's acceptable limits.

After refinement at the mill, the concentrate is loaded on trucks and shipped, and the waste from the milling process (i.e., tailings) are discarded. Most are deposited into a tailings pond that will be drained and covered once mining is exhausted. Remaining tailings are used to create a cement paste for filling some of the voids left by extraction, a technique known as backfilling.

Aside from providing a location for the disposal of tailings, backfilling is required in some areas to maintain structural integrity and to allow mining to continue. Prior to backfilling a void, it must first be prepared by sealing it with wooden panels and installing hoses that run from the underground void to the surface. Backfill paste, mixed in a plant above ground, is then pumped via the hoses down into the area being filled. A void, depending on its size, may require anywhere from a day to a month to fill. Once filling is complete, an additional 24 days are allowed for the paste cement to set, after which time extraction may begin on the adjacent blocks that require the support from the backfilled areas.

The quality of the schedule for the production process described previously is a significant driver of the mining operation's profitability. To produce a timely and coordinated production plan, the planner at Lisheen must consider the rate of extraction associated with each block's mining method, the dependencies between activities including backfilling requirements, and the size and grade of each block, so that the blend of ore reaching the mill is satisfactory. This difficult assignment is further complicated because mining follows a number of narrow veins of high-grade ore located along fault lines between plates of different rock strength. The interaction of these plates has produced an ore body with an inconsistent distribution of metal. Engineers highlight pockets of high-grade material by their choice of cut-off grade and these pockets form the basis for the creation of block shapes and the selection of extraction

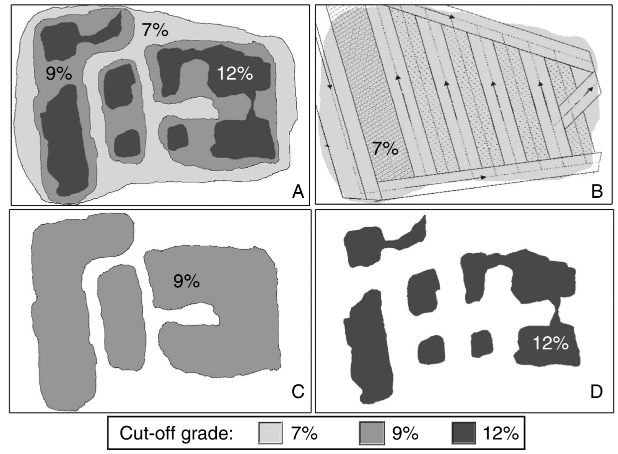

Figure 4: Cut-off grade selection is a significant factor in determining ore block shape and size. Box A shows a section of ore body with a nonuniform distribution of metal. Selecting a cut-off grade of seven percent results in a large area classified as ore (Box B), which can be divided into reasonably homogeneous blocks and extracted with a single mining method (e.g., the drift-and-fill process). At a cut-off grade of nine percent, the material classified as ore diminishes (Box C). Because drift-and-fill mining would now excavate too much waste rock, a different approach is used. Box D shows material classified as ore at a cut-off grade of 12 percent. With most of the area now defined as waste, a targeted mining method is suitable.

methods (see Figure 4). Hence, the grade of a block can vary greatly between adjacent blocks, and we sometimes find a high-grade ore block located beside a waste block. In addition, unlike open-pit mining, in which the ore body is subdivided into blocks of equal dimensions, the size and shape of the blocks can vary greatly. As a result, some blocks can be extracted in a day, while excavating others requires months. Consequently, each panel at Lisheen has such specific mining requirements that developing general mining rules is impossible, even at the zone level, which would make sequencing activities straightforward.

Schedule creation is particularly difficult because of the intertemporal effects on future mining activity associated with each extraction decision. These effects become especially important as mining reserves approach depletion and the variety in grade of the ore that is available to satisfy milling requirements also diminishes. With mine closure on the horizon, management must make critical decisions about which blocks to take and which blocks to leave behind.

In this paper, we show how we use integer programming (IP) to determine a near-optimal schedule for the Lisheen underground mine for its remaining years of operation. Arguably, the mine would have benefited greatly from an automated scheduling procedure more than a decade ago. However, we became involved with the project in only the past few years. As we describe in our paper, legacy limits the extent to which we can impact mining operations with our scheduling procedures. However, we are able to dictate a sequence in which blocks should be extracted to satisfy the profit-maximization goals for the remaining life of the mine, especially as those goals pertain to the extraction of haulage pillars along a critical retreat path out of the mine, and the ability of the mine to sustain an operationally acceptable level of production for as long as possible.

Literature Review

Williams et al. (1973) first demonstrated the potential role for optimization in underground mine production planning. Although their linear programming approach suits certain strategic mining problems involving blending or cost reduction, the method cannot incorporate the binary decisions required to enforce the logic to schedule production at an operational level. Chanda (1990) recognizes IP as a viable method for modeling discrete decision making. He combines mixed-integer programming (MIP) and simulation to produce a schedule for six consecutive work shifts for a copper mine in Zambia. A similar approach by Winkler (1996) employs integer variables, not only to model operational relationships, but also to model the costs as piecewise linear functions. She highlights the advantages of applying MIP to underground mine scheduling and illustrates the exponential complexity associated with this approach when generating a multiperiod schedule. Solving one single-period MIP model at a time, she relies on simulation to produce a multiperiod schedule for a German coal mine. Trout (1995) presents a more generalized MIP approach to underground mine production scheduling. Introducing variable restrictions to reduce the size of the model, he produces a schedule that maximizes the net present value of a copper mine for a 17-period horizon. Carlyle and Eaves (2001) apply a MIP model to schedule production at an underground platinum and palladium mine. They maximize discounted ore revenue by solving their model for a number of mine expansion scenarios.

Sarin and West-Hansen (2005) developed a Benders' decomposition technique to produce an exact solution for their MIP model. Applying their model to a coal mining case study results in a profit-maximizing 100-week schedule. Newman and Kuchta (2007) optimize long-term production at an underground iron ore mine in Sweden. With an objective of minimizing deviations from contracted production quantities, they use an aggregation heuristic to solve a MIP model for a 60-month horizon. This model forms the basis for Martinez and Newman (2011) to schedule both long- and short-term production for the same iron ore mine. That model provides greater resolution in the near term to satisfy detailed operational requirements. They produce solutions for a 48-month horizon using a decomposition heuristic.

Similar to the work we cite previously, we formulate, solve, and implement an IP model for production scheduling in an underground mining operation. However, although almost all underground mine production scheduling models have similarities (e.g., blocks, periods, sequencing constraints, resource restrictions), each mine operates in a specific manner. For example, objectives tend to differ between mines; an appropriate objective for a mine containing precious metal may be to maximize net present value, whereas the objective for a mine containing base metal may be to minimize deviations from long-term contracts. Some mines may stockpile to exploit blending and its associated quality of output; others may regard such a policy as a nuisance, requiring rehandling and its costs. The six standard underground mining methods have many variants. As such, each method necessitates the implementation of different operational policies. Therefore, writing one general model for underground mine production scheduling is difficult. Specific aspects relevant to our application are (1) a discounted objective involving metal (but no explicit economic parameters), (2) a complex and extremely ill-defined and nonhomogeneous (i.e., spatially, temporally, and regarding their metal content) set of areas, and (3) a mine that uses three underground mining methods and the associated rules governing extraction and backfill, where applicable.

We now describe the mine with respect to optimizing its production schedule in more detail.

Production Scheduling at Lisheen Mine

At the Lisheen mine, ore production will continue as long as the operation is profitable. However, as at other mines, a time will come when the metal that the mine produces will no longer be able to justify the costs of its extraction and, although ore blocks remain below ground, the mine will close. This time is called economic exhaustion. The challenge then is to extract the combination of blocks that realizes the most value from the mine before economic exhaustion is reached.

Production schedulers at Lisheen use iGantt (MineMax 2012) mine planning software to manually generate production schedules. This difficult and time-consuming task, performed semi-annually, requires the manual arrangement of approximately 2,000 extraction and backfilling activities on a Gantt chart (see Figure 5), such that planned ore production is sufficient to keep the mill running as close to capacity as possible.

To narrow the scope of the task, the planner adopts the following assumptions about the mining operation:

1. Financial aspects of the operation are ignored. Planners do not explicitly consider operational costs, mineral prices, and costs associated with mine closure when generating schedules.
2. The goal is to maximize metal production over the life of the mine. Although the mine planners do not explicitly consider financial information, they seek to produce as much metal up front in the remaining life of the mine as possible. The rationale behind this objective is that, because Lisheen sells metal on the spot market, too much risk is associated with scheduling large quantities of metal late in the mine's life when a drop in the spot price might render that metal uneconomical to mine.
3. The mining methods (i.e., room-and-pillar, longhole stoping, drift-and-fill) for each mining area are fixed. A change in the method of extraction for an area may require the planner to redefine the ore block shape, establish a new order of mining, or change the rate of mining for that area. As a result, for each mining activity, we also assume the same fixed

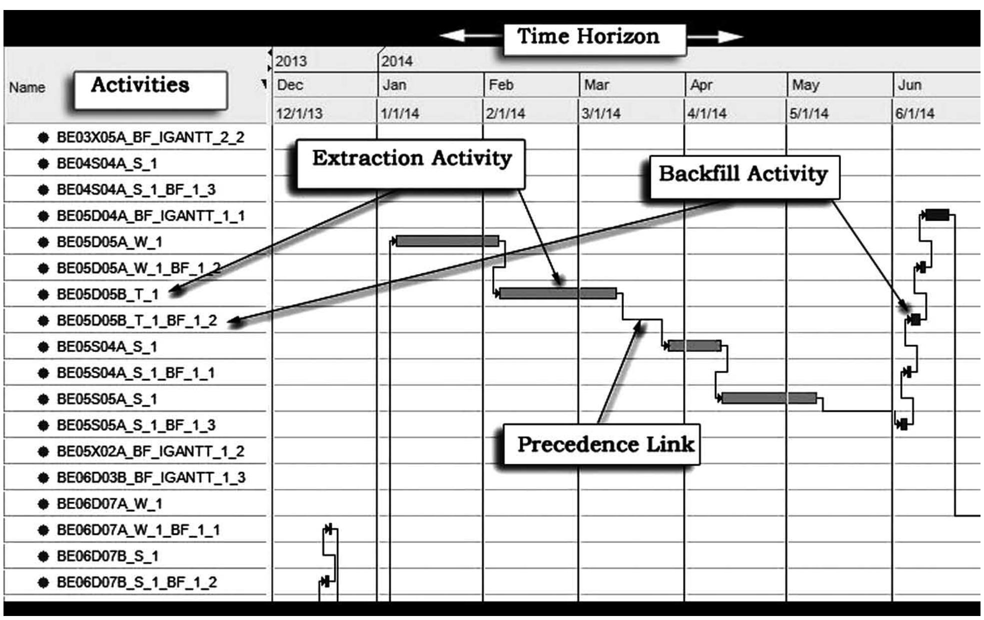

Figure 5: This example of manual scheduling using a Gantt chart shows extraction and backfilling activities as horizontal bars. The thin lines between the bars represent the precedence relationships between the activities. To meet production targets, planners arrange these activities on the timeline.

rates that Lisheen has used to generate its manual schedules.

4. The mine design and infrastructure are fixed. The planner does not consider developing new access drifts (i.e., passageways) to reach the ore blocks or haulage routes to connect panels with the crusher.

5. The cut-off grade that defines a block as ore or as waste based on the percentage of mineral content is fixed. During the engineering design of the mine, the planner selects a cut-off grade to maximize the mine value. Although economic conditions might indicate that a different cut-off grade could increase the mine's value, the change would also require the development of a new engineering plan (see Figure 4).

6. The rates for extraction and backfilling activities are fixed. Based on the block tonnage, excavation distance, or void volume, the planner can quickly calculate the time required to complete an activity. This is a standard assumption in most scheduling models, and changing the rates would be tantamount to changing a strategic decision (e.g., fleet size or work shift length).

7. Sufficient resources are in place to implement the schedule. Production at Lisheen is not limited by the availability of equipment or labor.

With these assumptions in mind, the mine planner must determine a start date for each extraction and backfilling activity, given the following constraints:

• Monthly tonnage targets for mine production cannot be exceeded. The mine has a production capacity constraint for each month based on the number of working areas that can be active at one time and the number of working days in that month. Mine production includes extraction of ore that feeds the mill, and necessary waste.

• Metal output from the mill must be maintained between monthly maximum and minimum levels. These blending constraints on metal production ensure that the quantity and head grade of the ore feeding the mill are maintained within operational limits.

• A monthly limit on cement paste for backfilling must not be surpassed. Paste availability depends on the tailings produced as waste by the mill and the capacity of the backfill plant that produces the paste cement. The production of tailings is a function of the quantity and grade of ore reaching the mill, which means that backfilling may be constrained by the production of ore in a previous period. However, because only a fraction of the voids created by extraction need to be filled at Lisheen, the availability of tailings does not limit paste production.

• Sequencing constraints must be observed. The extraction of a block or backfilling of a void must satisfy any sequencing rules that exist between that activity and any other activity. Often called precedences, these relationships may be defined implicitly as a consequence of the combination of mining methods chosen for a panel. Otherwise, precedences are defined explicitly to ensure that the schedule remains feasible.

• Once started, an activity must proceed continuously until completion. The engineering design defines blocks under the assumption that the entire block will be extracted. Partial extraction of blocks or partial backfilling of voids would result in structural instability.

Manual scheduling of ore production under these constraints is challenging. To satisfy the Lisheen business objectives, the planner adopts a trial-and-error approach to bring metal forward in the production plan, while simultaneously adhering to the ore production limits and grade blending requirements. For large problems, the combinational nature of this sequencing and blending problem would require the planner to enumerate a staggering number of alternative scenarios. Consequently, planners often have difficulty producing schedules that satisfy these constraints. Finally, the planner must be mindful that extraction decisions made today can have consequences for future mining activities. For example, tactical scheduling decisions to satisfy the next week's production quota may unintentionally prevent access to an area of ore and destroy its value.

As Lisheen approaches closure, a particularly difficult decision facing the planner is when, if ever, to extract ore blocks that form part of the mine's critical infrastructure. Throughout the mine, trucks carry ore from the panels to the crusher along haulage routes, which are supported by large ore blocks called haulage pillars. Moreover, these pillars are also candidates for extraction, and often contain valuable high-grade ore. However, once a haulage pillar is removed, the ore blocks that require the pillar to remain in place for their extraction can no longer be reached; these are termed sterilized reserves. Hence, the planner must consider the opportunity cost of sterilizing the ore associated with the extraction of a haulage pillar (see Figure 6).

Given the challenging nature of the task, manual scheduling can result in solutions that are far from optimal. Also, the time-intensive process of manual scheduling, which can require several weeks to complete at Lisheen, precludes any possibility of scenario analysis to evaluate the operation with respect to other parameter values (e.g., realization that poorer-than-expected ground quality in a panel would slow down mining in that area and necessitate an update of the mining rates for that zone). Finally, the planner cannot examine all of the schedule's permutations; therefore, measuring the quality of the solution is impossible.

By contrast, a mathematical optimization approach allows the planner to enumerate multiple schedules, select the one that results in the highest objective value, and then show a measure of the solution quality. Moreover, the planner can quickly generate an IP solution, and thus address any changes that may arise because of unexpected events.

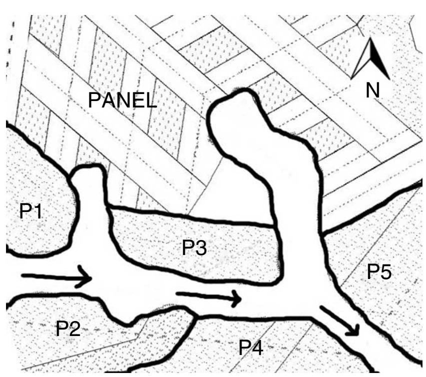

Figure 6: This example shows a mining panel adjacent to a haulage route. The pillars that support the haulage route are shown as P1–P5. The arrows indicate the direction of the mine exit. The haulage path collapses west of any pillar that is removed. Hence, the extraction precedence for these pillars would require that P1 be taken first and P5 last. Removal of pillar P3, P4, or P5 would result in the destruction of the haulage path in that area, eliminating access to the panel.

Integer-Programming Approach

Adopting an IP approach, we determine a near-optimal schedule for the Lisheen mine. We cast the problem mathematically, with the objective of recovering the maximum amount of discounted metal possible from the mine over its life. We also incorporate the same constraints and assumptions mentioned previously in relation to manual scheduling.

We apply a small discount factor to the objective solely to encourage a solution in which metal production is brought forward in the schedule. Specifically, without this discount factor, a solution with one tonne of ore extracted either in week 1 or week 2 would be equivalent. The addition of this factor produces a solution with the tonne of ore extracted in the first week. Such discounting can also improve the run-time performance of mathematical programs (Klotz and Newman 2013).

To measure value, we assume a constant metal price, which we normalize to 1 (see the formulation in the appendix). Multiplying the objective by a constant factor would not change the solution we produce, and would obfuscate our intent, which is to pull the blocks with the highest metal content forward in the schedule to realize the greatest metal sales (subject to operational constraints of the mine) in the shortest possible time. Indeed, lack of an explicit monetary factor happens occasionally in a base metal extraction operation (Newman and Kuchta 2007, Martinez and Newman 2011).

Although relatively straightforward to describe in mathematical terms, the complexity of the problem makes it difficult to solve in practice. The sequencing rules that govern the relationships between activities are idiosyncratic, complicated, and defined unclearly. In describing a challenge of solving integer programs for mine scheduling, Smith (1998) emphasizes the difficulty associated with complex precedence constraints. These precedences—encoded in iGantt during manual scheduling—are an output from that software. However, the iGantt precedence set provides few degrees of freedom, and solving our integer program with these encoded precedences would simply return the existing iGantt schedule. Inspection of the iGantt relationships using a computer-generated block model in Maptek's Vulcan (Maptek 2012) reveals that, although maintaining a feasible schedule requires many rules, some rules merely reflect a subjective choice in the order of mining certain ore blocks. Specifically, the precedence rules relating the extraction of infrastructural haulage pillars to other activities in dependent panels are unnecessarily restrictive. In most cases, the existing mining rule simply reflects the order in which the pillar was previously scheduled, often requiring extraction to wait until all dependent panel activities have finished. As previously mentioned, the timing of the extraction of these high-grade haulage pillars is a critical aspect of the mine schedule. Hence, we want to allow our model the freedom to choose when, if ever, to extract them. To this end, we reverse the existing logic, that a pillar must wait until the mining of the dependent panel has finished, and instead create a rule that prevents any further mining in the dependent panel once the haulage pillar has been extracted. For each panel, we identify the critical haulage pillar whose extraction would prevent any further activity in that panel. We then define a constraint that enforces this relationship, not just for that one critical haulage pillar, but for all blocks along the haulage route that, if extracted, would sterilize remaining ore in that panel. In this way, we delineate critical haulage routes throughout the mine that characterize the relationships between haulage pillars and mining panels (see Figure 7).

Another difficulty in constructing the sequencing constraints arises from activities for which no sequencing rules are defined. Including these orphaned activities in our model without defining rules for their execution would likely produce an IP solution that we could not implement. Without explicit and objective mining rules, we would have to spend considerable effort examining each precedence relationship between a pair of activities to determine if it leads to a valid constraint or whether it is merely a

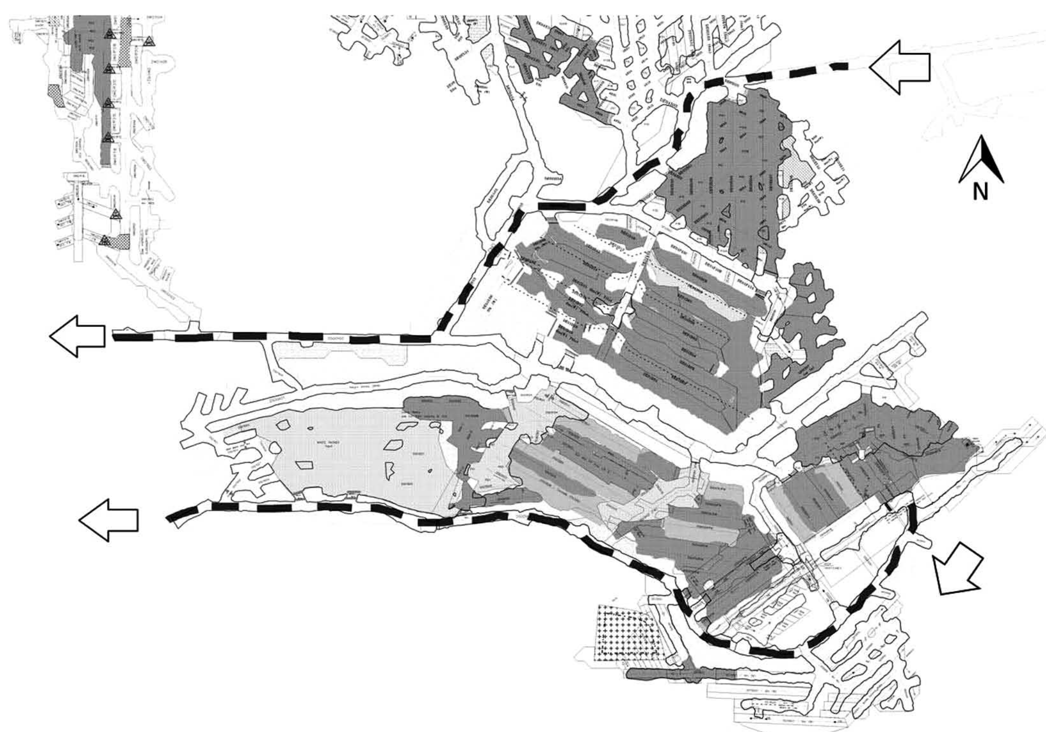

Figure 7: In this example of critical haulage routes, the dashed lines show haulage routes through a mining zone. These haulage routes are bordered by major panels containing blocks that are candidates for extraction. The pillars that support these haulage routes often consist of valuable ore blocks. However, once removed, we lose access to the area east of the extracted pillar.

subjective planning choice. Working with Lisheen, we redefine the precedence set by identifying and removing the subjective rules and by defining new precedents where needed.

A complication with our IP approach also arises from the heterogeneity of the ore blocks. The size of an ore block can vary greatly, which makes determining a standard time interval, or fidelity, for the problem difficult. The fidelity of the model defines the discrete times during which an activity can begin. In scheduling, we could assign an activity to start at a specific month, day, or hour. However, use of discrete time intervals results in schedules in which activities that, in practice, would take less than the interval to complete appear to require the full time interval before a dependent activity could begin. For example, a block that would require a day to extract in practice would be scheduled for a month if we were solving at monthly fidelity. Thus, the smaller the time interval, the closer together we can schedule activities. But, this benefit comes at a computational cost, because we must account for finer fidelity by defining variables for each activity and start-time combination.

A solution produced at monthly fidelity may work well for a mine that has homogeneous, large blocks that require more than a month to extract; however, as the histogram in Figure 8 illustrates, Lisheen has many small areas that can be mined in less than a day.

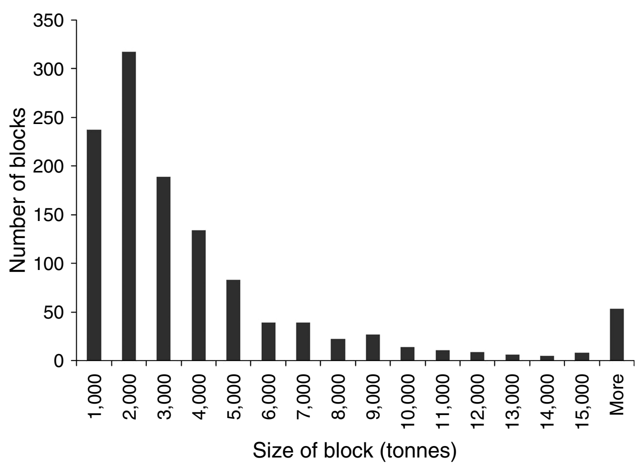

Figure 8: In this histogram of block size, smaller blocks correspond to a small access drift or a cross cut used to facilitate extraction of a larger block, such as a stope. The smallest blocks can be mined in a day or less. The largest require months to extract.

Therefore, applying a model with monthly fidelity results in schedule delays, because these small areas would be allotted a month for extraction, unnecessarily extending the extraction times of dependent areas. A frequently used approach to address such fidelity problems is to aggregate the smaller mining areas into larger clumps of similar size. However, such an approach reduces the resolution at which blending decisions can be made. As Smith (1998) points out, deposits rarely contain grade that is homogeneous enough to ignore grade blending requirements. Such is the case at Lisheen, where the high variation in the grades within adjacent blocks, in addition to the complexity of the mining precedences, precludes block aggregation.

To produce an implementable schedule, we require a solution at weekly fidelity for a two-year horizon (i.e., a horizon that corresponds to the projected life of the mine). The resulting problem contains 2,214 (extraction and backfilling activities) $\times$ 105 (weeks) $=$ 232,470 binary variables. We reduce this number by implementing start-date restrictions for each activity; these restrictions do not compromise the quality of the solution, but only serve to eliminate variables, which would necessarily assume a value of zero in an optimal, or even any feasible, solution. The manual schedule in iGantt provides a basis from which we develop early start-date restrictions for haulage pillar retreats. However, because we relax precedence rules between panels, using the same approach for panel activities is too restrictive. Consequently, for panel activities, we use a variant of an existing procedure for establishing early start dates for machine placement activities in a sublevel caving mine (Martinez and Newman 2011). Our variant considers blocks already scheduled to start, the prescribed sequence in which blocks must be extracted, and mining rates to determine that certain blocks cannot be extracted before a specific early start date. For example, if the following three conditions hold, (1) block A must be mined to obtain access to block B, (2) block A has started extraction in period 1, and (3) block A requires two weeks to extract, then we cannot begin to extract block B until at least the start of the third week. Correspondingly, we remove variables from the formulation that represent block B starting to be mined either in week 1 or in week 2. For sequences involving more than a single predecessor block, we simply add the durations of all predecessor blocks to the start date of the first block in the sequence to determine the start date of a given block; in these cases, all blocks in the sequence must be mined, and at a predetermined rate, which produces an exact result (i.e., a time before which a block at the end of a precedence chain cannot possibly begin to be extracted). This approach reduces the problem to 56,276 variables; however, with over a million constraints, the problem remains intractable. Therefore, not only do we use the exact approach of eliminating variables to reduce the problem size, but we also employ a heuristic on the reduced problem to produce a schedule in a reasonable amount of time (see Figure 9).

We can solve smaller instances of the (reduced) problem by including fewer activities in the model. We take advantage of this insight to find a solution that includes all the activities by decomposing the original problem into a series of smaller problems, which we solve in stages. With this approach, we begin by separating the extraction activities into sets of progressively lower metal content. In practice, we might use two or three sets for panel activities and three to five sets for pillar activities; however, to illustrate the approach here, we consider only two sets: high-grade and low-grade activities. Considering only the subset of high-grade activities, we solve the model (P), given in the appendix, for those activities at a weekly time fidelity for all weeks in the planning horizon. The solution yields a production schedule for the high-grade activities on a weekly basis. We then include the set of low-grade activities in the next model run, fixing those high-grade activities to the corresponding weekly start dates from the first solve. The net result is that we now obtain a solution that includes some subset of the low-grade activities, which are scheduled around the fixed high-grade activities. Note that to enforce precedence between successive solves, we must insert placeholders for the time required to extract the predecessors absent in the current solve whose successors are present in the current solve.

This approach alone will not produce a near-optimal solution for Lisheen, because the problem's complexity is determined largely by the number of haulage pillars included in the model. Haulage pillars complicate the problem because (1) they have few, if any, precedence requirements, and can consequently be scheduled for extraction at any point during the horizon, and (2) the value of adding a haulage pillar to the schedule must be weighed against the total value of the dependent panel blocks that would be sterilized as a result. If we exclude the haulage pillars, we can solve the model for the entire horizon in a matter of minutes; however, as we introduce haulage

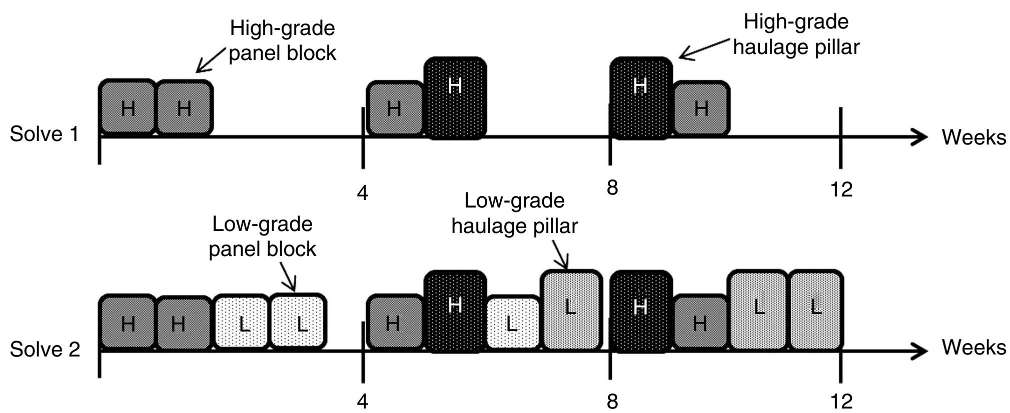

Figure 9: In this example of the heuristic solution method, we solve the problem twice at weekly fidelity. Solve 1 shows the solution for a 12-week horizon during which only high-grade panel and pillar activities are considered for scheduling. The times at which these blocks are extracted constitute the schedule at this point. The solution resulting from Solve 2 shows that blocks that were scheduled within Solve 1 are now required to remain scheduled at the equivalent week when solving this problem, but now low-grade panel blocks and pillars are also included.

pillars, the problem's complexity increases quickly until the problem becomes intractable. The key to solving the problem in stages is to find the right balance between the number of pillar activities and the number of other extraction and backfilling activities included in each stage of the heuristic. If the initial solve includes too few pillars, then that solution contains too much panel mining in the schedule, leaving less opportunity to accommodate remaining pillars during the next solve stage. If it includes too many pillars, the problem becomes intractable, or, in the tractable case with fewer panel blocks, the solution contains some pillars that are scheduled far too early, sterilizing dependent panel blocks and preventing them from entering the schedule during the next stage of the heuristic. For a given horizon, we use a trial-and-error approach to determine the number of activities and the ratio of pillar-to-panel activities to include at each stage in the heuristic. The number of stages depends on the length of the horizon, with longer horizons requiring a more gradual introduction of pillars and, thus, more stages.

Implementation

Our schedule, as it appears in the output file from the CPLEX solver (IBM 2011), is merely a list of scheduled activities and their corresponding start weeks. In this format, the schedule is difficult for Lisheen to validate and impractical to implement. Therefore, we integrate our solution with Lisheen's iGantt scheduling software to present our schedule in a way that is familiar to Lisheen's planners.

iGantt provides effective schedule visualization and reporting features that the planners use each day. The software displays the schedule graphically as a Gantt chart (see Figure 5), i.e., a timeline of scheduled activities and their interdependencies. In addition, iGantt can present a three-dimensional block model animation of the schedule, enabling the planner to quickly identify sequencing problems or other infeasibilities. Lisheen management uses iGantt's reporting capability extensively, for example, to produce short-term work schedules, to summarize historical backfill paste usage, or to forecast future resource requirements (e.g., the number of required cable bolting bits). Thus, implementing our schedule using iGantt has the added benefit of providing continuity in the mines' daily managerial operations.

We transfer our schedule to iGantt by creating an import file of the data, with rows of activities and columns of activity attributes (e.g., tonnage associated with an extraction activity). To this file, we add a start-date column, which we populate with the week during which each activity begins according to our solution. Importing this information into iGantt as a comma-separated-value file is straightforward, but the resulting schedule that iGantt creates is incorrect because the default functionality of the software chooses a start date for each activity that differs from the IP schedule's start date. We override these automated start dates by using a customized script that forces iGantt to adopt the IP activity start dates; these are given in terms of a start week, whereas iGantt's time fidelity is given in milliseconds. We convert the time at which an activity starts based on the beginning of the week in which it is scheduled in the integer program to milliseconds from the beginning of the horizon, and use that converted date in iGantt. We do not specify activity end dates to import into iGantt; instead, we let iGantt set them based on the activity rates. Admittedly, we are not scheduling at the level of fidelity that iGantt allows or that manual schedulers at the mine could use with the iGantt software. However, mine planners can use discretion to shift our dates within an allowable window, and a schedule at the level of milliseconds is impractical. Shifting activities outside of the allowable window (i.e., manually moving activities in iGantt forward to begin at the end-date of the preceding activity) would likely violate ore production and (or) milling constraints. Although our schedule may leave gaps in time (e.g., in our model, we record an activity that requires five days as requiring a week), our schedule provides a conservative estimate of the metal producible within the given horizon, and may be more realistic under the unanticipated conditions that necessarily prevail in an underground mining operation.

Review and validation are possible once we format our solution into an iGantt schedule. Using the iGantt visualizer, the planner examines our schedule for feasibility. At this point, our schedule may be infeasible from a practical standpoint because of missing or invalid sequencing constraints. The complexity of the sequencing rules makes it difficult to spot minor conflicts or omissions a priori; however, we highlight these mistakes in the animated iGantt schedule. With guidance from Lisheen, we correct any illogical or overlooked sequencing constraints in our integer program and continue the process of review and correction until we obtain a workable schedule.

Results

Solving our IP model for a 104-week horizon, we compare our solution with the manually generated schedule on the basis of, for example, metal output (from the mill), ore production (based on extracted material), and backfill paste usage (see Table 1). With respect to the objective of maximizing metal production over the life of the mine, the integer program solution shows 0.17 percent less metal assigned for recovery than the manual approach. When taken in isolation, this shortfall in total metal output is indicative of a poor result. However, we must consider a number of other measures before we can properly assess the quality of the IP solution.

Most significantly, the integer program discounts future production so that we can bring forward metal in the schedule. This satisfies management's desire to produce as much metal up front in the remaining life of the mine as possible. As we have mentioned, Lisheen management is not comfortable with scheduling large quantities of metal late in the mine's life because a low future spot price might preclude economic viability of the metal's extraction. Although the integer program may not increase total metal production, it does bring a significant quantity of metal forward in the schedule, increasing metal production by 10.51 percent in the first year (see Figure 10).

	Manual schedule	IP schedule	IP gain (%)	Comment
Metal output (tonnes)	321,154	320,621	−0.17	No IP improvement
Ore production (tonnes)	2,444,446	2,414,298	−1.23	IP improvement
Backfill paste used (cu.m)	947,104	810,273	−14.45	IP improvement
Waste mined (tonnes)	115,790	114,424	−1.18	IP improvement

Table 1: The table shows a comparison of the IP and manually generated solutions for a 104-week horizon.

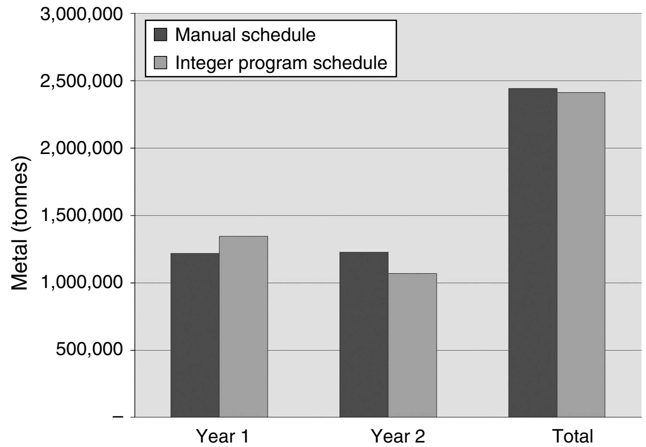

Figure 10: In this example of annual metal production, the integer program brings metal forward in the production schedule.

Following the IP schedule would also reduce costs because the decreased metal yield is accompanied by an even greater reduction in mining activity, with 1.23 percent less ore production and 1.18 percent less waste mining than the manual schedule generates. Therefore, the integer program schedules more efficient extraction; it also schedules significantly fewer backfill activities, corresponding to a 14.45 percent reduction in paste consumption over the remaining life of the mine. By releasing labor to more profitable extraction activities and by reducing wait time for backfill set up, pouring, and setting, this reduction in backfilling enables the planner to include valuable extraction activities that would otherwise have been precluded. In addition, although a backfilled void generally requires one month to set, at Lisheen, oxides often interact with the backfill paste, preventing or delaying hardening. Consequently, backfilling can sometimes fail to harden or can require two months or longer to complete, again forcing extraction activities outside of the schedule, resulting in a loss of value. Thus, scheduling less backfill activity reduces the potential for schedule delay.

When concurrently scheduled, backfill activities can also be responsible for large temporary drops in ore production which, for a number of reasons, are problematic. In particular, because it is highly inefficient for the mill to operate below a minimum ore threshold, low production levels result in mill shutdowns and costly restarts. To boost production in low-production periods, Lisheen planners use a trial-and-error approach to shift production activities in the manual schedule. However, the planner cannot be sure if rescheduling alone can rectify the drop in production. For Lisheen, production below approximately 80,000 tonnes of ore results in mill shutdown, and the manual schedule falls below this level for two months at the start of the second year (see Figure 11). By contrast, the IP model sets a lower bound for production; as a result, it suffers no mill-halting drops in ore production. Also, because lower production requires fewer workers, management would be forced to reassign miners to other duties or worse, lay off part of the workforce, a decision that would bring union action against Lisheen.

Not only does the IP schedule avoid dips in production, but it also schedules ore production more consistently than the manual method, with an average month-to-month change in production of 10,132 tonnes compared with 14,796 tonnes for the manual schedule. Thus, the IP solution enables mine management to create a more steady work flow than the manual plan does.

Scenario Analysis

The integer program produces schedules in fewer than 20 hours. This is significantly shorter than the manual approach, which can require several weeks to complete. Consequently, the IP method is useful for scenario analysis or for expediting the generation of new schedules to account for unforeseen events or changes in engineering design. Because the mine is approaching the end of production, we examine one important set of scenarios: alternative closure dates.

The closure date that defines the end of the mine's operational life is a best estimate from the manual scheduling process. It is the point in the schedule at which economic exhaustion occurs and the mine can no longer produce enough metal to justify its operational costs. The time at which Lisheen reaches economic exhaustion depends on future metal prices and costs, including significant mine closure and rehabilitation expenses. Although our model does not explicitly incorporate these financial components, it can generate feasible production alternatives, thus enabling management to understand how the schedule might change by shortening or lengthening the mine's life, and to make more informed mine closure

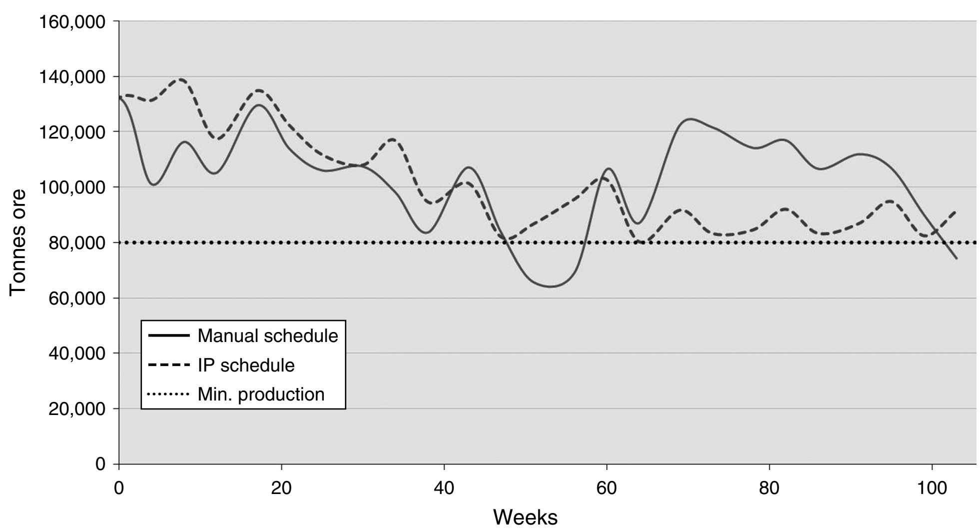

Figure 11: In this ore production example, the IP solution has more consistent month-to-month production than the manually generated schedule. Production in the IP schedule remains above 80,000 tonnes, allowing the mill to operate continually. By contrast, the manual schedule suffers from low production starting in week 48, which would result in a temporary mill shutdown and labor problems.

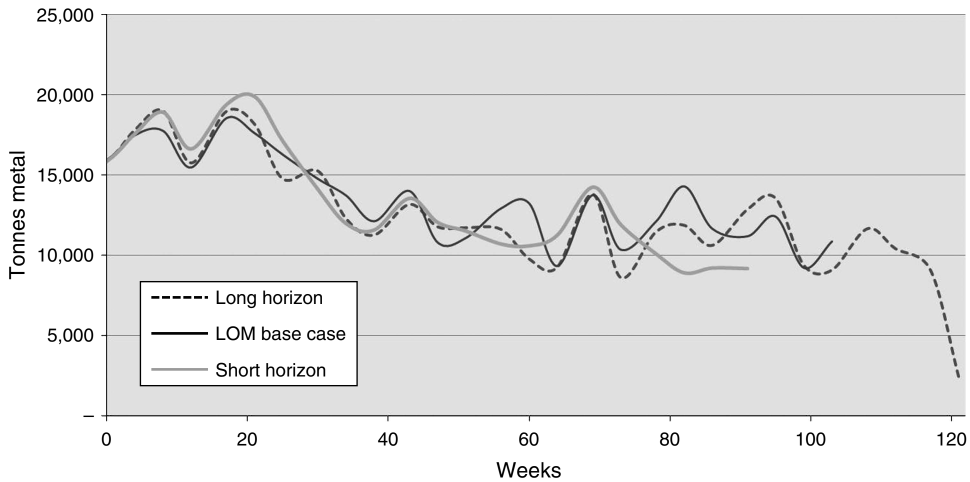

Figure 12: We compare IP solutions for varying time horizons against the IP solution for the current mine life. Increasing the mine life by another three months would be a viable option for Lisheen. Similarity in the production tonnage in each week during the first year is a consequence of the low degrees of freedom in the precedence rules. Later, when haulage pillars become candidates for extraction, weekly metal production differs more across scenarios, reflecting the more flexible precedence constraints on these pillars in our model.

decisions. To provide insight, we run scenarios with different horizons (see Figure 12).

With our previous results for a 104-week mine life providing a base case, we solve our integer program for two mine-life scenarios: (1) an early-closure option with a 92-week horizon, and (2) an extended horizon of 122 weeks. We compare the results of these scenarios to our base case on metal production (see Table 2).

The early-closure scenario is motivated by the observation that after the first 20 weeks of the 104-week IP schedule, ore production and metal output never approach maximum levels again. Consequently, Lisheen management is interested in examining the possibility that, if the base-case schedule has sufficient production or output slack, a solution for a shorter horizon might produce a similar or greater quantity of metal than the base-case scenario.

Scenario	92 weeks	104 weeks	122 weeks
Short horizon	280,095	n/a	n/a
Life-of-mine base case	288,165	320,621	n/a
Long horizon	279,532	311,511	342,614

Table 2: The table shows IP mine-life scenarios by comparing metal production.

Specifically, we shortened the horizon to 92 weeks because it corresponds to the end of a calendar year. Simply taking the base-case solution for the first 92 weeks would produce an infeasible schedule because it would omit the backfill activities required before closure. The schedule optimized over the 92-week horizon precludes these infeasibilities, and moves a small quantity of metal forward in the schedule. However, with this schedule, the total metal produced is 8,070 tonnes lower than the amount scheduled over the same period in the base-case solution. Shortening the horizon does not increase scheduled metal, thus illustrating Lisheen's limited production flexibility. With no significant gains using the early-closure option, we turn our attention to the extended horizon scenario.

The 122-week extended mine-life scenario proves to be a more interesting case. Although scheduled metal production is 9,109 tonnes less than the production amount corresponding to the base case at week 104, the extended schedule includes 21,993 tonnes more metal than the base case by the end of its horizon. However, two additional periods of low metal production accompany this gain in metal. Although these production drops are not low enough to stop the mill, and implementing the schedule would delay mine closure, the results show that management should consider a tradeoff between (1) the gains from a longer mine life with higher total metal production, and (2) the early-closure risks associated with dips in metal production and potentially low metal spot prices in the future.

Although our model can solve a new scenario instance in less than a day, we require additional time to make data changes, prepare the model, and import our solution into iGantt. Using our approach, altering parameter values is an easy and quick process; however, changing precedence relationships between activities, a more involved task, can take a number of days to update. For example, a change in the cutoff grade would require engineers to reexamine and, in some cases, redefine block shapes, mining methods, and precedence relationships between activities. Upon completion of this study, we would require up to a week to make the corresponding changes. For a commercial application, we could eliminate this preparation time by writing a software interface to instantly transfer changes made in iGantt to our IP model. If we add this graphical user interface to the optimizer, nontechnical users could perform scenario analysis on production scheduling. However, because our academic scope includes limited scenario analyses, developing a commercial interface for our model would be tangential to our study.

Insights

With the IP results in hand, the planners at Lisheen can see that the current mine design and precedence relationships do not allow for significant metal gains from scheduling changes alone. In particular, the IP schedules highlight the valuable blocks that would always be excluded from an economic production schedule, and no amount of manual rearranging of activities in iGantt would bring these blocks forward into the schedule. This insight has prompted planners to explore alternative methods for capturing these high-value blocks. These methods can involve changing the engineering design (e.g., selecting a different mining method), developing new access, or even blasting a new haulage route through a previously backfilled area (see Figure 13).

For the planners, the IP solutions also highlight myopic and subjective decisions made during manual production scheduling. For example, the integer

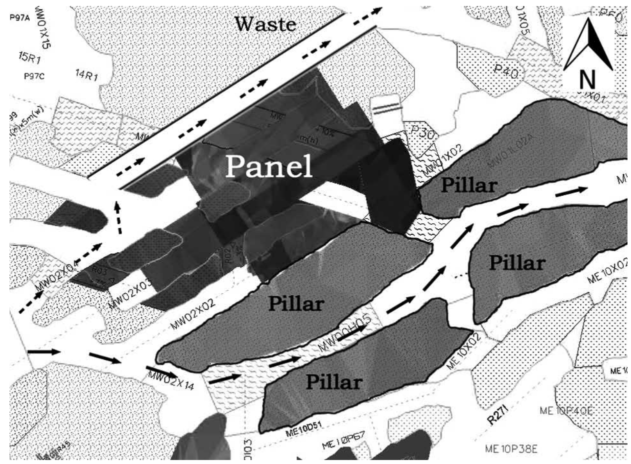

Figure 13: The IP schedule identifies four high-grade haulage pillars as blocks that would never be mined given the current mine design. The pillars form part of an existing main haulage route (solid arrows) connecting the crusher with a large ore zone to the west (not shown). Under the original plan, mining these pillars would only be possible once all activity in the western zone has finished. However, the IP results show that more profitable higher-grade pillars along the same haulage route would always be extracted in preference to these four pillars. Realizing this, Lisheen added a bypass (dashed arrows) to incorporate these pillars into the schedule.

program scheduled an area of very low-grade ore that Lisheen had determined was uneconomical to mine. Planners had overlooked that production from a number of very high-grade blocks in different panels was occasionally scheduled simultaneously, resulting in too high a head grade for ore feeding the mill. To maintain feasibility, low-grade ore is, at such times, required to balance the head grade. Additionally, a large decrease in ore production occurred at the start of the second year in the manual schedule. The IP schedule brought forward a retreat of haulage pillars, which were to be extracted late in the manual schedule to cover this deficit. The planner had not considered using these pillars to prevent the production decrease because of subjective reliance on a previous mine design. The mine configuration had since included new haulage routes and precedence rules, but the cognitive bias that those pillars could not be taken early remained in the planner's mind, preventing him from determining a solution for the shortfall in production.

Insights gained from the IP results have prompted Lisheen to take immediate action. In particular, the IP schedules have highlighted a number of panels that must be brought into production earlier to pull dependent blocks into the mining schedule. These areas are currently being mined in accordance with the IP schedule. Lisheen is in the process of changing the engineering design in many areas of the mine. When these designs are complete, we will incorporate the changes into the IP model and generate a new IP schedule. In this way, our work with Lisheen will continue to be an iterative process between mine design and schedule optimization.

Conclusions

Our IP approach produces near-optimal production schedules for a complex underground mine. Over a 104-week horizon, our schedule contains ore production that is more consistent with managerial goals than the mine's previous manual scheduling method was. Although both the manually produced and the IP-produced schedules are in iGantt for presentation purposes, our IP schedules add value to the mining operation by (1) shifting metal production forward in the schedule, (2) reducing waste mining and backfilling delays, (3) avoiding expensive mill-halting drops in ore production, and (4) enabling smoother workforce management. In addition, the integer program provides the mine with a scenario analysis capability, which has been instrumental in management's decision to revise much of the underground engineering design to increase the mine's operational life.

Although multiple open-pit optimization software packages are on the market, the idiosyncrasies of underground operations have, so far, made the creation of such a general software package for underground mining too challenging. However, the approach that we outline here can be adapted to other room-and-pillar underground operations, especially as they near the end of their operational lives.

Appendix. Model Formulation

The primary elements of the model are as follows.

Indices

a: Mining or backfilling activity, $a = 1, \ldots, n$.
$t, \hat{t}$: Period, $t = 1, \ldots, T$.

Sets

$\mathcal{A}^M$: Set of all extraction activities.
$\mathcal{A}^B$: Set of all backfilling activities.
$\mathcal{T}$: Set of all periods.
$\hat{\mathcal{T}}_a$: Restricted set of periods in which activity a can start.
$\bar{\mathcal{P}}_a$: Set of activities that must precede activity a.
$\tilde{\mathcal{P}}_a$: Set of activities, each of which must precede activity a, if it occurs.
$\hat{\mathcal{P}}_a$: Set of activities that must not precede activity a.

Parameters

$v_{at\hat{t}}$: Volume of ore obtained in period $\hat{t}$, given that we started extraction activity a at time t (tonnes).
g_a: Average percentage grade of the ore produced from extraction activity a.
$\bar{e}, \underline{e}$: Maximum and minimum allowable tonnage of ore excavated in a month, respectively (tonnes).
$\bar{g}, \underline{g}$: Maximum and minimum allowable metal produced by the mill in a month, respectively (tonnes).
t_a^m: Number of periods required for extraction activity a (months).
t_a^b: Number of periods required for backfilling activity a (months).
$p_{at\tilde{t}}$: Paste applied in period $\tilde{t}$, given that we started backfilling activity a at time t (cubic meters).
$\bar{p}$: Available paste for backfilling in each month (cubic meters).
r: Discount rate used to decrease the value of the metal produced in future periods.

Decision Variable

$$y_{at} = \begin{cases} 1 & \text{if activity } a \text{ starts during period } t, \\ 0 & \text{otherwise.} \end{cases}$$

Objective Function

$$\text{(P)} \qquad \max \quad \left\{ \sum_{a \in \mathcal{A}^M} \sum_{t \in \hat{\mathcal{T}}_a} \sum_{\hat{t} \in \mathcal{T}} g_a v_{at\hat{t}} y_{at} (1+r)^{-\hat{t}} \right\},$$

$$\text{subject to} \quad \sum_{t \in \hat{\mathcal{T}}_a} y_{at} \le 1 \quad \forall\, a \in \mathcal{A}^M \cup \mathcal{A}^B, \tag{1}$$

$$\underline{e} \le \sum_{a \in \mathcal{A}^M} \sum_{t \in \hat{\mathcal{T}}_a} v_{at\hat{t}} y_{at} \le \bar{e} \quad \forall\, \hat{t} \in \mathcal{T}, \tag{2}$$

$$\underline{g} \le \sum_{a \in \mathcal{A}^M} \sum_{t \in \hat{\mathcal{T}}_a} g_a v_{at\hat{t}} y_{at} \le \bar{g} \quad \forall\, \hat{t} \in \mathcal{T}, \tag{3}$$

$$\sum_{a \in \mathcal{A}^B} \sum_{t \in \hat{\mathcal{T}}_a} p_{at\tilde{t}} y_{at} \le \bar{p} \quad \forall\, \tilde{t} \in \mathcal{T}, \tag{4}$$

$$y_{at} \le \sum_{u \in \hat{\mathcal{T}}_{a'}:\, u \le t - t_{a'}^m - t_{a'}^b} y_{a'u} \qquad \forall\, a' \in \bar{\mathcal{P}}_a,\ a \in \mathcal{A}^M \cup \mathcal{A}^B,\ t \in \hat{\mathcal{T}}_a, \tag{5}$$

$$y_{at} \leq \sum_{u \in \hat{\mathcal{T}}_{a'}: u \leq t - t^m_{a'} - t^b_{a'}} y_{a'u} + \left(1 - \sum_{u \in \hat{\mathcal{T}}_{a'}} y_{a'u}\right) \quad \forall a' \in \tilde{\mathcal{P}}_a,\ a \in \mathcal{A}^M \cup \mathcal{A}^B,\ t \in \hat{\mathcal{T}}_a, \tag{6}$$

$$y_{at} \leq 1 - \sum_{u \in \hat{\mathcal{T}}_{a'}: u \leq t + t^m_a + t^b_a} y_{a'u} \quad \forall a' \in \hat{\mathcal{P}}_a,\ a \in \mathcal{A}^M \cup \mathcal{A}^B,\ t \in \hat{\mathcal{T}}_a, \tag{7}$$

$$y_{at} \text{ binary} \quad \forall a \in \mathcal{A}^M \cup \mathcal{A}^B,\ t \in \mathcal{T}. \tag{8}$$

We present the IP formulation for the original problem at monthly fidelity. The decision variables dictate whether we begin an activity during a specified period. In theory (i.e., in iGantt), these activities are scheduled at the beginning of the period. However, because most activities do not consume an integral number of periods to complete, in practice, the activity can begin at any point during the period in which it is scheduled to start, as long as it is finished in an integer number of periods no larger than the ceiling of the activity's duration. Our variable definition necessarily discretizes time, and at a fidelity that is arguably coarser than that at which the mine would implement the schedule. However, we require binary variables to enforce the precedence constraints, and the number of these variables is proportional to the product of the number of activities and the number of periods we consider. A finer time fidelity (approaching a continuum) might more closely approximate the way in which mine operations are conducted, but would render the model so intractable that it would produce no results quickly enough to be practical. Although not mathematically portrayed here, we can reduce the number of variables (from the product of the number of activities and the number of periods we consider) by employing early start dates for each activity, as discussed in the Integer Programming Approach section.

We maximize the discounted value of extracted metal over the horizon, where we discount consistent with the period in which we extract the metal. Note that we effectively assume an arbitrary constant metal price (given in dollars or euros per tonne), which we can omit from the objective without changing the optimal solution. As a result, our objective is given in tonnes, rather than in a monetary unit. Constraints (1) ensure that we never schedule a mining or backfilling activity more than once. Constraints (2) require that the scheduled production in any period is no more than the production capacity and no less than what the mill requires to operate for that period. Constraints (3) force the model to maintain metal output within the operational limits of the mill. Constraints (4) track the use of paste fill in each period and ensure that its use does not exceed its availability. Constraints (5)–(7) are the sequencing constraints that enforce precedence rules between activities.

The first set of sequencing constraints, (5), ensures that an activity a cannot begin until all of its predecessor activities $a' \in \bar{P}_a$ have completed. For example, block A must be mined to access block B, and block B must be mined to access block C. Therefore, we cannot mine block C without mining block A and block B first. The second set of sequencing constraints, (6), ensures that the order of mining in a panel be maintained, regardless of whether all predecessors of a block $a' \in \tilde{P}_a$ have been extracted. In this example, blocks A and B do not both need to be mined if we mine block C. Assume, without loss of generality, that we mine only block A. Then, we only impose the constraint that this block must be mined before block C. Although all of the blocks' precedences were initially characterized according to constraints (5), relaxing many to follow constraints (6) allows for a feasible, yet better, solution. Sequencing constraints (7) ensure that an activity a cannot occur once a haulage pillar $a' \in \hat{P}_a$, on which the activity depends, is extracted.

We model our integer program using the AMPL programming language, version 20090327 (Fourer et al. 2003) and solve it with the CPLEX solver, version 12.3 (IBM 2011) on a machine with two dual Intel Xeon X5570 quad core processors and 48 GB of RAM.

Acknowledgments

We thank the Lisheen Mine for providing us with this project and supporting the research. In particular, we thank Tom Bailey and Noeleen Fox for their time and commitment. We also thank Brian Hall of AMC Consultants for his comments and advice. Finally, we thank Dr. Tim Kaiser of the High Performance Computing Group at the Colorado School of Mines for his resources.

References

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

Chanda E (1990) An application of integer programming and simulation to production planning for a stratiform ore body. *Mining Sci. Tech.* 11(2):165–172.

Fourer R, Gay DM, Kernighan BW (2003) *AMPL: A Modeling Language for Mathematical Programming* (Thompson Learning, Pacific Grove, CA).

Hamrin H (1997) *Guide to Underground Mining Methods and Applications* (Atlas Copco, Stockholm).

IBM (2011) ILOG CPLEX. Accessed December 1, 2012, http://www-01.ibm.com/software/commerce/optimization/cplex-optimizer/index.html.

Klotz E, Newman AM (2013) Practical guidelines for solving difficult mixed integer linear programs. *Surveys Oper. Res. Management Sci.* 18(12):18–32.

Maptek (2012) Vulcan. Accessed December 1, 2012, http://www.maptek.com/products/vulcan.

Martinez MA, Newman AM (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.

MineMax (2012) MineMax iGantt. Accessed December 1, 2012, http://www.minemax.com/downloads/Minemax-iGantt-Brochure.pdf.

Newman A, Kuchta M (2007) Using aggregation to optimize long-term production planning at an underground mine. *Eur. J. Oper. Res.* 176(2):1205–1218.

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

Smith M (1998) Optimizing short-term production schedules in surface mining: Integrating mine modeling software with AMPL/CPLEX. *Internat. J. Surface Mining Reclamation Environ.* 12(4):149–155.

Trout L (1995) Underground mine production scheduling using mixed integer programming. *25th Internat. Sympos. Appl. Comput. Mineral Indust (APCOM), Brisbane, Australia*, 395–400.

Williams JK, Smith L, Wells PM (1973) Planning of underground copper mining. *10th Internat. Appl. Sympos. Appl. Comput. Mineral Indust. (APCOM), Johannesburg, South Africa*, 251–254.

Winkler B (1996) Using MILP to optimize period fix costs in complex mine sequencing and scheduling problems. *26th Internat. Appl. Sympos. Appl. Mineral Indust. (APCOM), University Park, PA*, 441–446.

Verification Letter

Tom Bailey, Chief Technical Services Engineer Lisheen Mine, Killoran, Moyne, Thurles, Co. Tipperary, Ireland, writes:

"Manually scheduling ore production is a difficult and labor intensive part of our mining operation and usually takes three or more weeks to complete. Maximizing metal extraction over the life of the mine becomes increasingly challenging as the number of available working areas declines. The production profile eventually decreases to a point beyond which it is no longer profitable to operate, and this defines the life-of-mine (LOM). At the end of the LOM, there is still ore remaining in the mine at varying Zinc grades. The optimizer generates a mining schedule that maximizes the metal extracted over the LOM, whilst minimizing the mineral resources left in the ground after closure.

"We are currently using the optimizer to assist in our production planning for the last two years of the mining operation at Lisheen. The optimized schedules generally support the decisions encoded in the manual schedule. However, in addition to the primary objective of maximizing metal extraction, the optimization tool provides us with the following benefits: (i) a way to quickly examine alternative end-of-mine-life scenarios, e.g., different mine closure dates (ii) less change in the production tonnage from one week to the next, enabling us to generate more consistent labor shifts; and (iii) the identification of manually scheduled infeasibilities, e.g., too much high-grade ore scheduled for simultaneous extraction.

"The optimizer has been particularly useful in identifying areas that potentially will not be mined no matter how quickly we mine in the neighboring vicinity. With this information now to hand, we have initiated new plans to accelerate the extraction in these areas. This has resulted in c.120kt at 13 percent ZnEq of extra ore now being included in our revised LOM Schedule.

"We will continue to employ the optimizer both to pull forward high grade areas and to identify areas that need additional mining to free them for mining before the mine closes."

Dónal O'Sullivan earned a BA in economics and a BA in engineering from Brown University, and he has completed a PhD in mineral and energy economics at the Colorado School of Mines. Before graduate school he specialized in electric power consulting and power market modeling. While at Mines, he taught an undergraduate course in economics and technology. He is an economic consultant with Galt and Company.

Alexandra Newman is an associate professor in the Division of Economics and Business at the Colorado School of Mines. She holds a BS in applied mathematics from the University of Chicago, an MS in operations research from the University of California, Berkeley, and a PhD in industrial engineering and operations research from the University of California, Berkeley. Prior to joining the Colorado School of Mines, she held the appointment of research assistant professor in the Operations Research Department at the Naval Postgraduate School. Professor Newman's primary research interests lie in optimization modeling, particularly as it is applied to logistics, the energy and mining industries, and military operations. She is serving as associate editor of *Operations Research* and *Interfaces*, and is an active member of INFORMS.