

Chapter 1

LONG- AND SHORT-TERM PRODUCTION SCHEDULING AT LKAB'S KIRUNA MINE

Alexandra M. Newman*

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

Mark Kuchta

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

Major Michael Martinez

*Department of Mathematical Sciences
United States Air Force Academy
Colorado Springs, Colorado, 80840 USA
micmarti@mines.edu*

Abstract LKAB's Kiruna mine is an underground sublevel caving mine located above the Arctic circle in northern Sweden. The iron ore mine currently uses a long-term production scheduling model to strategically plan its ore extraction sequence. In this chapter, we describe how we modify this model to consider several different levels of time resolution in the short- versus long-term, and provide guidance for increasing model tractability. We demonstrate numerically the increase in schedule quality and model tractability as a result of these modifications.

*Corresponding author

Keywords: integer programming, production scheduling, underground mining, applications

Introduction and Background

LKAB's Kiruna mine, located above the Arctic circle in northern Sweden, has been operational for more than a century. The mine currently employs about 600 workers and produces approximately 24 million tons of iron ore per year in the form of three raw iron ore products: *B1*, *B2*, and *D3*. These raw products are used to supply planned production quantities at four ore post-processing plants, or mills. The ore products are classified according to their phosphorus content, and are processed into fines and pellets, both of which are used as raw materials in the manufacture of steel end-products.

For about half a century, iron ore at Kiruna was extracted exclusively via surface methods, but about 1960, the pit deepened to such an extent that it became more cost-effective to mine underground. The underground mining method Kiruna currently employs is known as large-scale sublevel caving. This method is used for extracting ore from vertically positioned, fairly pure, large, vein-like deposits. The mine is divided into ten main production areas, about 400 meters (m) to 500 m in length. Each production area has its own group of ore passes; such a group is also known as a shaft group. A shaft group is located at the center of each production area, and extends down to the main level. Each production area consists of about 10 sublevels, and entry to these sublevels is gained via access ramps. One or two 25-ton-capacity electric Load Haul Dump Units (LHDs) operating on a sublevel within each production area transport the ore from the crosscuts (from which the ore is extracted) to the ore passes, where loaded trains haul the ore to a crusher. At the crusher, the ore is broken into pieces small enough to be hoisted to the surface via vertical shafts. Up to 18 LHDs can operate daily throughout the mine; however, the allowable number of LHDs within each shaft group is restricted to about two or three to prevent LHD drivers from driving over and damaging LHD cables.

The site on which each LHD operates is also referred to as a machine placement. The number of machine placements that can be started in a given time period is restricted due to the availability of the crew that prepares the machine placement to start to be mined. The number of active machine placements, i.e., machine placements currently being mined, is also restricted due to LHD availability. Each machine placement belongs to a unique shaft group. A machine placement averages 200 to 500 m in length and contains from one to three million tons of ore and waste rock. A machine placement possesses the same height as the mining sublevel and extends from the hangingwall to the footwall. Between one and five smaller (100 m) entities known as production blocks constitute a machine placement. About one month is required to mine each production block. If a machine placement is left partially mined,

old explosives (which only have a life of about 30 days) must be replaced to reblast the solidified cave rock. This requirement, coupled with the aggravation of tracking partially-mined machine placements, results in operational restrictions that require continuous production within a machine placement until all available ore has been removed. Whether a machine placement can (or must) be mined depends on the relative position of machine placements where mining has already begun. Specifically, certain machine placements beneath a given machine placement cannot start to be mined until some portion of the given machine placement has been mined, and machine placements to the right and left of a given machine placement must start to be mined after a specified portion of the given machine placement has been mined (to prevent blast damage on adjacent machine placements). These operational constraints are referred to as vertical and horizontal sequencing constraints, respectively. Each machine placement possesses a series of notional drawdown lines, consisting of several production blocks each. Within a machine placement, the order in which production blocks must be mined is regulated by this series of drawdown lines, which also helps to enforce continuous mining of a machine placement. These drawdown lines cut horizontally or at a 45 degree angle through several blocks within the machine placement and preclude production blocks in a drawdown line underneath a given drawdown line from being extracted until all ore in a given drawdown line is extracted. This mining pattern is necessary to correctly execute the sublevel caving method so that the mined out areas do not collapse on top of ore that is yet to be retrieved. Minimum and maximum production levels per month govern the rate at which the blocks within a machine placement are mined. These rates ensure continuous mining of machine placements, as discussed above, as well as adherence to production capacity restrictions. Because of vertical and horizontal sequencing constraints and the relative positions of machine placements and the production blocks within them, there are only certain time periods in which these can be mined.

Figure 1.1 illustrates the relationship between the ore body, machine placements, production areas, ore passes, levels and sublevels, vertical shafts, shaft groups, the crusher, crosscuts (i.e., production drifts), the hangingwall (H.W.) and footwall (F.W.). Figure 1.2 shows the relationship between production blocks, machine placements, and drawdown lines. See Topal (2003) for a more detailed description of the Kiruna mine and its characteristics.

1. Literature Review

Winkler (1996), among others, identifies the importance of using integer programming models to determine discrete mine production schedules. However, many researchers are unable to solve realistic mining scenarios, e.g., Trout (1995), Smith (1998), Smith *et al.* (2003). Earlier attempts at the Kiruna




Figure 1.1. The main body in the figure depicts the relationship between the production areas, ore passes, shaft groups, sublevels, and vertical shafts to the surface. The enlargement in the upper left hand corner shows a plan view of a machine placement, consisting of various production drifts, within a production area. The inset in the upper right hand corner shows a load haul dump unit, while the inset in the lower right hand corner depicts a train hauling ore to the crusher.



Figure 1.2. Machine placement MP849.29-30 consists of eight columns of production blocks. Notional drawdown lines pass through a “row” of blocks. The ore body is mapped onto an underlying network of production drifts, which connects to the ore pass.

mine failed to yield a production schedule of requisite length in a reasonable amount of time, e.g., Almgren (1994), Topal (1998), Dagdelen *et al.* (2002). Rather than determining an optimal schedule, these authors resort to shortening the schedule time horizon and/or to sacrificing schedule quality. Carlyle and Eaves (2001) present a tractable model that maximizes revenue from Stillwater’s platinum and palladium mine, as do Sarin and West-Hansen (2005) for an underground coal mine; however, both models differ from ours in their objective and, because of the mining method (sublevel stoping, and longwall, room-and-pillar and retreat mining, respectively), in the constraints. Kuchta *et al.* (2003) provide a more comprehensive list of references.

2. Scheduling Model

Kiruna currently uses a long-term production scheduling model (Kuchta *et al.* 2004) to strategically plan its ore extraction sequence. Prior to the existence of the model, a mine planner would have to develop by hand over the course of a week or more a schedule that would determine when to start mining each machine placement. Because of the vast number of choices and all the scheduling rules, a planner could easily narrow his options unacceptably by choosing to mine certain machine placements that would later preclude sequencing constraints from being satisfied and/or that would result in unacceptably high deviations from demands. As a result, LKAB opted to employ optimization techniques to more quickly generate better-quality production plans. The model is a mixed-integer program that contains thousands of binary variables representing whether or not to mine a specific machine placement in each month of the planning horizon. The model considers the physical limitations of the mine, while meeting as closely as possible the planned production quantities of each raw ore product.

We have modified the original model to comprise several levels of detail. At the coarser (original) level of detail, decision variables consist of which machine placements to start mining each month. For machine placements that are already being mined, we model decisions at a finer level of detail with variables that represent the amount to mine from each production block in each month. Correspondingly, the model also tracks which drawdown lines have been completely mined. This extra level of detail allows the mine planner to more closely control the amount and types of ore that are extracted from each machine placement in each time period. We show in the numerical results section how this extra flexibility helps us to achieve better-quality schedules than those using the long-term model.

A formulation of our combined long- and short-term model follows:

PARAMETERS:

- LHD_t = number of machine placements that can start in time period t

- LHD_v = maximum number of active machine placements in shaft group v
- d_{kt} = demand for ore type k in time period t (ktons)
- $r_{at'tk}$ = reserves of ore type k available at time t in machine placement a given that the machine placement started to be mined at time t' (ktons)
- R_{bk} = reserves of ore type k contained in block b (ktons)
- \bar{C}_{at} = maximum production rate of machine placement a in time period t (ktons per time period)
- \underline{C}_{at} = minimum production rate of machine placement a in time period t (ktons per time period)
- $\rho_{at't} = \begin{cases} 1 & \text{if machine placement } a \text{ is being mined at time } t \text{ given that} \\ & \text{it started to be mined at time } t' \\ 0 & \text{otherwise} \end{cases}$

VARIABLES:

- \bar{z}_{kt} = amount mined above the demand for ore type k in time period t (ktons)
- \underline{z}_{kt} = amount below the demand for ore type k in time period t (ktons)
- x_{bt} = amount mined from production block b in time period t (ktons)
- $w_{lt} = \begin{cases} 1 & \text{if we finish mining all blocks contained in drawdown line } l \\ & \text{by time period } t \\ 0 & \text{otherwise} \end{cases}$
- $y_{at} = \begin{cases} 1 & \text{if we start mining machine placement } a \text{ at time period } t \\ 0 & \text{otherwise} \end{cases}$

FORMULATION:

(P):

$$\min \sum_{k,t} \underline{z}_{kt} + \sum_{k,t} \bar{z}_{kt}$$

subject to

$$\sum_a \sum_{t' \leq t} r_{at'tk} y_{at'} + \sum_b \frac{R_{bk}}{\sum_{\hat{k}} R_{b\hat{k}}} x_{bt} + \underline{z}_{kt} - \bar{z}_{kt} = d_{kt} \quad \forall k, t \quad (1.1)$$

$$\sum_a \sum_k \sum_{t' \leq t} r_{at'tk} y_{at'} + \sum_b x_{bt} = \sum_k d_{kt} \quad \forall t \quad (1.2)$$

$$\sum_a \sum_{t' \leq t} \rho_{at't} y_{at'} + \sum_l (1 - w_{l\hat{t}}) \leq LHD_v \quad \forall v, t, \hat{t} \quad (1.3)$$

$$\sum_a y_{at} \leq LHD_t \quad \forall t \quad (1.4)$$

$$\sum_t x_{bt} \leq \sum_k R_{bk} \quad \forall b \quad (1.5)$$

$$w_{lt} \leq w_{l,t+1} \quad \forall l, t \quad (1.6)$$

$$\sum_b \sum_{u \leq t} x_{bu} \geq \sum_b \sum_k R_{bk} w_{lt} \quad \forall l, t \quad (1.7)$$

$$\sum_{u \leq t} x_{bu} \leq \sum_k R_{bk} w_{lt} \quad \forall b, l, t \quad (1.8)$$

$$\sum_b x_{bt} \leq \bar{C}_{at} \quad \forall a, t \quad (1.9)$$

$$\sum_b x_{bt} \geq \underline{C}_{at} (1 - w_{lt}) \quad \forall a, l, t \quad (1.10)$$

$$w_{lt} \geq y_{at} \quad \forall a, l, t \quad (1.11)$$

$$\sum_{t \leq \hat{t}} y_{at} \geq w_{l\hat{t}} \quad \forall a, l, \hat{t} \quad (1.12)$$

$$\sum_{t \leq t'} y_{at} \geq y_{a't'} \quad \forall a, a', t', a' \neq a \quad (1.13)$$

$$\sum_{t' \leq t} y_{a't'} \geq y_{at} \quad \forall a, a', t, a' \neq a \quad (1.14)$$

$$\sum_t y_{at} \leq 1 \quad \forall a \quad (1.15)$$

$$\bar{z}_{kt}, \underline{z}_{kt}, x_{bt} \geq 0 \quad \forall b, k, t; \quad w_{lt}, y_{at} \text{ binary} \quad \forall l, t, a \quad (1.16)$$

The objective function measures the tons of deviation. Note that we could weight this either by ore type, or by time period, or both. Constraints (1.1) record for each ore type and time period the amount in excess or deficiency of the required amount of ore. Constraints (1.2) require that for each time period in the short term, the total amount of ore required, regardless of ore type, is mined. This prevents the postprocessing mills from sitting idle. Constraints (1.3) limit the maximum number of active machine placements in each shaft group and time period. The index t belongs to the set of time periods in which draw-down line l can finish being mined, and \hat{t} is the time period by which all blocks in drawdown line l must finish being mined. Constraints (1.4) constrain the number of long-term machine placements that can be started in a time period. Constraints (1.5) preclude mining more than the available reserves. Constraints (1.6) indicate that once a drawdown line has finished being mined, it has finished for the horizon. Constraints (1.7) relate the finish of mining a drawdown

line to mining the blocks within that drawdown line. Constraints (1.8) preclude a block b in a drawdown line from starting to be mined unless all blocks in constraining drawdown lines (l) have been mined, and holds for all time periods in which drawdown line l can finish being mined. Constraints (1.9) and (1.10) enforce maximum and minimum production rates, respectively, for each machine placement and time period. Constraints (1.11) and (1.12) enforce vertical and horizontal sequencing, respectively, between machine placements in the short term. Note that in (1.11), the drawdown line, l , in a constraining machine placement controls access to a constrained machine placement, a . (The relationship is reversed in (1.12).) Constraints (1.13) and (1.14) enforce vertical and horizontal sequencing, respectively, between machine placements in the long term. In these two constraints, a' belongs to the set of machine placements whose access is restricted vertically, or forced by adjacency, respectively, to machine placement a . Constraints (1.15) allow a machine placement to start to be mined at most once during the time horizon. Finally, nonnegativity and integrality are enforced, as appropriate. In the interest of brevity, we have omitted the use of a large number of sets, specifying, for example, the machine placements in a shaft group or the eligible time periods in which a machine placement can start to be mined. Martinez *et al.* (2005) give a detailed formulation.

The model is unique in several respects: (i) it does not account for the difference in costs from mining various machine placements due, for example, to their location in the mine, (ii) the objective does not consider the net present value of ore, (iii) model instances are not necessarily solved for the life of the mine, which may result in undesirable end effects, and (iv) there is no allowance for holding inventory, or stockpiling. With respect to the first issue, we assume that all ore will be mined eventually, and hence, total mining costs are sunk. Therefore, we need not consider discrepancies in costs between mining various machine placements. The second aspect is explained by the difference between the markets for iron ore and precious metals. Precious metals such as gold and silver are traded on, for example, the Commodity Exchange of New York. These metals are bought and sold worldwide, and the strategy of mines extracting these metals is to maximize profits by producing as much as is economically viable given current market prices. By contrast, markets associated with base metals such as iron ore are regionalized, as transportation costs are high relative to the value of the commodity. Within these markets, steel companies enter into a contract with an iron ore producer, settling on a price commensurate with the chemical and physical characteristics of the iron ore. Large buyers tend to influence prices in contracts between other buyers and iron ore producers. The negotiated prices generally hold for about a year, and iron ore producers are obligated to supply a certain amount of iron ore to each buyer with whom they hold a contract. Therefore, iron ore mines like Kiruna are concerned with meeting contractual demands as closely as possible. With respect to the third

point, given the future uncertainty in the iron ore composition of each block and the computational time currently required to solve smaller models, attempting to produce life-of-mine schedules is impractical. To mitigate end effects, we solve the model on a rolling horizon basis, updating the ore type composition of each machine placement as the information becomes available.

Finally, company policy does not allow LKAB to stockpile iron ore. A traditional inventory constraint would not apply in this setting at any rate. Specifically, there is physically no space in which to store more than about 50 ktons of extracted iron ore. Furthermore, because LKAB's goal is to meet demand as closely as possible in each time period so as to regulate the amount of ore processed at the mills, a shortage in one time period cannot be compensated by a surplus in, say, the following time period. We could instead recommend using the results from our original model as follows. For each ore type and time period in which there exists production excess (positive deviation), we add to a stockpile until the limit of 50 ktons has been reached. Correspondingly, for each ore type and time period in which there exists a shortage (negative deviation), we draw up to 50 ktons from this stockpile, decrementing the total amount stored, as appropriate. However, too many successive periods of overproduction prevent a significant amount of the overproduced ore from being stored in the stockpile, and too many successive time periods of underproduction deplete the 50 kton buffer without overproducing to replenish it. We find that, for Kiruna's current scenarios, there are few instances of alternating excess and under production between time periods.

3. Solution Techniques

Because the model is large, we use techniques to eliminate all variables that would necessarily assume a value of zero in the optimal solution, and, in fact, in any feasible solution. We capitalize on a modified version of a resource-constrained critical path model to determine earliest and latest possible start dates for each machine placement in the scheduling horizon, allowing us to eliminate a portion of the y_{at} binary variables. The modification manifests itself in that not only are there vertical sequencing constraints that dictate the duration of an "activity," i.e., mining a machine placement, but there exist also horizontal sequencing constraints that require, rather than allow, a subsequent activity to be started only after a given activity has started (or has been completed). LHD availability is the resource. We can also assign a late start date to each machine placement. This late start date would eliminate decisions to start to mine a machine placement so late that adjacent and underlying machine placements eventually become "locked in," thereby increasing the amount of deviation between the actual and the planned production quantities beyond values otherwise obtained in an optimal solution. We use information about

active machine placements together with horizontal sequencing constraints to determine the latest time period in which mining a subset of machine placements can start. We can then assign late start dates to machine placements whose start dates are affected by machine placements within this subset.

We can also eliminate variables associated with mining a drawdown line before an earliest finish date or after a late finish date. We can determine an early finish date for a drawdown line based on the principle of a critical path model by comparing the tonnage available in each drawdown line with the tonnage that can be mined in each time period. Similarly, the latest time at which the mining of a drawdown line could finish is the time at which the first drawdown line in the machine placement finishes being mined added to the longest amount of time it would take all drawdown lines overlying the given drawdown line to be mined. Note that we can use similar principles to establish early start and late finish dates for production blocks to eliminate x_{bt} variables corresponding to mining a production block before its earliest start date or after its latest finish date. However, because the variables associated with mining a production block are continuous, the direct benefit of eliminating such variables is small. However, an indirect benefit of an early start date for each production block is its use in establishing an early start date for a drawdown line, which is simply the earliest early start date among all blocks in a drawdown line. Early start dates for a drawdown line help to eliminate irrelevant terms in constraint (1.3). Martinez *et al.* (2005) provide details regarding the early and late start, and the early and late finish algorithms, as well as the early start algorithm itself.

In previous research addressing only the long-term model, we have used not only variable elimination based on early and late start and finish dates, but also an optimization-based heuristic, which we term the *aggregation procedure*, to eliminate all but a reasonably good set of starting times for each machine placement. This allows us to restrict the model to a subset of start date choices beyond the restrictions we determine with the early and late start algorithms. To date, we have found that this procedure is useful only for eliminating the y_{at} variables because the loss of fidelity inherent in the procedure would be unacceptable for short-term decisions. We refer the interested reader to Newman and Kuchta (2004) for more details.

4. Numerical Results

We demonstrate the benefits of our solution procedures, as well as the improvements we gain in the superiority of the solution by using the combined long- and short-term model over simply using the long-term model for production planning. We conduct our numerical experiments using the AMPL programming language (Fourer, *et al.*, 2003; and AMPL Optimization LLC, 2001) and the CPLEX solver, Version 9.0 (ILOG Corporation, 2004), and use

the CPLEX parameter setting which applies its relaxation induced neighborhood search heuristic every 40 nodes. We run all model instances on a Sunblade 1000 computer with 1 GB RAM. The scenario we use possesses current data from LKAB's Kiruna mine. The data set contains three ore types and spans 24 months.

We summarize the reduction in model size as a function of the number of continuous and binary variables, and the number of constraints when we employ each of the early start, late start, early finish and late finish algorithms *independently of each other*. We apply each algorithm to all relevant model entities: that is, the early start algorithm applies to machine placements, draw-down lines, and production blocks. The late start algorithm applies only to machine placements. The early finish algorithm applies to drawdown lines, while the late finish algorithm applies to both production blocks and drawdown lines. The monolithic model contains over 4000 binary variables and over 8000 constraints. Applying the early start algorithm gives an approximately 25% reduction in both the number of binary variables and in the number of constraints. The number of continuous variables decreases by about 15%. Applying the late start algorithm gives about a 3% reduction in the number of binary variables, and about a 1% reduction in the number of constraints. Applying the early finish algorithm reduces the number of binary variables and the number of constraints by more than 10%. Applying the late finish algorithm reduces the number of binary variables and the number of constraints by about 40%, and the number of continuous variables by 15%. Using all four algorithms in conjunction with each other yields a model with about 20% of the original number of binary variables, continuous variables, and constraints as found in the monolith.

We also make comparisons regarding the quality of the solutions from the long-term and combined models. We apply the variable reduction techniques, mentioned above, to both models, as applicable. Additionally, we weight the objective functions both in the long-term and combined models so as to penalize deviations in earlier time periods more heavily. In contrast to the reduced combined model (with presolve), which has 672 binary variables, 579 continuous variables, and 1702 constraints, the reduced version of the long-term model (including presolve) possesses 416 binary variables, 144 continuous variables, and 734 constraints. The long-term model solves to within 5% of optimality in 4 seconds, whereas the combined model requires 4350 seconds of solution time to reach the same gap.

The extra fidelity in the combined model increases its size and decreases its tractability. However, the combined model yields a 69.6% reduction in deviation compared with that corresponding to the schedule generated by the long-term model. The long-term model only makes decisions at the machine placement level. Once a machine placement starts to be mined, the monthly production quantities within that machine placement are fixed, and mining

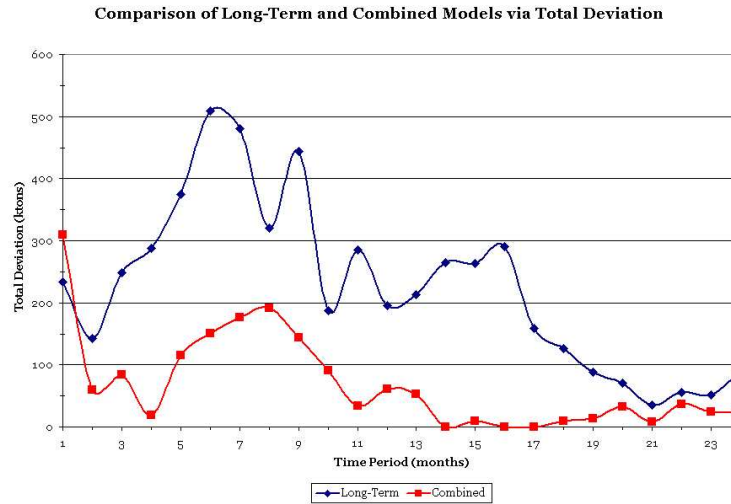


Figure 1.3. Depiction of total deviation (ktons of ore) as a function of the monthly time periods in the planning horizon for both the long-term and combined models. Although the graphs appear qualitatively similar, the amount of deviation in the former graph far exceeds that in the latter graph for all but one time period.

must occur according to that fixed sequence. With the combined long- and short-term model, production rates are allowed to vary between set minimum and maximum values for each production block within a machine placement, thereby allowing partial mining of a monthly production block in order to more closely meet demands for the three ore types.

Figure 1.3 shows a comparison of total deviation (both under- and over-production for all three ore types) for each month in the horizon. Currently active machine placements, i.e., initial conditions, cause high deviations in the early time periods in the solutions of both models relative to the deviations in the later months of the horizon.

5. Conclusions

We present a model that considers both short- and long-term production scheduling for LKAB's Kiruna mine. The benefits of this combined model are its short-term fidelity in directing miners at an operational level which ore to extract, and its long-term clairvoyance showing mine planners at a strategic level the mining areas to develop. Adding short-term fidelity to the original long-term model improves the objective of meeting demands for each ore type and time period. Because the model contains thousands of binary variables and

constraints over just a two-year horizon, we present methods for reducing the size of the model, hence increasing its tractability. Specifically, we develop several algorithms to determine eligible time periods in which a machine placement, a production block, and a drawdown line can be mined. This, in turn, allows us to eliminate variables whose values would equal zero in the optimal solution. Future research entails developing additional methods to enhance model tractability, to enable the generation of production schedules over, say, a four- or five- year horizon. Ultimately, life-of-mine schedules, though the varying availability of data means they can only serve as estimates, would be attractive to mine developers.

Acknowledgments

The authors would like to thank LKAB for the opportunity to work on this challenging project and for permission to publish these results. The authors would specifically like to thank LKAB employees Hans Engberg, Anders Lindholm, and Jan-Olov Nilsson for providing information, assistance, and support. We also thank Professor Kevin Wood of the Naval Postgraduate School, Professor Candi Yano of University of California, Berkeley, and Professors F.S. and J.K. Newman of the University of Illinois at Urbana-Champaign for helpful comments on earlier drafts.

References

- Almgren, T. (1994). *An approach to long range production and development planning with application to the Kiruna Mine, Sweden*, Luleå University of Technology, Doctoral Thesis number 1994:143D.
- AMPL Optimization LLC. (2001). AMPL, Version 10.6.16. www.ampl.com.
- Carlyle, M. and B.C. Eaves. (2001). "Underground planning at Stillwater Mining Company," *Interfaces* **31**(4): 50-60.
- Dagdelen, K., M. Kuchta, and E. Topal. (2002). "Linear programming model applied to scheduling of iron ore production at the Kiruna Mine, Kiruna, Sweden," *Transactions of the Society for Mining, Metallurgy, and Exploration*, **312**: 194-198.
- Fourer, R., D. Gay, and B. W. Kernighan. (2003). *AMPL: A Modeling Language for Mathematical Programming*, Thompson Learning, Pacific Grove, CA.
- ILOG Corporation. (2003). CPLEX, Version 9.0. www.ilog.com.
- Kuchta, M., A. Newman, and E. Topal. (2003). "Production scheduling at LKAB's Kiruna Mine using mixed integer programming," *Mining Engineering*, April, 35-40.
- Kuchta, M., A. Newman, and E. Topal. (2004). "Implementing a production schedule at LKAB's Kiruna Mine," *Interfaces* **34**(2): 124-134.
- Martinez, M., A. Newman, and M. Kuchta. (2005) "Using decomposition to optimize long- and short-term production planning at an underground mine," Working Paper, Division of Economics and Business and Mining Engineering Department, Colorado School of Mines, Golden, CO, November.
- Newman, A., and M. Kuchta. "Using aggregation to optimize long-term production planning at an underground mine," *European Journal of Operational Research*, to appear.

- Sarin, S. and J. West-Hansen. (2005). "The long-term mine production scheduling problem," *IIE Transactions* **37**(2): 109-121.
- Smith, M.L. (1998). "Optimizing short-term production schedules in surface mining: Integrating mine modeling software with AMPL/CPLEX," *International Journal of Surface Mining* **12**(4): 149-155.
- Smith, M.L., Sheppard, I. and G. Karunatillake. (2003). "Using MIP for strategic life-of-mine planning of the lead/zinc stream at Mount Isa Mines," in *Proceedings, 31st International Symposium on the Application of Computers in the Mineral Industry*, Capetown, South Africa: 465-474.
- Topal, E. (1998). *Long and short term production scheduling of the Kiruna iron ore mine, Kiruna, Sweden*, Master of Science Thesis, Colorado School of Mines, Golden, CO.
- Topal, E. (2003). *Advanced underground mine scheduling using mixed integer programming*, PhD Dissertation, Colorado School of Mines, Golden, CO.
- Trout, L.P. (1995). "Underground mine production scheduling using mixed integer programming," in *Proceedings, 25th International Symposium on the Application of Computers in the Mineral Industry*, Brisbane, Australia: 395-400.
- Winkler, B.M. (1996). "Using MILP to optimize period fix costs in complex mine sequencing and scheduling problems," in *Proceedings, 26th International Symposium on the Application of Computers in the Mineral Industry*, Pennsylvania State University, University Park, PA: 441-446.