

Implementing a Production Schedule at LKAB's Kiruna Mine

Mark Kuchta

Mining Engineering Department, Colorado School of Mines, 1500 Illinois Street, Golden, Colorado 80401, mkuchta@mines.edu

Alexandra Newman

Division of Economics and Business, Colorado School of Mines, 1500 Illinois Street, Golden, Colorado 80401, newman@mines.edu

Erkan Topal

Mining Engineering Department, Dicle University, Diyarbakir, Turkey, etopal@dicle.edu.tr

LKAB's Kiruna Mine, located in northern Sweden, produces about 24 million tons of iron ore yearly using an underground mining method known as sublevel caving. To efficiently run the mills that process the iron ore, the mine must deliver planned quantities of three ore types. We used mixed-integer programming to schedule Kiruna's operations, specifically, which production blocks to mine and when to mine them to minimize deviations from monthly planned production quantities while adhering to operational restrictions. These production schedules save costs compared to schedules produced manually by meeting desired production quantities more closely and reducing employee time spent on preparing schedules.

Key words: production and scheduling; applications.

History: This paper was refereed.

Planning an underground mine is a complex procedure consisting of five stages: (1) determining the geometry and grade (or quality) distribution of the ore body, (2) deciding how to mine the ore, that is, by surface or underground mining, (3) designing the mine infrastructure, that is, how to lay out the mine to mine and retrieve ore efficiently, (4) planning how to mine and process the ore, and finally (5) decommissioning the mine and restoring the site to an environmentally acceptable state. The mining and processing (or *production*) phase requires detailed scheduling. Specifically, a production schedule must provide a mining sequence that takes into account the physical limitations of the mine and, to the extent possible, meets the demanded quantities of each raw ore type at each time period throughout the life of the mine. Mines use the schedules as long-term strategic-planning tools to determine when to start mining a production area and as short-term operational guides.

LKAB's Kiruna Mine, located above the Arctic Circle in northern Sweden, is the second largest underground mine in the world today. The ore body is a high-grade magnetite deposit approximately four kilometers long and about 80 meters wide on average, and it lies roughly in a north-south direction with a dip of about 70 degrees from the horizontal plane. In 1878, English metallurgists Sidney Thomas and Perry Gilchrist discovered how to process high-quality steel from iron ore with a high

phosphorous content, and the economic value of the Kiruna site became evident. In about 1890, the company Loussavaara-Kiirunavaara Aktiebolag (LKAB) was formed, and eight years later, it began mining operations at Kiruna. Today, the company employs about 3,000 workers, and the Kiruna Mine produces approximately 24 million tons of iron ore per year.

The two main in situ ore types differ in their phosphorous content. About 20 percent of the ore body contains a very high-phosphorous (P), apatite-rich magnetite known as *D ore*, and the other 80 percent contains a low-phosphorous, high-iron (Fe) content magnetite known as *B ore*. For the entire deposit, the best quality B ore is about 0.025 percent P and about 68 percent Fe. The D ore varies considerably and has average grades of about two percent P.

From the two main in situ ore types, the mine obtains three raw ore types used to supply four post-processing plants, or mills (Table 1). Phosphorus is the main ore contaminant and the amount of phosphorous present in the ore determines the ore type. The B1 type contains the least phosphorus, and the mills transform the ore into high-quality fines (of the granularity of fine sand) simply by crushing and grinding the ore and removing the contaminants using magnetic separation. B2 ore contains somewhat more phosphorous, and D3 ore has the highest phosphorous content. The mills process both B2 and D3 into ore pellets approximately spherical in shape by crushing and grinding the ore into a

Type	Percent Fe	Percent P	Use
B1	68.0	0.06	Fines production
B2	—	0.20	Medium phosphorous-content pellets production
D3	—	0.90	High phosphorous-content pellets production

Table 1: From the two main in situ ore types, the mine obtains three raw ore types. The B1 type contains the least phosphorous (contaminant) and is used to produce high-quality fines (of the granularity of fine sand). B2 ore contains a medium amount of phosphorous, and D3 ore has the highest phosphorous content. The mine processes both B2 and D3 ores into ore pellets, approximately spherical in shape.

finer consistency than the B1 ore, then adding binding agents and other minerals, such as olivine and dolomite, and finally firing the resulting product in large kilns to form hard pellets. Pelletizing plants not equipped with flotation circuits remove excess phosphorus to manufacture pellets from B2 ore, and pelletizing plants equipped with such flotation circuits produce pellets from D3 ore. The mine generally cannot extract the B2 type directly from the ore in situ but rather produces it almost entirely by blending high-phosphorous D ore with low-phosphorous B ore during extraction (Kuchta 1999).

Trains transport the fines and pellet products from the mills to harbor facilities in Narvik, Norway and Luleå, Sweden. The company ships most of the production to steel mills in Europe but ships some product to the Middle East and to the Far East. Iron ore fines and pellets are used as raw materials in the manufacture of various steel products, such as kitchen appliances, automobiles, ships, and buildings.

The method a mine uses to extract ore depends on how deep the deposit lies and on the geometry of the deposit, as well as on the structural properties of the overlying and surrounding earth. Mining companies use surface, or *open pit* mining when deposits are fairly close to the surface. They first remove the surface soil and overlying waste and then recover the ore by drilling and blasting. As the pits deepen, they may need to decrease the slopes of the pit walls to avoid pit failure, that is, waste material sliding down into the active area of the mine. The deeper the pit, the more complicated the haulage routes become and the higher the probability of encountering underground water. Eventually the pit becomes too costly to operate. The mining company then either shuts down the mine or begins mining underground.

The variety of underground mining methods are categorized as self-supported methods, supported methods, and caving methods. Kiruna currently uses sublevel caving, an underground caving method applied to vertically positioned, fairly pure, large, vein-like deposits. Miners first drill ore passes that extend vertically from the current mining area down

to the bottom of a new mining area where a transportation level is located. They then create horizontal sublevels on which to mine and access routes that run over the length of the ore body within a sublevel. Finally, miners drill self-supported horizontal crosscuts through the ore body perpendicular to the access routes. Kiruna spaces sublevels about 28.5 meters apart and spaces crosscuts about 25 meters apart. The crosscuts are seven meters wide and five meters high. From the crosscuts, miners drill near-vertical rings of holes in a fan-shaped pattern. Each ring contains around 10,000 tons of ore and waste. The miners place explosives in the holes and blast the rings in sequence, destroying the ceiling on the blasted sublevel, to recover the ore. Miners recover the ore on each sublevel, starting with the overlying sublevels and proceeding downwards. Within each sublevel, they remove the ore from the hanging wall to the forefront of the mining sublevel, or the foot wall. As the miners recover the ore from a sublevel, the hanging wall collapses by design and covers the mining area with broken waste rock (Figure 1). Initially, Kiruna used surface-mining methods, but in 1952, it began underground mining operations (Figure 2).

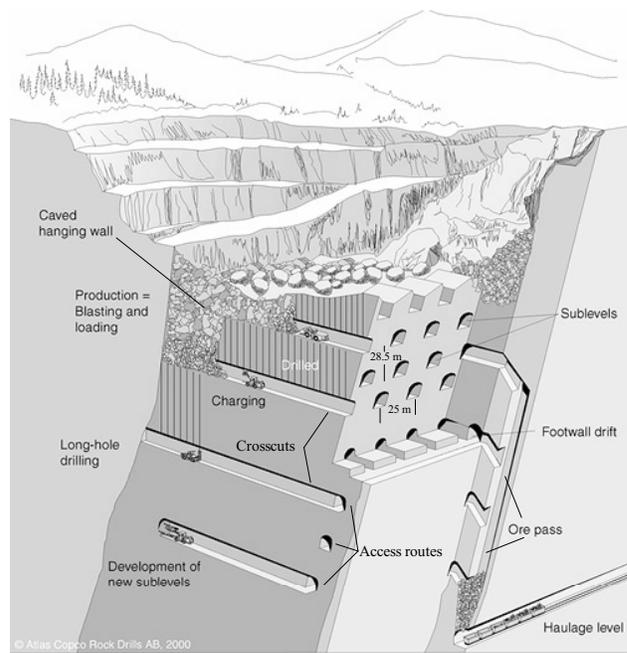


Figure 1: Depicted here in this sublevel caving operation are an ore pass extending vertically down to horizontal mining sublevels, access routes running the length of the ore body within a sublevel, and crosscuts drilled perpendicular to the access routes. Kiruna geometries are superimposed on the figure, that is, the spacing between crosscuts is 25 meters, between sublevels is 28.5 meters, and between access routes is 28.5 meters. Within each sublevel, ore is removed from the hanging wall to the foot wall, after which the hanging wall on that sublevel collapses into the working area by mine design. Railcars transport the ore from the mined area to a crusher. (Source: Atlas Copco 2000)

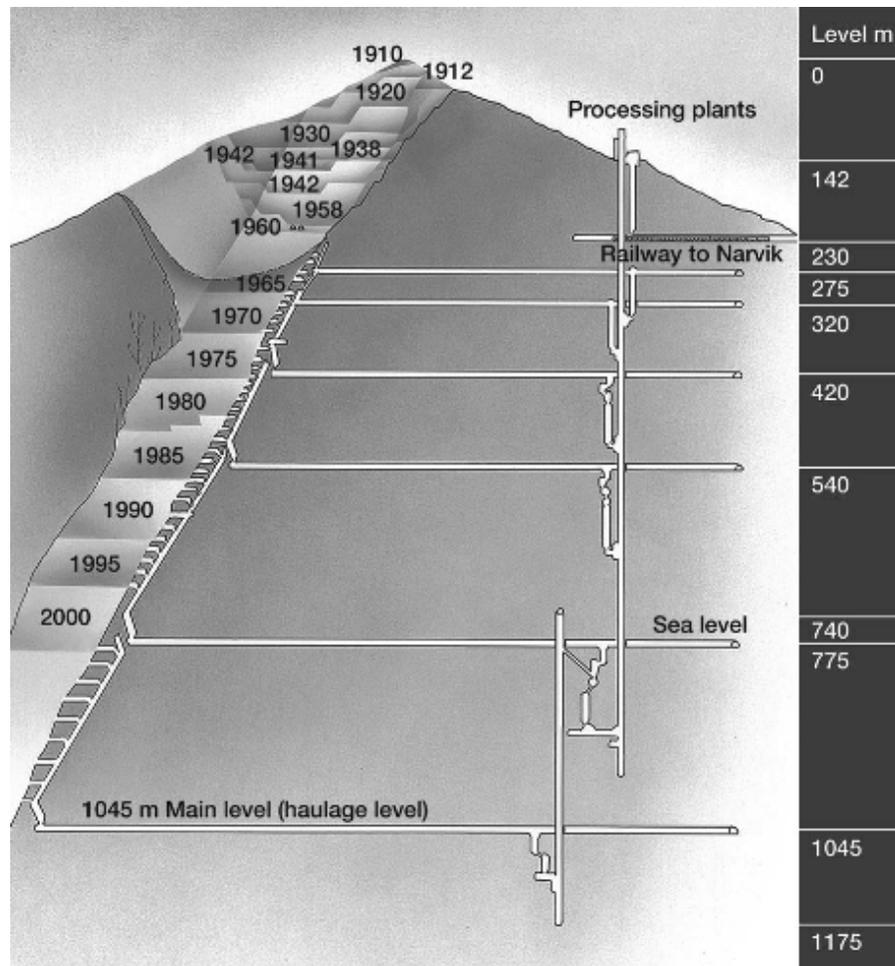


Figure 2: Kiruna began its mining operations around 1900 using surface methods. In 1952, underground operations began. Today, Kiruna is exclusively an underground mine. Miners are currently extracting ore from the 1,045 meter level. (Source: LKAB 2001)

The mine is divided into 10 main production areas, extending from the uppermost mining level down to the current main 1,045 meter level. These production areas are about 400 to 500 meters in length, each with its own group of ore passes, also known as a *shaft group*, located at the center of the production area and extending down to the 1,045 meter level. One or two 25-ton-capacity electric load-haul-dump units (LHDs), that is, vehicles that load, transport, and unload the ore, operate on a sublevel within each production area, and transport the ore from the crosscuts to the ore passes. Large trains operating on the 1,045 meter level transport the ore from the ore passes to a crusher, which breaks the ore into pieces four inches or less in size for subsequent hoisting to the surface through a series of vertical shafts (Figure 3).

The site on which each LHD operates is also referred to as a *machine placement*. Depending on production requirements, up to 18 LHDs can be operating daily in various parts of the mine. Each machine

placement is usually 200 to 500 meters long and contains from one to three million tons of ore, equivalent to between 10 and 12 *production blocks*, which are the same height as the mining sublevel (about 28.5 meters) and extend from the hanging wall to the foot wall. Once the mine has started mining a production block, to conform to mining restrictions, it must maintain continuous production of the production blocks within a machine placement until it has removed all the available ore.

Kiruna calculates iron ore reserves contained in each machine placement in two ways. Using the first method, it calculates the reserves for long-term strategic planning in the undeveloped areas of the mine. Geologists estimate the quantities of B and D ore from an in situ geologic block model. Geologists develop this model by drilling and sampling the deposit and then extrapolating from the samples to estimate the ore types present. Next, mine planners use the quantities of B and D ore from the in situ block model

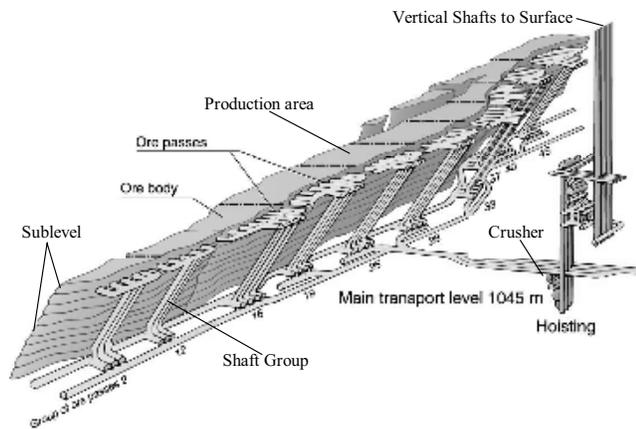


Figure 3: Depicted here is the ore body, consisting of various production blocks and with ore passes located at the center of each production area. Trains transport the ore from the ore passes to a crusher, which breaks the ore into pieces of four inches or less, a size that can be lifted by a bucket on a rope to the surface through a series of vertical shafts. (Source: LKAB 2001)

that the geologists develop to calculate the expected quantities and grades of the three ore types, B1, B2, and D3, for production blocks 100 meters in length, extending from the hanging wall to the foot wall, and the height of the planned sublevel. In their calculation, the planners account for blending the raw ores and the gravity flow of broken rock (Kuchta 1999).

Kiruna uses the second method of calculating iron ore reserves in each machine placement when it has more information about the production blocks, specifically, when it has finished developing a production area. It can calculate the in situ tons and grades and hence the expected tons of B1, B2, and D3 ore ring-by-ring, making a more precise estimate of production quantities. Mine planners use these calculations, in turn, to produce monthly production estimates for each machine placement (Figure 4).

Once a month mine planners update their estimates of the tons available for all active machine placements using the mine database and postprocessing system (Kuchta and Engberg 2002). These estimates give the mine scheduler appropriate “initial conditions” with which to begin planning the next month’s schedule.

Early applications of optimization to problems in the mining industry concern open-pit mining (Lerchs and Grossman 1965, Wilke and Reimer 1977, Underwood and Tolwinski 1998). Williams et al. (1973) and Jawed (1993) used linear programs to plan sublevel stoping (a self-supported underground mining method) in a copper mine and room and pillar (also a self-supported underground mining method) in a coal mine, respectively.

Linear-programming models cannot capture discrete decisions about whether to mine a production block or not. Tang et al. (1993), Tolwinski and Underwood (1996), and Winkler (1998), among others, combined linear programming with simulation or used heuristics to address discrete decisions. Kuchta (2002) developed a manual scheduling model, essentially a computer-aided heuristic, that the Kiruna mine used. The operator scheduled when production blocks should be mined, and the model tracked the outcomes of these decisions, updating block availability and the quantities of the various ore types Kiruna mined per time period as necessary.

Although these models are attempts to capture discrete decisions, none produce solutions that are guaranteed to be of reasonable (near-optimal) quality. Winkler (1996) pointed out the importance of directly capturing discrete decisions and associated constructs (for example, fixed costs and logical conditions) with exact solution methods, that is, mixed-integer-programming models. However, because of the large size of such scheduling models and hardware and software limitations, he declared that the theoretical complexity of mixed-integer programs precludes their use for multiperiod mine scheduling. Others lent credibility to Winkler’s statement, for example, Winkler and Griffin (1998) and Smith (1998), in trying to solve a model for a silver and gold surface mine, and Trout (1995), when attempting to solve a mixed-integer multiperiod production-scheduling model for underground stoping operations for a base metal (copper sulphide). Several researchers in prior attempts at Kiruna mine failed to produce production schedules of requisite length in a reasonable amount of time (Almgren 1994, Topal 1998, Dagdelen et al. 2002). Instead of solving for an optimal schedule, they resorted to shortening the schedule time frame or sacrificing schedule quality. Because of these shortcomings, Kiruna never adopted these schedules. The closest work to ours was an integer-programming model to plan a production schedule for a sublevel stoping operation at Stillwater Mining Company (Carlyle and Eaves 2001). The model provides near-optimal solutions to maximize revenue from mining platinum and palladium; however, the authors do not describe any special techniques to expedite solution time.

We state our problem as follows: Given monthly demands for the three ore types, when do we start to mine each production block (or machine placement), each containing specified quantities of the three ore types, to minimize deviation from these monthly demands subject to mine operational restrictions?

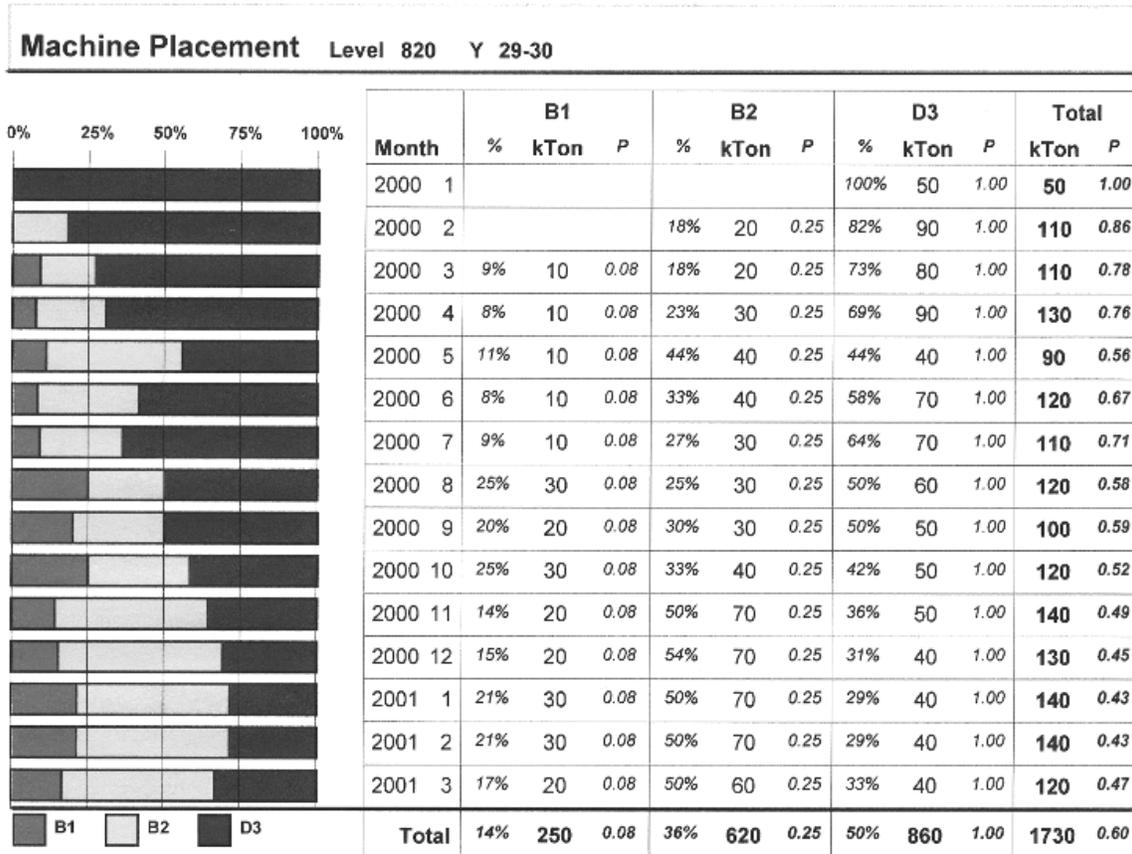


Figure 4: Mine planners use estimates for 100 meter blocks to generate data covering expected monthly production for each machine placement. Shown here are production estimates from January 2000 to March 2001 of B1, B2, D3 and total production estimates for a block at the 820 meter level from y-coordinate 29 to y-coordinate 30 (which spans a distance of 200 meters). Each column corresponds to an ore type, and each row to a month's worth of production. Each cell contains information for a month and the combination of production quantities for each ore type, specifically, the percentage of that ore type present, the number of kilotons, and the percentage of phosphorous.

Model Description and Previous Manual Scheduler

To create a production schedule, the planner must determine the start dates for the various machine placements such that the mine can produce the tons of B1, B2, and D3 ores required each month. The mine supplies one mill with B1 ore, two mills with B2 ore, and one mill with D3 ore. Because it can stockpile only about 6,000 tons of ore each day, the mine must meet production demands at the four mills almost exactly so that the mills can meet their requirements for production. We minimize the deviation from the specified production levels for each ore type in each month. Moreover, the mine must observe the following operational constraints:

- The amount of each ore type mined in each month minus surplus and plus shortage must equal the demand for each ore type.

- Vertical mine-sequencing constraints preclude mining machine placement *b*, which is under machine placement *a*, until at least 50 percent of machine placement *a* has been mined.

- Horizontal mine-sequencing constraints require that machine placements adjacent to a given machine placement and on the same sublevel be mined after 50 percent of the given machine placement has been mined.

- Shaft-group constraints restrict the number of active LHDs within a shaft group at any one time to a predetermined maximum, usually two or three.

Rullplan, the program developed for scheduling production manually according to these specifications, is a database application written in Microsoft Access 97. It includes a user interface consisting of various forms for data entry and program control, and it produces reports. All data is stored in the mine's central relational database, and a schematic overview tracks available machine placements by

shaft group and mining sublevel throughout the relevant planning horizon.

The scheduling process can be characterized as a computer-assisted manual heuristic. The scheduler first establishes production targets for the three raw ore types for each month within the planning horizon. The scheduler initializes a five-year schedule by adding to the schedule all active machine placements. As the miners deplete the ore at the active machine placements, they can no longer meet production targets. The scheduler then selects an available machine placement that he or she thinks will best meet demand while adhering to mine-sequencing constraints. The scheduler adds the machine placement to the schedule by entering a start date for that machine placement. The scheduling program then assigns start dates for all the production blocks within that machine placement according to the constraints on sequencing and shaft groups and displays the ore tonnages that will result. The scheduler continues to assign start dates for selected available machine placements month by month until he or she obtains a five-year schedule. Using Rullplan, the scheduler takes five days to devise a five-year schedule.

Furthermore, these schedules are clearly myopic, that is, they do not incorporate the effects on availability of machine placements even a few time periods into the future. Without foresight, the scheduler may produce schedules that are far from optimal and may even be infeasible. In some instances, the scheduler backtracks and chooses different start dates for machine placements to induce feasibility. As the scheduler works on scheduling periods late in the time horizon, however, this effort becomes more costly and has an increasingly small chance of producing a feasible schedule. Therefore, a final schedule may easily contain infeasibilities, especially in the “out-years.”

Current Mathematical Programming Scheduler at Kiruna Mine

Even for instances in which heuristic (manual) algorithms produce useable production schedules, schedulers have no easy way to judge the quality of these schedules relative to the best (for example, the cost-minimizing) schedule. We use a mathematical-programming technique, mixed-integer programming (MIP), to produce optimal production schedules for underground mines. The use of MIP has been hindered because models of real-world problems must often incorporate a large number of decision variables, many of which must assume integer values. Because of the large number of integer variables, solution times may be unacceptably long for practical

planning purposes. By preprocessing the production data and formulating the model carefully, we reduced the number of integer variables in our multiperiod production model and thus greatly reduced solution times.

The main advantage of our formulation over previous attempts at Kiruna was this reduction in the number of variables and the resulting dramatic improvement in model tractability. By developing a new database and formulating the model carefully, we aggregated perhaps 12 production blocks into a single machine placement. Specifically, we can replace the binary variable indicating whether production block b is mined in time period t with a binary variable indicating whether machine placement a starts to be mined in time period t . For a five-year horizon scheduled month by month, we reduced the number of binary variables in our model from about 60 (time periods) * 1,100 (production blocks) = 66,000 binary variables, to about 60 (time periods) * 60 (machine placements) = 3,600 binary variables.

We can further reduce the number of integer variables by assigning earliest and latest possible start dates to machine placements based on the logic that (1) because of sequencing and shaft-group constraints, Kiruna cannot start mining a machine placement before it starts mining the requisite number of machine placements surrounding it, and (2) based on demand constraints and bounds on a reasonable amount of deviation between demand and production, Kiruna must start mining a machine placement early enough that it does not lock in underlying machine placements, preventing production of the required amount of ore. (Newman and Kuchta 2003 give details.)

When the scheduler makes further modifications (adding tightening constraints and active machine placements) and runs the model with appropriate hardware and software, he or she produces a near-optimal schedule for a five-year time horizon in minutes (Appendix).

Results

Over a five-year horizon, we obtain a near-optimal integer-programming solution. Figure 5 depicts the first year of such a five-year schedule. In this solution the ratio of the tons of iron ore mined constituting a deviation from planned production (that is, the amount over or under the amount planned) to the total tons of iron ore mined is less than five percent. This ratio is significantly higher, perhaps 10 to 20 percent, for the manually generated schedules. It is difficult to accurately compare the solution quality of the schedules prepared manually and automatically because despite the fact that deviations from

Level Y	Shaft		201	202	203	204	205	206	207	208	209	210	211	212	total
765 42- 43	45	1													690
765 44- 46	45	2													1680
792 14- 15	16	1													1630
792 19- 21	19	2													130
792 33- 34	33	2													540
792 35- 37	37	1													420
792 39- 42	40	1													1760
792 43- 46	45	2													750
818 16- 17	16	2													650
820 16- 17	16	2													150
820 18- 19	19	1													1160
820 23- 26	25	1													1120
820 27- 28	28	1													1250
820 31- 32	33	1													1650
820 33- 34	33	2													1600
820 35- 37	37	1													1690
820 38- 41	40	1													490
849 27- 28	28	1													130
849 29- 30	28	2													1210
849 31- 34	33	1													820
Devel. ore															1560
Devel. waste															996
kton total			1913	1863	1973	1913	1843	1783	1863	1953	1813	1683	1673	1803	22076
kton/day total			61.7	66.5	63.6	63.8	59.5	59.4	60.1	63.0	60.4	54.3	55.8	58.2	
kton/day B1			11.0	10.7	11.0	11.3	10.3	11.0	11.0	11.3	10.8	10.0	10.0	9.0	
kton/day B2			27.5	29.4	27.2	28.8	27.2	28.8	28.2	28.8	27.3	22.4	22.1	25.3	
kton/day D3			23.2	26.4	25.5	23.7	21.9	19.7	21.0	22.9	22.3	21.9	23.7	23.9	
month			201	202	203	204	205	206	207	208	209	210	211	212	

Figure 5: This figure depicts the first year of a complete five-year schedule obtained with the optimization model. The row headings specify, for each machine placement in the schedule, the level, the y-coordinates (or horizontal span of distance), and the shaft-group number. The column headings give the year and month of the schedule. For example, the heading 201 represents January 2002. Each cell graphically depicts the monthly amounts of the three ore types, B1, B2, and D3, contained in each machine placement in the production schedule. At the bottom of the figure are monthly production totals.

planned production are higher for the manually generated schedules, they often violate mine-sequencing constraints. However, the deviations that the manually generated schedules imply serve only as a lower bound on the deviations that actually exist, rendering the manually generated schedules even less desirable than they would appear.

Because the mine does not stockpile iron ore, if it produces less ore than desired, the mills are forced

to produce less final product, and they lose sales. Because the mills operate at a constant rate, Kiruna must often leave excess ore in the mine until the mills can process it. Therefore, we can estimate cost savings as the tons of absolute deviation in desired ore quantities multiplied by the current profit per ton of ore.

Using our model, the scheduler can obtain a complete five-year schedule with integer programming in 300 seconds. Model preparation is not time consum-

ing, as the required data are readily available and can easily be imported as data files. Furthermore, the scheduler can work on other projects while the model is running. The entire activity of preparing the schedule takes only a few hours at most. By contrast, creating a five-year schedule manually takes about five man-days; overall, the scheduler spends about 25 percent of the time preparing these long-term schedules and the shorter-term monthly and yearly schedules. Therefore, we estimate the savings in cost for the time spent generating schedules as about 25 percent of the scheduler's salary.

Finally, producing a single schedule takes so much time that schedulers seldom produce alternate schedules. However, mine planners may be interested in various production schedules so that they can plan for changing demands or other contingencies. Planners might also be interested in the ramifications of alternate mining strategies related to designing mine infrastructure.

Kiruna planners have integrated our production-scheduling system with the mine's existing computer-assisted manual planning system, Rullplan. Our integer program provides long-term strategic schedules, which mine planners generate monthly and use without modification over the two- to three-year period during which mine operators develop a production area. The schedules insure that the machine placements required for production are ready when needed. Mine planners also base the next short-term production schedule on the first month of the long-term schedule. For this purpose, they may adjust the integer-programming schedule or the actual production quantities for the coming month to account for unpredictable events, such as a sudden change in demand at the mills or the breakdown of a LHD.

Conclusions

Our optimization model for long-term production scheduling at LKAB's Kiruna Mine uses a new database with a new block-data format for which we preprocess production data for a mining area into monthly production quantities for the three raw ore types. With manual methods, it is difficult to visualize the interactions among production areas far ahead in time, so planners commonly abandon partially completed schedules and start over from the last good starting point, that is, a point in the schedule at which no operational constraints have been violated and no constraint violations seem imminent. With backtracking, a planner can take a week or more to develop a complete five-year schedule. We developed an exact solution procedure using integer programming, which cut the time needed to generate a schedule, and produces schedules of high quality.

Extensions to this work have included developing methodology to reduce the solution time for large problems (Newman and Kuchta 2003). Mine planners now want the ability to develop short-term (that is, monthly) production schedules with time fidelity of days. The challenges lie in developing a tractable model and in integrating this model with the long-term model we present in this paper.

Appendix: Model Formulation

The formulation follows:

Indices

- a = machine placement.
- b, b' = production block.
- k = ore type, i.e., B1, D2, D3.
- t = time period (month).
- v = shaft group, i.e., 1, ..., 10.

Sets

T_b = set of eligible time periods in which production block b can be mined (restricted by production block location and the start date of other relevant production blocks).

B_t = set of eligible production blocks that can be mined in time period t .

B_a = set of production blocks within machine placement a .

B_b^v = set of production blocks whose access is restricted vertically by production block b .

B_b^R = set of production blocks whose access is forced by right adjacency to production block b .

B_b^L = set of production blocks whose access is forced by left adjacency to production block b .

A_v = set of machine placements contained in shaft group v .

Parameters

r_{bk} = amount of ore type k in block b (tons).

d_{kt} = demand for ore type k in time period t (tons).

\underline{t}_b = earliest start date for production block b .

\bar{t}_b = latest start date for production block b .

T = length of the planning horizon.

LHD_v = the maximum number of simultaneously operational LHDs in each shaft group v .

$$P_{abv} = \begin{cases} 1, & \text{if block } b \text{ of machine placement } a \\ & \text{is in shaft group } v, \\ 0, & \text{otherwise.} \end{cases}$$

Decision Variables

$$y_{bt} = \begin{cases} 1, & \text{if we start mining production block } b \\ & \text{in time period } t, \\ 0, & \text{otherwise.} \end{cases}$$

\bar{z}_{kt} = amount mined above the desired demand of ore type k in time period t (tons).

z_{kt} = amount mined below the desired demand of ore type k in time period t (tons).

Objective Function

$$\begin{aligned} & \text{Min} \sum_{k,t} z_{kt} + \bar{z}_{kt} \\ & \text{Subject to} \sum_{b \in B_t} r_{bk} y_{bt} + z_{kt} - \bar{z}_{kt} = d_{kt} \\ & \qquad \qquad \qquad \forall k \text{ and } t \in T_b, \quad (1) \\ & \sum_{t \in T_b} y_{bt} \geq y_{b't'} \quad \forall b, b' \in B_b^v, t' \in T_{b'}, b \neq b', \quad (2) \\ & \sum_{t \in T_b} y_{bt} \leq y_{b't'} \quad \forall b, b' \in B_b^R, t' \in T_{b'}, b \neq b', \quad (3) \\ & \sum_{t \in T_b} y_{bt} \leq y_{b't'} \quad \forall b, b' \in B_b^L, t' \in T_{b'}, b \neq b', \quad (4) \\ & \sum_{a \in A_v} \sum_{b \in B_a} \sum_{t \in T_b} P_{abv} y_{bt} \leq \text{LHD}_v \quad \forall v, \quad (5) \\ & \bar{z}_{kt}, z_{kt} \geq 0 \quad \forall k, t, \\ & y_{bt} \text{ binary} \quad \forall b, t. \end{aligned}$$

We minimize the deviations from the planned quantities of B1, B2, and D3 ores for each month. Constraints (1) track the tons of B1, B2, and D3 ore mined per time period and the corresponding deviations from the specified production levels. Constraints (2) control vertical sequencing between mining sublevels. Constraints (3) and (4) enforce horizontal sequencing between adjacent production blocks. Constraints (5) ensure that no more than the allowable number of LHDs is active within a shaft group. Finally, we enforce nonnegativity and integrality of variables as appropriate.

This is the formulation of the original (intractable) model. The improved model differs as follows: rather than defining a separate binary variable indicating whether production block b is mined in each time period t , i.e., y_{bt} , as presented in the formulation above, we define a binary variable indicating whether machine placement a starts to be mined in time period t , i.e., y_{at} . We can make this variable change because all production blocks within a machine placement must be mined continuously and in a specific order. Thus, we obtain no extra fidelity by modeling the mining of production blocks as opposed to machine placements; rather, we unnecessarily make the model intractably large. However, using variables that represent machine placements, rather than production blocks, requires nontrivial accounting to consider the time required to mine each production block and, given the number of production blocks in each machine placement, the time to mine (some portion of) the machine placement. We capture this

extra bookkeeping by constructing sets that contain the indices of summation and the indices over which we qualify each constraint.

Several additional constraints, while redundant with the original constraints, restrict the search space for the optimal solution, which reduces solution time. Specifically, we add constraints that (1) require block b to start being mined at some point during the time horizon if its late start date falls within the time horizon (constraints (6)), and (2) allow block b to start being mined at some point during the time horizon if its late start date occurs beyond the time horizon (constraints (7)). These constraints appear as follows:

$$\sum_t y_{bt} = 1 \quad \forall b \mid \bar{t}_b \leq T, \quad (6)$$

$$\sum_t y_{bt} \leq 1 \quad \forall b \mid \bar{t}_b > T. \quad (7)$$

Finally, we can add those production blocks to the schedule that are currently active, i.e., already being mined. In this case, the early and late start dates are equal. The constraint is a special case of constraints (6) and is as follows:

$$y_{bt} = 1 \quad \forall b \mid \underline{t}_b = \bar{t}_b. \quad (8)$$

Naturally, we could rewrite constraints (6)–(8) in terms of machine placements, rather than production blocks, using the appropriate indices. We have, however, kept the notation consistent with that in the formulation above, i.e., the objective function and constraints (1)–(5).

To set the earliest possible start date for a given machine placement, we use an exact algorithm to account for sequencing constraints, i.e., constraints (2)–(4), and for shaft group constraints, i.e., constraints (5), successively updating the start date for each machine placement as necessary, based on the early start times of machine placements whose mining must precede that of the given machine placement. We use a heuristic algorithm to determine a tolerance for the amount of deviation by ore type and month based on the associated demand and availability and then establish a late start date for each machine placement to preclude underlying machine placements from being locked in, thus preventing our meeting demand within the prespecified tolerance.

We implement our mixed-integer program using the AMPL programming language (Fourer et al. 2003) and the CPLEX solver, Version 7.0 (ILOG 2001). The number of time periods and production blocks for our scenario would have required about 66,000 binary variables with the old database in which production blocks were not aggregated into machine placements. With the new database, about 3,600 binary variables are required. By placing the active machine placements in the schedule, i.e., fixing variable values, we

can reduce the number of binary variables by 900. Using early and late start dates to restrict the eligible time periods in which a machine placement can be mined, we reduce the number of integer variables to about 700.

With the original model formulation, planners could not obtain a schedule guaranteed to be within 15 percent of optimality in three days on a Sunblade 1000 with 1024 MB RAM. By contrast, with the new formulation, we obtain an optimal schedule in about five minutes on a Sun Ultra 10 machine with 256 MB RAM.

Acknowledgments

We thank LKAB for the opportunity to work on this challenging project and for permission to publish these results and LKAB employees Hans Engberg, Anders Lindholm, and Jan-Olov Nilsson for providing information, assistance, and support in this endeavor. We also thank several anonymous referees for their comments on a prior version of this paper.

References

Almgren, T. 1994. An approach to long range production and development planning with application to the Kiruna Mine, Sweden. Doctoral thesis 1994:143D, Luleå University of Technology, Luleå, Sweden.

AMPL Optimization LLC. 2001. Version 10.6.16, Bell Laboratories, Murray Hill, New Jersey.

Atlas Copco. 2001. Atlas Copco Rock Drills AB. www.atlascopco.com. Retrieved April 25, 2002 <http://sg01.atlascopco.com/SGSite/Default.asp?cookie%5Ftest=1>.

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

Dagdelen, K., M. Kuchta, E. Topal. 2002. Linear programming model applied to scheduling of iron ore production at the Kiruna Mine, Kiruna, Sweden. *Trans. Soc. Mining, Metallurgy, Exploration, Inc.* 312 194–198.

Fourer, Robert, David M. Gay, Brian W. Kernighan. 2003. *AMPL: A Modeling Language for Mathematical Programming*. Thompson Learning, Pacific Grove, CA.

ILOG Inc. 2001. ILOG CPLEX 7.0. Reference manual and software. Incline Village, NV.

Jawed, M. 1993. Optimal production planning in underground coal mines through goal programming—A case study from an Indian mine. *24th Internat. Appl. Comput. Oper. Res. in Mineral Indust. (APCOM) Sympos. Proc.*, Montreal, Quebec, Canada, 43–50.

Kuchta, M. 1999. Resource modeling for sublevel caving at LKAB's Kiruna Mine. *28th Internat. Appl. Comput. Oper. Res. in Mineral Indust. (APCOM) Sympos. Proc.*, Colorado School of Mines, Golden, CO, 519–526.

Kuchta, M. 2002. A database application for long term production scheduling at LKAB's Kiruna Mine. *30th Internat. Appl. Comput. Oper. Res. in Mineral Indust. (APCOM) Sympos. Proc.*, Phoenix, AZ, 797–804.

Kuchta, M., H. Engberg. 2002. A modular based integrated design approach to computerized mine planning at LKAB. *30th Internat. Appl. Comput. Oper. Res. in Mineral Indust. (APCOM) Sympos. Proc.*, Phoenix, AZ, 341–352.

Lerchs, H., I. F. Grossmann. 1965. Optimum design of open pit mines. *Trans. Canadian Inst. Mining* 68 17–24.

LKAB. 2001. Internal promotional material. Luossavaara-Kiirunavaara Aktiebolag, Kiruna, Sweden.

Newman, A., M. Kuchta. 2003. Eliminating variables and using aggregation to improve the performance of an integer programming production scheduling model for an underground mine. Working paper, Colorado School of Mines, Golden, CO.

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

Tang, X., G. Xiong, X. Li. 1993. An integrated approach to underground gold mine planning and scheduling optimization. *24th Internat. Appl. Comput. Oper. Res. in Mineral Indust. (APCOM) Sympos. Proc.*, Montreal, Quebec, Canada, 148–154.

Tolwinski, B., R. Underwood. 1996. A scheduling algorithm for open pit mines. *IMA J. Math. Appl. in Bus. and Indust.* 7(3) 247–270.

Topal, E. 1998. Long and short term production scheduling of the Kiruna Iron Ore Mine, Kiruna, Sweden. MS thesis, Colorado School of Mines, Golden, CO.

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

Underwood, R., B. Tolwinski. 1998. A mathematical programming viewpoint for solving the ultimate pit problem. *Eur. J. Oper. Res.* 107(1) 96–107.

Wilke, F. L., T. H. Reimer. 1977. Optimizing the short term production schedule for an open pit iron ore mining operation. *15th Internat. Appl. Comput. Oper. Res. in Mineral Indust. (APCOM) Sympos. Proc.*, Brisbane, Australia, 425–433.

Williams, J. K., L. Smith, P. M. Wells. 1973. Planning of underground copper mining. *10th Internat. Appl. Comput. Oper. Res. in Mineral Indust. (APCOM) Sympos. Proc.*, Johannesburg, South Africa, 251–254.

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

Winkler, B. M. 1998. Mine production scheduling using linear programming and virtual reality. *27th Internat. Appl. Comput. Oper. Res. in Mineral Indust. (APCOM) Sympos. Proc.*, Royal School of Mines, London, U.K., 663–673.

Winkler, B. M., P. Griffin. 1998. Mine production scheduling with linear programming—Development of a practical tool. *27th Internat. Appl. Comput. Oper. Res. in Mineral Indust. (APCOM) Sympos. Proc.*, Royal School of Mines, London, U.K., 673–681.

Hans Engberg, Luossavaara-Kiirunavaara AB, 981 86 Kiruna, Sweden, writes: "The long-term mine scheduling of LKAB's Kiirunavaara iron ore deposit has always been a time consuming part of the mine planning process. The major part of this job has previously been done manually with some help from different spreadsheet applications. In order to meet different and changing customer requirements and market situations these schedules in the past had to be updated or redone up to four-five times a year.

"The initial work to update the database and prepare other important parameters still has to be done manually with the new system. The real achievement of this new application is that the actual planning

phase of the five-year plan has shortened from several days to only a few minutes! This fact alone probably gives us the greatest advantage in means of encouraging the fantasy and creativity of the mine planner.

“In the old way of doing things, the planner would have to work very intensely for five to 10 days with the same schedule. To immediately afterwards start to do an alternative plan may have been psychologically

tough, especially when knowing that at least the last years of a five-year schedule are very approximate. With this new program, we feel that we have taken several steps closer to reach optimal solutions in our long-term schedules. These schedules should allow for more efficient use of resources that ultimately will result in reduced mining and overall costs for LKAB in the future.”