Interfaces

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

Medium-Term Rail Scheduling for an Iron Ore Mining Company

Gaurav Singh, Rodolfo García-Flores, Andreas Ernst, Palitha Welgama, Meimei Zhang, Kerry Munday

Interfaces
Vol. 44, No. 2, March–April 2014, pp. 222–240
ISSN 0092-2102 (print) | ISSN 1526-551X (online)

informs

http://dx.doi.org/10.1287/inte.1120.0669

Medium-Term Rail Scheduling for an Iron Ore Mining Company

Gaurav Singh, Rodolfo García-Flores, Andreas Ernst
CSIRO Mathematics, Informatics and Statistics, Clayton, Victoria 3168, Australia
{gaurav.singh@csiro.au, rodolfo.garcia-flores@csiro.au, andreas.ernst@csiro.au}

Palitha Welgama, Meimei Zhang, Kerry Munday
Rio Tinto Iron Ore, Operations Center, Perth Domestic Airport, Western Australia 6105, Australia
{palitha.welgama@riotinto.com, meimei.zhang@riotinto.com, kerry.munday@riotinto.com}

In mineral supply chains, medium-term plans are made for scheduling crews, production, and maintenance. These plans must respect constraints associated with loading and unloading, stockyard capacities, fleet capacities, and maintenance and production requirements. Additionally, compliance with grade quality depends on blending minerals from different sources. In this paper, we present an optimization tool developed for a major multinational iron ore mining company to manage the operations of its supply network in the Pilbara region of Western Australia. The tool produces plans for time horizons from a few weeks to two years, while addressing the nonlinearities that blending introduces. The plans our tool produces allow the company to ship a higher amount of iron ore than it did when it followed the plans obtained by its former manual approach. The company's planners now rely solely on our tool because it has enabled them to schedule up to one million additional tonnes of material per annum and has reduced the planning time from five hours to less than one hour.

Keywords: scheduling; rail planning; medium-term planning; blending; mixed-integer nonlinear programming.
History: This paper was refereed. Published online in *Articles in Advance* February 21, 2013.

Australia is the largest iron ore exporter in the world. In 2009, it produced 394 million tonnes of iron ore and exported 362 million tonnes, generating 30 billion Australian dollars (AUD). Of all the iron ore in Australia's mines, 97 percent comes from the Pilbara region of Western Australia. Rio Tinto Iron Ore (RTIO), with an operating capacity of approximately 240 million tonnes per annum, is one of the major producers of iron ore in this region. Its freight rail network, the largest privately owned network in Australia, currently connects 14 mines and three shipping terminals. RTIO has plans to further expand its capacity to 330 million tonnes by 2015. Figure 1, a map of the Pilbara region, shows the region's mines and ports and the relative size of its iron ore reserves.

Figure 2 shows a simplified schematic of the RTIO operations in the Pilbara region. After the iron ore is removed from the ground, it is dumped onto the mine's stockpiles for loading onto trains. This extracted mineral is classified by its size as lump or fines. Both have a number of subproducts, each with its own specific composition or grade. When the trains arrive at the ports, car dumpers transfer the iron ore to the port's stockpiles, and shiploaders transfer the ore into ships for export.

As Figure 2 shows, most of the mines and ports within RTIO's operations have two types of stockpiles, live and bulk. The live stockpiles are considered part of the production line, whereas the bulk stockpiles are used mostly for buffering and storage. Both live and bulk stockpiles have their own maximum storage capacities that are determined by the available area, allocated, for the storage. These storage areas are, accordingly, referred to as live yard and bulk yard. The figure also shows outloaders, which are conveyors that move material from the live to the bulk stockpiles, and inloaders, which move material in the opposite direction. RTIO transfers ore between live and bulk stockpiles based on the following rationale. The raw ore is moved from pit to ship via a series of intermediate transportation subsystems, such as rail networks, inloaders, outloaders, car dumpers, and shiploaders.

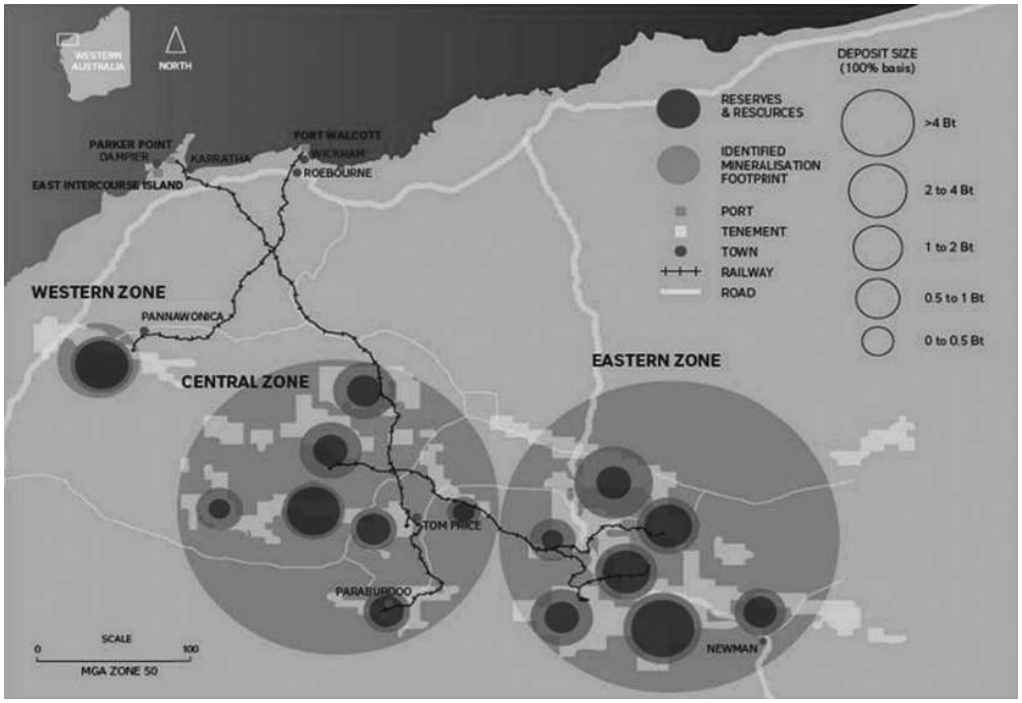

Figure 1: The map shows the operations of Rio Tinto Iron Ore (RTIO) in the Pilbara region. The dark grey bubbles represent the relative size of the reserves and resources of the mines, and the light grey bubbles are proportional to the estimated size of the deposit, according to the scale shown on the top right corner. The light areas represent RTIO's tenement of land.

Bulk stockpiles serve as buffers between these subsystems to mitigate the impact of subsystem variability on the quality of the shipped material. Stockpiles are also used for blending lumps and fines of different compositions into finished products (i.e., blends) that, when shipped, have the composition that RTIO's customers expect (i.e., target composition). The finished product is considered of good quality if it satisfies this target composition. In summary, transfer between stockpiles occurs when the stockpiles' levels exceed

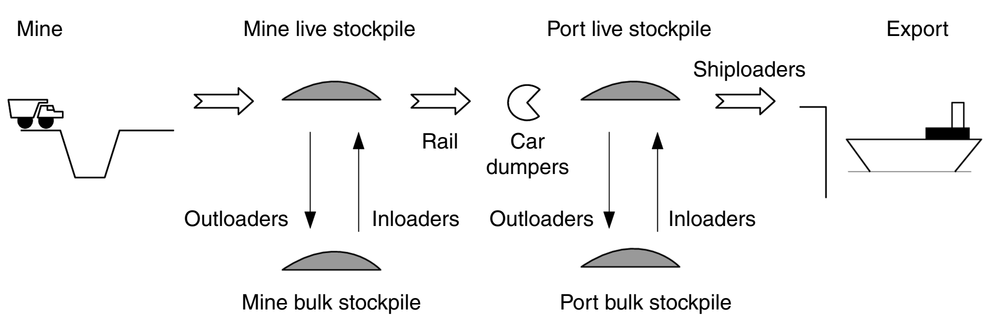

Figure 2: The graphic shows a simplified schematic of RTIO operations in the Pilbara region. Material from the mines is dumped onto the live stockpiles. In both mines and ports, outloaders move iron ore from live to bulk stockpiles, and inloaders move the material in the opposite direction. The ore in the mines' live stockpiles is transported by rail to the ports where it is dumped onto the ports' live stockpiles by car dumpers. Shiploaders put the ore onto ships for export.

the company's preferred margins (i.e., when the stockpiles' tolerance limits have been reached), or when the iron ore grade needs to be modified. However, RTIO prefers not to transfer material to the bulk stockpiles very often, because this increases costs and may also produce delays.

RTIO aims to maximize iron ore throughput while managing grade requirements and system capacities. Fulfilling its contractual obligations with other partner mining companies is an additional concern for the company. These obligations require it to transport minimum amounts of mined products from partner mines to certain ports during the contractual periods. Quality control, capacities, and contractual obligations are nontrivial requirements that, to the best of our knowledge, the literature has not reported as integral parts of medium-term rail planning (i.e., two weeks to two years).

The manual process of determining the required time and the number of trains to be sent from mines to ports across the rail network, while considering all the operational constraints (i.e., determining a schedule), was cumbersome and time consuming. The company's planners made all the required calculations using spreadsheets, which made gaining insight into the problem's complexity difficult. An additional disadvantage of the manual approach was the difficulty of doing what-if analyses. To ameliorate this situation, RTIO sought the assistance of CSIRO's operations research group. We developed an optimization tool that finds the optimal allocation of trains to mines, assists in medium-term operations planning, and enables the assessment of alternative scenarios.

The optimization tool solves a mathematical problem that assumes that management's plan for the next three to five years is the backbone for determining medium-term operations. The model translates resource availabilities, as foreseen by the tactical plan, into operational targets. To facilitate scenario analysis, we incorporated different incentives into the model for planners to use in testing. These incentives act as "knobs" to guide algorithms toward more desirable solutions. They include bonuses for removing material from stockpiles at mines, increasing the stockpile levels at ports, and increasing the number of train trips. In 15 minutes of execution time, our tool produces plans with at least as high a throughput of iron ore as those constructed using five hours (and occasionally more) of manual planning.

The grade requirements introduce an additional complication into the model, that is, nonlinearities. These arise because the target compositions are expressed as fractions whose denominators are the amount of mass in the stockpiles, which is one of the decision variables. See *Nomenclature* in Appendix A (*Formulation*) for the full list of decision variables. The equations are explained in detail in the *Constraints on Grades* subsection and in Appendix B (*Iron Ore Grades*). The grades calculated by our model deviate only marginally from the target composition. Henceforth, we will refer to the differences between calculated and target compositions as grade deviations.

In this paper, we describe a medium-term planning tool for allocating trains to mines and emphasize the challenges posed by target compositions, contractual commitments with partner companies, and product movement between live and bulk stockpiles. We organize the remainder of the paper as follows: following the *Literature Review* section, we give a detailed description of the iron ore transportation system in *Problem Description*. *The Model* discusses the constraints that must be satisfied. In *Implementation*, we discuss issues of practical interest to RTIO, including incentives. *Solution Approaches* introduces the optimal solution method and two additional heuristic methods that we used to provide RTIO with solutions of good quality in a short time. *Results* presents some numerical outcomes of the tool and demonstrates that the proposed model produces solutions with a higher throughput of iron ore than those obtained manually. Finally, we summarize our work and present the company's requests for future research in the *Conclusions* section.

Literature Review

The focus in this paper is on maximizing the iron ore throughput of RTIO's operations by means of a rail freight schedule, while ensuring that the composition of all blends falls within the target ranges, and that all the company's operational and contractual requirements are met. With respect to rail scheduling, several medium-term transportation planning projects have been reported in the literature. An exhaustive review is beyond the scope of this paper; however, the interested reader may refer to the reviews by Caris et al. (2008) and Macharis and Bontekoning (2004), which focus mostly on planning intermodal freight transport, or Cordeau et al. (1998) which focuses only on rail. Newman and Yano (2000) solve a discrete-time scheduling problem to minimize costs for the rail linehaul portion of an intermodal network in the United States and compare it to a variety of decentralized scheduling approaches. They report that decentralized scheduling produces better results. Kuo et al. (2010) present a multiline freight rail scheduling problem that considers demand elasticity. The model minimizes operating costs incurred by carriers and delays incurred by shippers, while ensuring that the schedules and demand levels remain consistent.

With respect to blending and production planning in mining, most research papers address these problems separately. Common examples are quarry production scheduling problems. In this type of problem, the objective is usually to maximize the net present value of a project under precedence, capacity, and blending constraints for individual mines (e.g., Kumral 2011, Rehman and Asad 2010). On a larger scale, previous studies have addressed the challenges of simultaneously optimizing product composition and production schedules in the supply chains of extractive industries. Fröling et al. (2010) propose a planning tool for a company that operates four zinc recycling plants and must allocate zinc-rich steel scrap from different sources to these sites for processing. Their model considers linear input-output functions by multiple linear regression as a way of addressing blending in the planning problem. Sandeman et al. (2010) discuss a mineral export optimization model coupled with a discrete-event simulation model used to fine-tune the grades of mineral and present a gold export operation as a case study. Ulstein et al. (2007) formulate a model for the Norwegian oil industry, where the mole fraction of different hydrocarbons is considered when maximizing the profit obtained from oil and gas distribution plans.

With respect to blends of multiple components, Liu and Sherali (2000) and Shih (1997) calculate optimal plans for the shipping and blending of coal to be used as fuel by electricity companies. These plans must consider electricity demand, coal quality, and price. The authors model composition as tolerance limits of ash, nitrous oxide, and other components, rather than as actual percentages, as we do in our case. Regarding blending and planning for iron ore production, Everett (2001) lists algorithms and simulation methods for estimating iron ore composition at different stages of the production chain to aid in production-scheduling decisions. Everett (2007) gives an excellent description of the role stockpiles play in the iron ore production system and introduces Excel-based software to aid daily ore selection and maintain target composition. García-Flores et al. (2011) outline the iron ore planning and blending problem that we discuss at length in this paper.

With respect to our general modelling approach, Bilgen and Ozkarahan (2007) report a model to calculate shipping schedules for the export of grain blends. These authors also use an integer decision variable to schedule ships in a problem constrained by blending requirements; however, unlike our problem, their two-component blending model is linear. We tackle the nonlinear blending problem using successive linear programming, which consists of using linear approximations and an iterative procedure to calculate product compositions. This method is simple and is commonly used in the petrochemical industry (e.g., Méndez et al. 2006). Audet et al. (2004) review alternative methods to successive linear programming for tackling the blending problem, including bilinear and quadratic programming. We also introduce two sliding time-window heuristics, which are common in manufacturing problems, but have only recently been introduced in open-pit mining problems (Cullenbine et al. 2011).

Some work in the literature has separately addressed facets of the iron ore medium-term rail scheduling problem; however, no previous work has combined blending multiple components and using strategic plans as a guideline for tactical planning. Including these aspects in the formulation has been indispensable for the success of our decision-making tool.

Problem Description

The aim of RTIO planners, and therefore of the tool we present in this paper, is to devise an optimal plan for allocating trains to mines; we wish to maximize revenue, while satisfying all grade requirements and operational capacity constraints. In this section, we describe the main characteristics and constraints of the problem.

RTIO's freight rail network currently consists of approximately 1,400 kilometres of track. The company owns two train fleets, the Robe Valley fleet with five trains of 160 wagons each, and the pooled fleet with 30 trains of 233 wagons each. The former has a capacity of 18 kilotonnes per train and serves the three mines in the Robe Valley region; the latter serves the remaining 11 mines and has a typical capacity of 25 kilotonnes per train. The mines served by the pooled fleet are also divided into regions because there are limits on the number of trains available to serve some of these regions. The cycle time for each trip (i.e., the time it takes for an empty train to go from a port to a mine, load iron ore, and come back to a port) varies between 20 and 40 hours depending on the mine and the product. The average train length is 2.6 kilometres. Figure 3 shows a map of the rail network.

The mines average 350 kilotonnes of live yard capacity; depending on their production and loading capacity, they can send up to 30 trains per day. The material-handling scheme in the mines can be either first-in, first-out (FIFO) or last-in, first-out (LIFO). The scheme determines the order in which the ore enters and leaves the stockpiles; consequently, it affects the composition of the ore conveyed between stockpiles.

The ore extracted in a mine is first placed at the live yard, which has a physical limit on the amount of material it can store. In addition, the live stockpiles have desirable minimum and maximum levels. The stockpile closing level should preferably fall between these tolerance limits at the end of each planning period. Bulk stockpiles have only a maximum capacity limit. The amount of product that can be transported between live and bulk stockpiles in any period is limited by the inloaders' and outloaders'

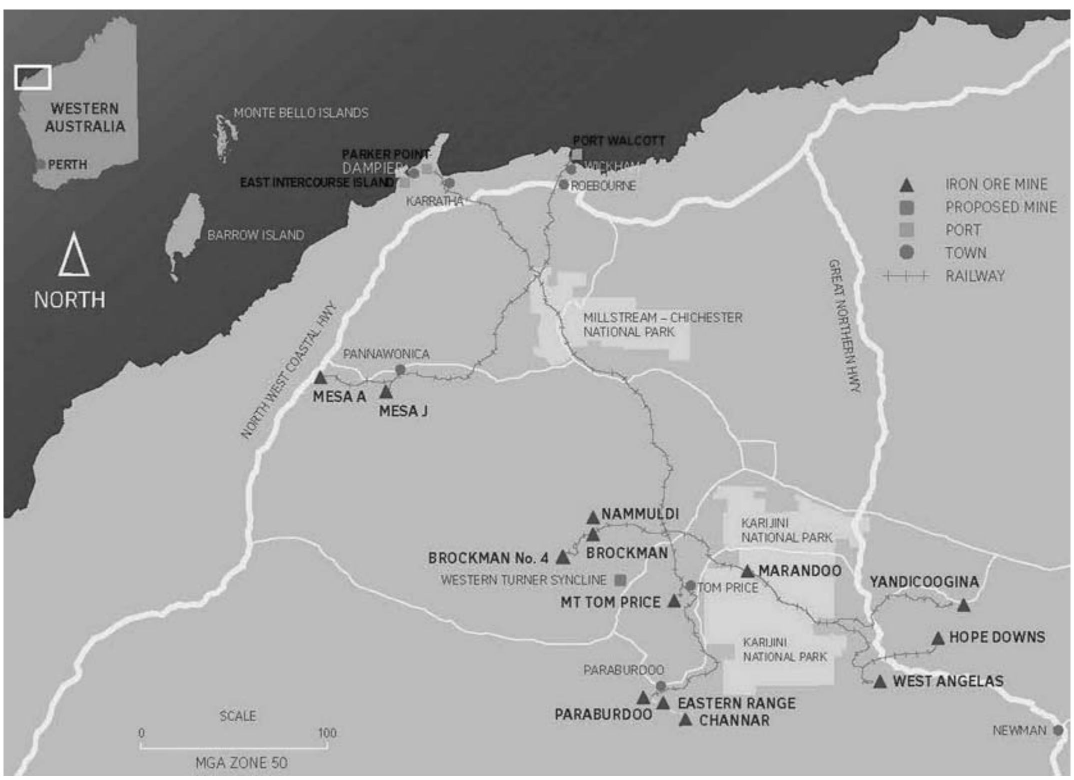

Figure 3: The map shows the Pilbara rail network. It also shows the road network, and the existing and planned mines, ports, and main settlements. The Robe Valley fleet serves the western branch of the rail network (i.e., the Mesa A and Mesa J mines).

capacities at that period. In general, frequently transferring product between live and bulk stockpiles is undesirable because of the costs and delays associated with the transfer operation. Similar stockpile capacity constraints and inloader and outloader capacity constraints apply at the ports.

The shiploaders and car dumpers also have capacity limits expressed as constraints in the model. In any period, the total tonnes of material loaded onto a ship from the port cannot be more than the ship's loading capacity for that period. Similarly, the total number of trains emptied by the car dumper cannot be more than car dumper's capacity for that period.

Some car dumpers may be given preference to unload material originating from certain mines or regions. The dumped iron ore is blended in the ports' live stockpiles to attain shipped product grade. An additional operation, called "lump rescreening" is needed in the ports to comply with final product specifications: when the lump product is loaded onto the ships, a certain proportion of this lump, which is undersized, is returned to a fines stockpile. This returned material is referred to as return fines. Hence, the final composition of the fines product also depends on the composition of the lump product. The contractual obligations specify targets per period on the number of trains delivered from the relevant mines, within a certain tolerance.

Incorporating all the constraints into the manual calculation of a schedule was a tedious process that consumed up to five hours of the RTIO planners' time, without complete certainty that all the constraints were being met. To complicate matters, the demand for iron ore increased steadily from 57 million tonnes in 1999 to 230 million tonnes in 2010, and RTIO is planning an expansion to a total capacity of 330 million tonnes per annum in 2015. As a consequence, the planners expect the number of wagons to increase from the current 8,000 to 12,000 and the number of locomotives to increase from 130 to 180. This translates into up to 20 additional trains to further complicate the problem, making the need for automated planning and scheduling even more acute.

The Model

We consider a finite planning horizon in which each period represents a certain number of days, typically one week. In our model, we do not require the number of days in any two periods to be the same. Mines are divided into five regions, four of which are served by the pooled train fleet and one by the Robe Valley fleet. The main decision variable is the number of trains to be sent from mines to ports during each period. Other important variables are the live and bulk stockpile levels in mines and ports, and the tonnages transferred from (to) live stockpiles to (from) bulk stockpiles.

A summary of the model follows. For each period in a given time horizon, determine the number of trains sent to each mine and the amount shipped from each port to maximize total revenue. Revenue is expressed as the sum of the sales profit minus the total cost of violating the soft constraints, subject to operational constraints, grade constraints, contractual obligation constraints, and car dumper preference constraints. In the next section, we present the soft constraints incorporated in the model. We cannot release the full formulation of the problem for confidentiality reasons, but we present an outline of the model in Appendix A (*Formulation*). García-Flores et al. (2011) provide further details.

Operational Constraints

Capacity constraints form the core of the problem. For the stockpiles, these constraints determine the closing level at the end of each period. The amount of material that a stockpile can store is limited. Both live and bulk stockpiles have a maximum yard limit that cannot be violated; we model it as a hard constraint. Live stockpiles also have minimum and maximum desirable levels, modelled as soft constraints. If the closing level does not fall within these limits, we penalize the difference. The material at mines and ports is stored primarily in the live stockpiles and is transferred to bulk when the maximum desirable live level is reached. Bookkeeping constraints track the inventory levels in the stockpiles. For example, the ending live inventory at a mine in a given period equals its opening inventory, minus the material transferred to the bulk, plus the material from the bulk, plus the production, and minus the material transported from the mines in this period. This is Equation (A18) in *Formulation* in Appendix A. We assume that the amount of material lost because of handling is negligible.

The only stockpile operation that needs further explanation is lump rescreening, introduced in the *Problem Description* section. This operation occurs during the loading of iron ore into the ships at the ports, and consists of removing a fixed proportion of the finer material from the lumps and adding it to the fines stockpile. We also consider lump rescreening in the bookkeeping constraints of the live stockpiles at the ports. Through these constraints, the model considers that compositions of shipped fines also depend on the composition of lump product. Therefore, modelling lump rescreening is important for calculating the composition of shipped product.

We list other operational constraints that we consider in the model.

1. The total amount shipped from the ports in any period cannot exceed the capacity of the shiploaders for that period at that port.

2. The total number of trains sent to a region (or mine) in any period cannot exceed the region's (or mine's) maximum allowed number of trains for that period. Similar constraints exist for car dumpers and inloaders.

3. The total number of fleets and fleet hours used should not exceed the available capacity of a fleet in a given period. We model these as soft constraints. We do not estimate the available fleet hours and cycle times. RTIO uses a detailed simulation model to estimate average cycle times to all regions. We use these averages as an input to the model we present in this paper.

Constraints on Grades

We express the composition of lumps and fines as a mass percentage of iron, silica, alumina, phosphorus, manganese, sulphur, and other components. The tool must consider not only the quantities of iron ore transported by trains and ships and in the live and bulk mine and port stockpiles, but also their composition. To illustrate how this is done, we describe next the expressions for live stockpile grades and grades of material transported by rail.

We calculate the live stockpile grades at the end of a period by finding the total fraction of a component that is left at the live stockpile of a mine after all the transfer operations. The mass balance of the component considering these operations is the sum of the amount of the component that remains in the mine at the end of the previous period, plus the amount in the mined ore as measured by the company, plus the amount transferred from the bulk stockpile, minus the amount transferred to the bulk stockpile, minus the amount transported by rail to the ports. To calculate the fraction of the component, we divide this mass by the total live stockpile level.

To calculate the grades of material transported by rail, we first determine the mass of material loaded onto the trains. This mass is the amount of component in the mined ore as measured by RTIO, plus the amount of the component in the iron ore taken from the live and bulk stockpiles. The material handling scheme affects the order in which the iron ore is loaded onto the trains. If the mine operates according to a LIFO scheme, the current period's mined ore is first loaded onto the trains, and then the ore from the live stockpiles is loaded. If this amount is still not sufficient to fill the train, material is loaded from the bulk stockpile. If the mine operates according to a FIFO scheme, material from the live stockpile is loaded first onto the trains and then the material from the current period's mined ore is loaded; if this is still not sufficient to fill the train, iron ore is loaded from the bulk stockpile. To calculate the mass fraction of each individual component, we divide the component's mass by the train's total storage capacity.

We model the bulk stockpile grades and grades of material to be shipped in a similar fashion. We note that this approach ignores the kind of partial mixing that occurs when materials are placed in a stockpile and replaced again; however, we adopt it as a valid assumption on the advice of RTIO's planners. As the reader can verify in Appendix B (*Iron Ore Grades*), the denominators of the live stockpile grades, Equation (B1), and grades of material transported by rail, Equation (B6), are expressed in terms of decision variables, which introduce nonlinearities in the model. Following Méndez et al. (2006), we use Algorithm 1 in Appendix B to obtain linear approximations to the nonlinear constraints. In the first step, we solve the relaxed problem without grade constraints and without restricting the number of trains to be integer. In the following steps, we use the previous iteration's solution values for stockpile levels and numbers of trains in the grade equations as estimates for the denominators. We repeat this procedure for a fixed number of iterations. As we show in the *Results* section, we need only a few iterations for this process to produce solutions whose penalty because of grade deviations is small.

Constraints on Contractual Obligation

Some of the mines must comply with both annual and per-period delivery quotas. The per-period quota includes a soft target that represents the total number of trains to be sent from each of these mines in the corresponding period. For the annual quota, planners define a cumulative target for each period to ensure that the annual delivery targets are being met, within a certain tolerance, as the end of the year draws closer. The cumulative target is reset at the beginning of the calendar year. We model the per-period constraints as soft constraints and the cumulative targets as hard capacity constraints.

Constraints on Preferences

RTIO prefers that specific car dumpers serve specific mines. We model this as a soft constraint and penalize the amount by which one of these car dumpers cannot comply with unloading the products from the mines in question.

Constraints on the Lump-Fine Ratio

RTIO prefers to maintain a specific ratio of lump to fines, within tolerance, at each port. We also model this as a soft constraint and penalize the amount by which this ratio is violated.

Implementation

We implemented the software application in C++ (C++ Standards Committee 2012) with the capability of interfacing to either Gurobi (Gurobi Optimization 2012) or CPLEX (IBM 2012) solvers. The user interacts with the application through a Visual Basic-supported (Microsoft Corporation 2012) interface. The interface gives structure to the problem by storing the optimization parameters (e.g., incentives and penalties) in separate sheets from resource data. The tool also presents results in a structured manner that allows the user to easily identify bottlenecks by highlighting the most- and least-used resources.

To facilitate scenario analysis, we provide additional control over the model by adding extra knobs in the form of incentives. Incentives are additional terms in the objective function that guide the model toward a solution that recommends the transfer of as much iron ore as possible from the mines to the ports, regardless of whether or not it can be shipped. Without incentives, the tool prescribes the use of just enough train capacity to transfer the material that can be shipped. That is, the purpose of the incentives is to use train capacity to its maximum. We can achieve this using the four incentives we describe here.

1. Encouraging the transport of more material from the mines to the ports. This incentive is for increasing the number of trains; its value is proportional to a fraction of bulk handling costs.
2. Discouraging the accumulation of material at the mine sites. This incentive is for reducing the levels of live and bulk stockpiles at the mines; its value is proportional to a fraction of the bulk handling costs.
3. Favoring the accumulation of finished product at the ports. This incentive is for increasing the levels of live stockpiles in ports; its value is proportional to a fraction of shipped profit.
4. This incentive is similar to incentive 1; however, its value is proportional to a fraction of shipped profit.

Incentives are artificial in the sense that no real cost is incurred for not taking these actions. They are part of the model only if the user activates them; the interface verifies that at most one incentive is active when running the optimizer. We discuss their effect in the *Incentives* subsection of the *Results* section.

Solution Approaches

We implement three methods for solving the medium-term rail-planning problem. The method we refer to as optimal throughout this paper is the linearized mixed-integer nonlinear programming (MINLP) described in the *The Model* section for which the search was terminated with a guarantee that the solution lies within a specified percentage gap of optimality. This method finds the number of trains that maximizes profit over the full-time horizon, and produces the best possible solution, considering all constraints.

In addition to the optimal method and at the request of RTIO's planners, we implemented two sliding-time-window heuristics to provide the user with good quality solutions using less CPU time. We denote these heuristics as H1(I) and H2(I', i'); Appendix C (*Heuristics*) provides details of the heuristics and the meaning of the parameters. These heuristics split the time horizon into subintervals using this rationale: the farther back in time a schedule was realized, the smaller the effect it should have on present plans. For example, optimizing the plan for six months into the future is less important than optimizing it for the next month.

We implement these heuristics because, by dividing the time horizon into intervals, they simplify the problem into smaller subproblems and thereby reduce memory requirements. The heuristics also reduce the computational load on the solver and run at a comparable speed or faster than the optimal method. Finally, limited influence of adjacent time intervals on periods far in the past or in the future is consistent with the planners' intuition about building schedules. We regard all these as valid reasons to use heuristic methods (Zanakis and Evans 1981). We discuss the performance of these methods in the *Comparison of Methods* subsection of the *Results* section.

Results

We obtained the following results using CPLEX 12.1.1 and Gurobi 4.5.1 in a 64-bit Intel Xeon CPU with two processors of eight cores (2.27 GHz) each and 48 GB of RAM. We calculated all the results presented here without using incentives, except for the results presented in the *Incentives* subsection. We justify the use of the solvers in Appendix D (*Optimizer Parameters and Flags*).

Validation

To validate the results, we compare the iron ore throughput prescribed by the optimal and manual schedules. Figure 4 shows the differences between the total amount of material transported as calculated by the optimizer and the supply plan produced manually by the planner. The optimizer used H2(10, 5) and 0.2 percent gap for each time window.

The optimizer scheduled 89 kilotonnes more than the manual plan in the last five months of 2010, and 416 kilotonnes more than the manual plan for all of 2011. Therefore, 505 kilotonnes of additional iron

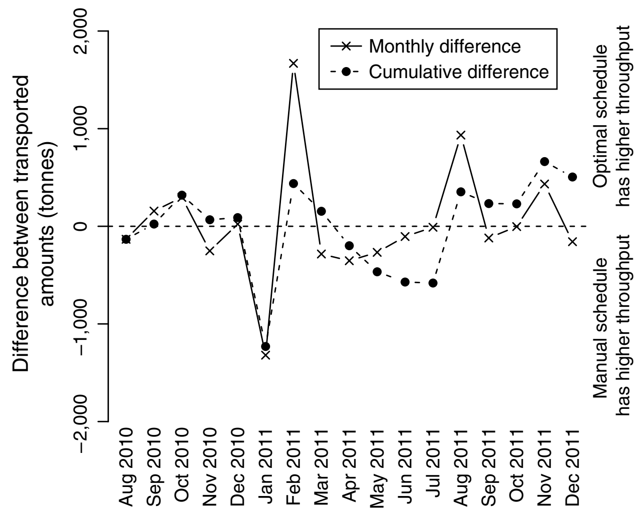

Figure 4: The graph compares tonnage transported using the optimal and manual schedules. The continuous line shows the difference between the monthly transported amounts for these schedules. The dashed line shows the difference in cumulative tonnage. The cumulative difference at the end of the planning period is positive, which means that the optimizer outperforms the manual schedule.

ore were sent to the ports for the entire planning horizon, corresponding to $114 million of additional income. The average amount of iron ore transported per month is 20,196 kilotonnes. The optimizer shows a drop in performance in January 2011 and a recovery in the following month, as a result of confidential information known to the planner about the possibility of compensating with additional train trips in February for a shortage in January; this information was not part of the input to the heuristic. A close inspection of both schedules shows that the manual plan used 58 more trains in January 2011, but the optimizer used 57 more trains in February 2011, which roughly evens out the amount of material transported in these two months. Nevertheless, the cumulative differences clearly show that the annual amount of transported material scheduled by the optimizer is higher.

The tool consistently produces plans with higher iron ore throughput than the manual approach. RTIO has been using our tool for actual rail planning since November 2011. Planners have abandoned the manual approach and now rely solely on our software, which has given them the ability to plan shipping of

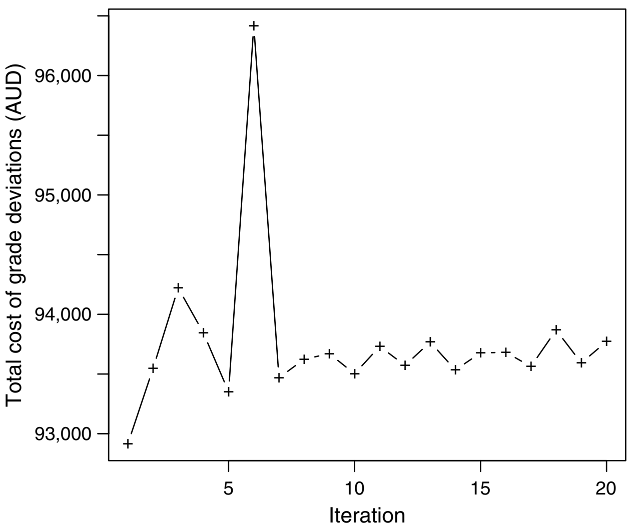

Figure 5: The graph shows the cost of iron ore grade deviations as a function of the number of iterations of Algorithm 1. The results show that after only eight iterations, the algorithm produces decision variable values that are stable enough and close enough to the target compositions to produce only small variations in the total cost of grade deviation from one iteration to the next.

potentially up to one million tonnes per annum more in a much shorter planning time.

Grade Compliance

The results in Figure 5 for 20 iterations of Algorithm 1 (the grade preservation) show that the grades do converge after eight iterations. That is, the algorithm produces decision variable values that are stable enough and close enough to the target compositions to produce only small variations in the cost of grade deviations.

The total grade deviation costs are a very small fraction of the total costs (approximately 10^{-7}), while most of the grade deviation penalties are approximately 12.5×10^6 dollars per kilo tonne, which is a high but arbitrary value chosen in consultation with RTIO planners to make it extremely costly to deviate from the desired compositions (Winston 1994). This combination of high penalties and small contribution of grade deviation costs to the total costs indicates that the deviations from the desired grades are very small.

Incentives

Table 1 shows the differences between the solutions obtained using the four implemented incentives

Incentive	Parameter value	Objective function value minus reference (AUD)	Number of trains minus reference	Transported amount minus reference (kt)
Incentive 1	0.50	1.64e+14	−162	−4.14
Incentive 2	0.50	−8.84e+13	−8	−0.20
Incentive 3	0.01	5.38e+09	5	0.13
Incentive 4	0.10	2.99e+10	113	2.91

Table 1: The table shows results using different incentives (gap = 3%). It shows the objective function values, number of trains, and transported amounts obtained using each incentive, minus the corresponding value obtained when solving the problem without incentives (i.e., the reference solution). Incentives 1 and 2 do not increase the amount of iron ore transported to the ports. Incentives 3 and 4 produce a higher number of train trips than the reference solution without incentives.

discussed at the end of the *Implementation* section and the results of the reference solution, that is, the solution without active incentives. Reiterating, incentives 1 and 4 encourage an increased number of trains, incentive 2 discourages the accumulation of material in the mines' live stockpiles, and incentive 3 encourages the accumulation of material in the ports' live stockpiles. All incentives are artificial and, strictly speaking, do not capture actual operational requirements. However, we consider them useful because they aim to increase train capacity usage; for this reason, we compare them to the reference solution, which captures all operational requirements. Table 1 shows clearly that when using the parameter values shown, incentives 1 and 2 do not produce solutions with higher iron ore throughput than the solution generated without using incentives.

The positive terms in the objective function of the problems that use incentives 1, 3, and 4 boost the value of the decision variables; we can see their contribution in that the objective values of the solutions when using these incentives are greater than the objective value of the reference solution. However, the term in the objective function that corresponds to incentive 2 is negative because this incentive strives to reduce the inventory of iron ore in the mines' stockpiles; this is reflected in a negative difference of objective function values. The results also show that this does not necessarily imply that more trains will be sent to keep the mines' stockpiles low: the solution that uses incentive 2 sends eight trains fewer than the reference solution.

By contrast, incentives 3 and 4 do increase the amount of iron ore transported to the ports. The solution that uses incentive 3 sends five trains more than the reference solution, increasing the amount of material stored in the ports' live stockpiles. Incentive 4 produces a much larger increase in transported ore (113 additional trains) by adding to the objective function a small fraction of the shipped profit times the number of trains. Because the solution that uses incentive 4 produces the largest increase in the amount of iron ore transported by rail, the planner should clearly prefer the solution obtained by using this incentive over the reference solution and all the other incentives. This conclusion is valid when using the parameter values listed in Table 1; however, these parameters can be changed in the Excel interface to explore the solution space.

Comparison of Methods

Figure 6 presents the percentage deviations calculated as the value of the objective function obtained using the optimal method minus the value of the objective function calculated by the heuristics, divided by the value from the optimal method, and multiplied by 100. The grey-shaded bars indicate the value of the percentage deviations after a fixed number of iterations of the grade preservation algorithm. In this figure, the columns on the left of the dotted line correspond to the results of H1, whereas the columns on the right correspond to the results of H2.

We call an algorithm more reliable if its final objective value does not vary with the increase in iterations of the grade preservation algorithm. Using this definition, Figure 6 shows clearly that H2 is more reliable than H1. Observe that the deviations of H1(15) and H1(25) from the objective value, as calculated by the optimal method, can be greater than two percent, whereas all instances of H2 remain less than one percent. As we terminate the execution of the optimal method with a gap of two percent, in some cases, the percentage deviation is a small negative number, which implies that the heuristic can find solutions whose objective value is greater than that of the optimal method. We find it interesting that although the heuristics are myopic and, in the H1 case, do not consider that train numbers are always integral or, in the H2 case, that the full-time horizon parameters are disregarded, they manage to find very good solutions;

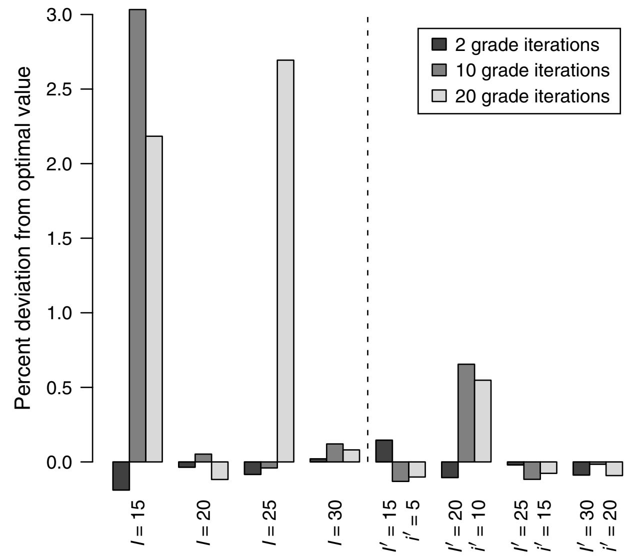

Figure 6: The graph shows the deviation from the optimal value for the heuristics as a function of the number of iterations used in the grade preservation algorithm. The columns on the left of the dotted line correspond to H1; the columns on the right correspond to H2. For details on the heuristics and the meaning of parameters I, I', and i', refer to Appendix C (*Heuristics*). The deviation is calculated as the value of the objective function obtained using the optimal method minus the value of the objective function from the corresponding heuristic, divided by the value from the optimal method, and multiplied by 100. Heuristic H1 is less reliable than H2 with respect to grade preservation because H2's range of percent deviation (0.79%) is lower than H1's (3.22%).

these solutions are occasionally better than those that the optimal method generates.

Figure 7 shows the trade-off between percent gap and CPU time for all the implemented solution methods. The position of a point in this plot shows the CPU time required to achieve a solution using a given method and the solution's gap. We can clearly see the compromise for the optimal method: the solid dots in the figure show that the solution improves as the time limit increases.

Figure 7 shows that H1 generally produces higher revenues in shorter execution times than does H2. We would select H1 as the method of choice if the composition results obtained using it were more reliable. Unfortunately, Figure 6 shows that H1 is less reliable than H2, making it difficult to claim that either heuristic method performs better. Despite this, we recommend the use of either H2(45, 15) or the optimal method with a gap of three percent. The main

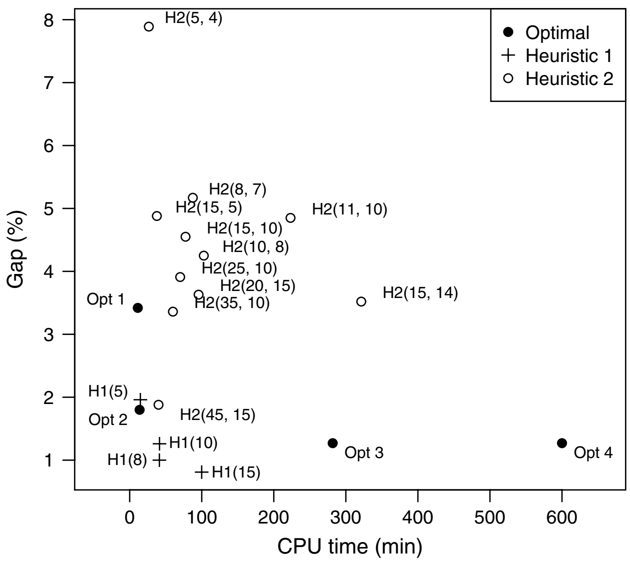

Figure 7: The graphic compares performances in the optimal method and heuristics H1 and H2 in terms of gap percentages and CPU times. Although it may appear from this figure that H1 produces better solutions in shorter times, H2 is preferable because it produces more reliable results, that is, its results are more consistent when combined with Algorithm 1 for preserving grade compositions. H2(45, 15) is an acceptable tradeoff. The points for the optimal method represent runs with different execution time limits; the longer the execution time, the lower the percent gap.

reasons for this choice are that (1) H2 is more reliable in terms of composition accuracy, (2) H2 is more practical because it does not consider the entire planning horizon at each iteration, and (3) by considering subintervals in each calculation step, H2 reflects better the influence of recent actions on decisions that must be taken in the near future. The farther back in time a schedule was realized, the smaller the effect it should have on present plans.

Figure 7 also shows that the speed of a heuristic is generally not affected by its input parameters. This means that heuristic execution times tend to be short, regardless of the gap. We expect this because the heuristics' intervals are smaller; therefore, their optimization subproblems have fewer variables and constraints. By contrast, the accuracy of the optimal method is clearly affected by the execution time limit we impose on it.

Comparison of Solvers

As we noted earlier, we designed the tool to use either the CPLEX or Gurobi solvers. Table 2 shows a comparison between the performances of CPLEX 12.1.1 and Gurobi 4.5.1 using data from one typical production that we set up as an example without incentives, and with execution flags as described in Appendix D (*Optimizer Parameters and Flags*). The table shows that the Gurobi solver returns a smaller gap; thus, it gives a better solution to the maximization problem in a much shorter time. The Gurobi solver created tighter cuts for the problem (i.e., the lower bound after cuts was higher) and found better solutions to the subproblems of the heuristics, which translates into better overall performance and plans with higher throughput. This result was generated although the first feasible solution that the Gurobi solver found was worse than the solution that the CPLEX solver produced.

Conclusions

RTIO's planners successfully implemented our tool to manage the operations of the rail network in the Pilbara region of Western Australia; they have abandoned the manual approach and now rely solely on our software. The tool consistently produces plans with a higher iron ore throughput than the manual approach they used previously. It has reduced the time planners spend determining schedules from five hours to less than one hour, and sometimes to under 15 minutes. The software has increased the amount of material transported to the ports by one million tonnes in a typical planning horizon, which represents an increase in sales of over $100 million.

Solver	Objective value	Gap (%)	Time	LB after cuts	Root solution	First feasible solution gap (%)
Gurobi	-8.01120×10^{10}	2.30	227 min	7.82710×10^{10}	7.81290×10^{10}	54.40
CPLEX	-8.20807×10^{10}	4.71	606 min	7.82143×10^{10}	7.81286×10^{10}	51.01

Table 2: The table compares the Gurobi and CPLEX solver performance. The example uses data from one typical production setup without using incentives. LB denotes lower bound.

Our tool does more than assist in rail network planning. By design, it enables the translation of tactical plans at RTIO into shorter-term operational plans by optimizing the use of resources, including trains, rail tracks, and stockpiles. The user has the option to enable and adjust incentives to transfer material to and from the stockpiles, or to increase the number of train trips. By giving the planners control over the incentives, they can experiment with what-if scenarios, which in some cases produce schedules with higher iron ore throughput than running the optimizer without any incentives.

The company is particularly concerned with (1) preserving the quality of its shipped products, and (2) fulfilling its contractual obligations. To ensure that the resulting schedules fulfill grade requirements, we add equations to the model, which have the unfortunate effect of making it nonlinear. We handle this extra complexity by introducing an iterative algorithm to approximate the actual compositions. Our results show that the deviations from the target composition values are negligible. We also model contractual obligation constraints using annual and per-period targets, which can be met within a user-defined tolerance.

In addition to the optimal method, we implement and test two heuristic methods, labelled H1 and H2. Their purpose is to provide RTIO with usable, good-quality solutions in short execution times. From our results, H1 might initially seem to be the best-performing heuristic because it has a shorter execution time and smaller optimality gap than H2. However, H2 is the preferred heuristic because it has a smaller range of deviation from target grades than H1, as the results of the grade preservation algorithm show. Additionally, the overlapping time horizons between consecutive iterations of H2 better represent the influence of recent actions on near-future decisions: the farther back in time a schedule was realized, the smaller the effect it should have on present plans. This observation is consistent with the planners' intuition about building schedules; thus, RTIO readily accepted H2 as an alternative to the optimal solution method. For these reasons, we recommend the use of either H2(45, 15) or the optimal method with a gap tolerance of three percent.

The tool generates comprehensive output reports that have been useful in operations by helping planners understand uncertainties and identify bottlenecks. It has released valuable staff time for use in other business priorities, and has facilitated the formalization of rules of thumb and operational policies. It has also encouraged staff members to maintain clean and updated resource data.

RTIO started testing the software in late 2010; guided by the planners' feedback, we have since made many improvements. However, some enhancements remain in the pipeline. These include shortening the planning horizon and taking into account ship arrival times to make the tool more relevant to day-to-day operations, incorporating costing and profit information, and providing automatic access to RTIO's databases for increased integration with existing software systems. RTIO and CSIRO will continue to collaborate to address some of the outstanding research challenges.

Appendix A. Formulation

This appendix summarizes all the relevant features of the model (García-Flores et al. 2011). A typical production problem spans 62 periods representing 28 months, and consists of 67,158 variables and 31,024 constraints. We cannot release the full formulation of the problem for confidentiality reasons.

Nomenclature

Sets

- C The set of all components of a product.
- D_{mrp} The set of all car dumpers in port $r \in R$ receiving mined product $p \in P$ from mine $m \in M$.
- F The set of all train fleets.
- G The set of all regions; a region may comprise a single mine.
- M The set of all the mines.
- M_f The set of all mines serviced by fleet $f \in F$, $M_f \subset M$.
- M_g The set of all mines belonging to region $g \in G$, $M_g \subset M$.
- P_m The set of all mined products for mine $m \in M$.
- R The set of all the ports.
- S The set of all unique shipped product, i.e., $S = \bigcup_{m\in M\ p\in P_m} S_{mp}$.
- S_r The set of all shipped products from port $r \in R$, $S_r \subset S$.
- S_{mp} The set of all shipped products for mine $m \in M$ and product $p \in P$, $s_{mp} \subset S$.
- T Total number of planning periods (weeks) in the model.

Model parameters

μ_{ft} Additional cycle time of fleet $f \in F$ at period $t \in T$.
π_4 Fraction of shipped profit for incentive 4.
BPM_{mp} Cost of handling mined product $p \in P_m$ in mine $m \in M$ to/from bulk stockpiles.
BPR_{rs} Cost of handling shipped product $s \in S_r$ in port $r \in R$ to/from bulk stockpiles.
CT_{mpt} Cycle time of a train carrying mined product $p \in P_m$ from mine $m \in M$ during period $t \in T$.
GPI_{rsc} Penalty (per tonne) for deviating from target grade of component $c \in C$ in shipped product $s \in S_r$ at port $r \in R$.
GPO_{rsc} Penalty (per tonne) for violating control limits of component $c \in C$ in shipped product $s \in S_r$ at port $r \in R$.
IOP_grade_{mpt} Actual mined grade (fraction) of product $p \in P$ produced at mine $m \in M$ in period $t \in T$.
IOP_{mpt} Amount of product $p \in P_m$ produced at mine $m \in M$ in period $t \in T$.
MFT_{ft} Available train numbers in fleet $f \in F$ at period $t \in T$.
$MSVL_{mp}$ Penalty for violating limits in live stockpile content of product $p \in P_m$ in mine $m \in M$.
MT_{gt} The maximum number of allowed trains in region $g \in G$ at period $t \in T$.
PFT_{ft} Available pooled hours of fleet $f \in F$ at period $t \in T$.
PP_{md} Penalty for not fulfilling a preference, for example, dumping all product coming from mine $m \in M$ in a specific car dumper $d \in D_m$.
$PRSP_{rt}$ Penalty for violating the specified lump-to-fines ratio at port $r \in R$ in period $t \in T$.
$PSVL_{rs}$ Penalty for violating limits in live stockpile of product $s \in S_r$ in port $r \in R$.
$S^{\max}_{mpt}$ Maximum live stockpile level at mine $m \in M$ for product $p \in P_m$ in period $t \in T$.
$S^{\min}_{mpt}$ Minimum live stockpile level at mine $m \in M$ for product $p \in P_m$ in period $t \in T$.
SP_s Profit per sold tonne of shipped product $s \in S$.
TP_f Penalty for exceeding train trips of fleet $f \in F$.
TS_{mp} Capacity of a train in tonnes transporting mined product $p \in P_m$ from mine $m \in M$.
$U^{\max}_{rs}$ Maximum tonnes of product $s \in S_r$ that can be stockpiled at port $r \in R$.
$W^{\max}_{rst}$ Maximum live stockpile level at port $r \in R$ for product $s \in S_r$ in period $t \in T$.
$W^{\min}_{rst}$ Minimum live stockpile level at port $r \in R$ for product $s \in S_r$ in period $t \in T$.
$Y^{\max}_{mp}$ Maximum tonnes of product $p \in P_m$ that can be stockpiled at mine $m \in M$.
$Z^{\max}_{rt}$ Maximum shipping capacity at port $r \in R$ in period $t \in T$.

Decision variables

$\alpha^{\downarrow}_{mpt}$ Amount by which minimum live stockpile limits are violated at mine $m \in M$ for mined product $p \in P_m$ in period $t \in T$.
$\alpha^{\uparrow}_{mpt}$ Amount by which maximum live stockpile limits are violated at mine $m \in M$ for mined product $p \in P_m$ in period $t \in T$.
$\beta^{\downarrow}_{rst}$ Amount by which minimum live stockpile limits are violated at port $r \in R$ for shipped product $s \in S_r$ in period $t \in T$.
$\beta^{\uparrow}_{rst}$ Amount by which maximum live stockpile limits are violated at port $r \in R$ for shipped product $s \in S_r$ in period $t \in T$.
δ_{rt} Penalty variable for deviations in the preferred ratio of shipped fines to shipped lump at port $r \in R$ in period $t \in T$.
γ_{ft} Amount by which the maximum number of allowed trains of a fleet $f \in F$ was exceeded at period $t \in T$.
κ_{mdt} Amount by which a selected car dumper $d \in D_m$ cannot comply with serving mine $m \in M$ in period $t \in T$.
B_tns_{mpct} Amount of component $c \in C$ transferred from the bulk stockpile of product $p \in P$ at mine $m \in M$ in period $t \in T$.
BM_{mpct} Bulk stockpile grade of component $c \in C$ in product $p \in P_m$ at mine $m \in M$ in period $t \in T$.
ei_{rsct} Excess variable to penalise when grade of component $c \in C$ of shipped product $s \in S_r$ from port $r \in R$ at time $t \in T$ is different from target.
eo_{rsct} Excess variable to penalise when grade of component $c \in C$ of shipped product $s \in S_r$ from port $r \in R$ at time $t \in T$ is outside the control limits.
IOP_tns_{mpct} Amount of component $c \in C$ in mined product $p \in P$ at mine $m \in M$ in period $t \in T$.
L_tns_{mpct} Amount of component $c \in C$ left in the stockpile of product $p \in P$ at mine $m \in M$ at the end of period $t \in T$.
LM_{mpct} Live stockpile grade of component $c \in C$ in product $p \in P_m$ at mine $m \in M$ in period $t \in T$.
R_tns_{mpct} Amount of component $c \in C$ in product $p \in P$ at mine $m \in M$ in period $t \in T$ railed to the ports.
RG_{mpct} Railed grade of component $c \in C$ in product $p \in P_m$ from mine $m \in M$ in period $t \in T$.
s_{mpt} Live stockpile level at mine $m \in M$ for mined product $p \in P_m$ in period $t \in T$.
si_{rsct} Slack variable to penalise when grade of component $c \in C$ of shipped product $s \in S_r$ from port $r \in R$ at time $t \in T$ is different from target.
so_{rsct} Slack variable to penalise when grade of component $c \in C$ of shipped product $s \in S_r$ from port $r \in R$ at time $t \in T$ is outside the control limits.

u^+_{rst} Transfer to bulk stockpiles at port $r \in R$ for shipped product $s \in S_r$ in period $t \in T$.

u^-_{rst} Transfer from bulk stockpiles at port $r \in R$ for shipped product $s \in S_r$ in period $t \in T$.

w_{rst} Live stockpile level at port $r \in R$ for shipped product $s \in S_r$ in period $t \in T$.

x_{mpdst} Number of trains to mine $m \in M$ for mined product $p \in P_m$ at car dumper $d \in D_{mrp}$ for shipped product $s \in S_{mp}$ in period $t \in T$.

y^+_{mpt} Transfer to bulk stockpiles at mine $m \in M$ for mined product $p \in P_m$ in period $t \in T$.

y^-_{mpt} Transfer from bulk stockpiles at mine $m \in M$ for mined product $p \in P_m$ in period $t \in T$.

z_{rst} Amount of product $s \in S$ shipped from port $r \in R$ in period $t \in T$.

Objective Function

The objective is to maximize revenue while minimizing the deviation from the product quality specification.

$$\begin{aligned}\text{Revenue} &= \text{Total profit} - \text{Cost of live stockpile violations}\\ &\quad - \text{Cost of bulk stockpile violations}\\ &\quad - \text{Cost of contractual commitment violations}\\ &\quad - \text{Cost of grade noncompliance}\\ &\quad - \text{Cost of bulk handling}\\ &\quad - \text{Cost of extra trains}\\ &\quad - \text{Cost of violating preference requirements}\\ &\quad - \text{Cost of violating lump-to-fines ratios}\\ &\quad + \text{Incentives.}\end{aligned} \tag{A1}$$

The total profit for shipping product is

$$\sum_{s\in S} SP_s \sum_{t\in T}\sum_{r\in R} z_{rst}, \tag{A2}$$

where SP_s is the profit per shipped tonne of product s and z_{rst} is the amount of product s shipped from port r in period t. The cost of live stockpile limit violations is

$$\sum_{m\in M}\sum_{p\in P_m} MSVL_{mp}\sum_{t\in T}(\alpha^{\uparrow}_{mpt}+\alpha^{\downarrow}_{mpt}) + \sum_{r\in R}\sum_{s\in S_r} PSVL_{rs}\sum_{t\in T}(\beta^{\uparrow}_{rst}+\beta^{\downarrow}_{rst}), \tag{A3}$$

where $MSVL_{mp}$ represents the limit violation penalty for live stockpiles of mined product p in mine m, $PSVL_{rs}$ is the limit violation penalty for live stockpiles in the ports, $\alpha^{\uparrow}_{mpt}$ and $\alpha^{\downarrow}_{mpt}$ are the amounts by which the maximum and minimum live stockpile limits are violated in the mines, respectively, and $\beta^{\uparrow}_{rst}$ and $\beta^{\downarrow}_{rst}$ are the amounts by which maximum and minimum live stockpile limits are violated in the ports, respectively. An analogous term is needed for the cost of bulk stockpile limit violations.

The cost of violating contractual agreements in a given period is

$$\sum_{m\in M^v} JVP_m \sum_{t\in T}(\delta^{\uparrow}_{mt}+\delta^{\downarrow}_{mt}), \tag{A4}$$

where JVP_m is the penalty for violating per-period contractual commitments, and $\delta^{\uparrow}_{mt}$ and $\delta^{\downarrow}_{mt}$ are the excess and slack variables for contractual targets by period, respectively.

The cost of deviating from the target compositions of the shipped products is

$$\sum_{c\in C}\sum_{r\in R}\sum_{s\in S_r}\left[GPI_{rsc}\sum_{t\in T}(si_{rsct}+ei_{rsct}) + GPO_{rsc}\sum_{t\in T}(so_{rsct}+eo_{rsct})\right], \tag{A5}$$

where GPI_{rsc} is the penalty for violating the target grade of component c, GPO_{rsc} is the penalty for violating the composition control limit, si_{rsct} and ei_{rsct} are the slack and excess variables to penalize when c is off target, respectively, and si_{rsct} and ei_{rsct} are the slack and excess variables to penalize when c is outside the control limit, respectively. The cost associated with moving material from (to) bulk stockpiles is

$$\sum_{m\in M}\sum_{p\in P_m} BPM_{mp}\sum_{t\in T}(y^+_{mpt}+y^-_{mpt}) + \sum_{r\in R}\sum_{s\in S_r} BPR_{sr}\sum_{t\in T}(u^+_{rst}+u^-_{rst}), \tag{A6}$$

where BPM_{mp} and BPR_{rs} are the handling costs of products at mines and ports, respectively, y^+_{mpt} and y^-_{mpt} are the transfers to and from bulk stockpiles at mines, and u^+_{rst} and u^-_{rst} are the transfers to and from bulk stockpiles at ports.

There is a limit on the number of trains available per fleet (see constraint (A20) later). The cost of violating this limit is

$$\sum_{f\in F} .TP_f \sum_{t\in T}\gamma_{ft}, \tag{A7}$$

where γ_{ft} is the amount by which the number of allowed trains of fleet f was exceeded, and TP_f is the corresponding penalty.

The cost of not respecting a preference, such as dumping a product using a nonpreferred car dumper, is

$$\sum_{m\in M}\sum_{d\in D_m} PP_{md}\sum_{t\in T}\kappa_{mdt}, \tag{A8}$$

where PP_{md} denotes the penalty for not fulfilling the preference, and κ_{mdt} is the amount by which a selected car dumper d cannot comply with serving mine m in period t.

The penalties for violating the preferred ratio of shipped lump to fines are

$$\sum_{r\in R} PRSP_r \sum_{t\in T}\delta_{rt}, \tag{A9}$$

where $PRSP_{rt}$ is the penalty for violating the specified ratio, and δ_{rt} is the penalty variable for deviations in the ratio.

Finally, we illustrate the use of penalties in the model. The incentive for a higher number of trains as a fraction π_4 of shipped profit is expressed as

$$\pi_4 \sum_{s\in S} \mathrm{SP}_s \sum_{m\in M} \sum_{p\in P_m} \sum_{d\in D} \sum_{t\in T} x_{mpdst}\,, \tag{A10}$$

where the decision variable x_{mpdst} is the number of trains sent from mine m carrying mined product p to car dumper d to produce shipped product s at period t. All other incentives (not shown) are expressed in a similar fashion. Only one of $0 < \pi_i < 1$, $i = 1, \ldots, 4$, is greater than 0 when running the algorithms.

Operational Constraints

The main constraints that must be followed to ensure that the schedule is operationally feasible are:

1. The total amount of live product stacked, both at the mines and the terminal, should not exceed the allocated product capacities and the live stockpile capacity. Live stockpiles are part of the main production line, and bulk stockpiles act as buffers.

$$S^{\min}_{mpt} - \alpha^{\downarrow}_{mpt} \le s_{mpt} \le S^{\max}_{mpt} + \alpha^{\uparrow}_{mpt} \quad \forall\, m \in M,\; p \in P_m,\; t \in T, \tag{A11}$$

$$W^{\min}_{rst} - \beta^{\downarrow}_{rst} \le w_{rst} \le W^{\max}_{rst} + \beta^{\uparrow}_{rst} \quad \forall\, r \in R,\; s \in S_r,\; t \in T, \tag{A12}$$

where s_{mpt} and w_{rst} are the live stockpile levels at mines and ports, respectively, $S^{\min}_{mpt}$ and $S^{\max}_{mpt}$ are the minimum and maximum live stockpile levels at mine m, and $W^{\min}_{rst}$ and $W^{\max}_{rst}$ are the minimum and maximum live stockpile levels of shipped product s at port r in period t. Similar constraints apply to bulk product.

2. Inloaders and outloaders cannot service more than a specified maximum capacity in tonnes per period:

$$0 \le \sum_{s\in S} z_{rst} \le Z^{\max}_{rt} \quad \forall\, r \in R,\; t \in T, \tag{A13}$$

$$0 \le s_{mpt} \le Y^{\max}_{mp} \quad \forall m \in M,\; p \in P_m,\; t \in T, \tag{A14}$$

$$0 \le w_{rst} \le U^{\max}_{rs} \quad \forall r \in R,\; s \in S_r,\; t \in T, \tag{A15}$$

where $Z^{\max}_{rt}$, $U^{\max}_{rs}$, and $Y^{\max}_{mp}$ are the maximum tonnes of product that can be shipped, stockpiled at ports, and stockpiled at mines, respectively.

3. The amount of material in the live stockpiles must not exceed the site's yard limit. For mines and ports that have a bulk stockpile, material accumulates in the live stockpile until this stockpile reaches its maximum control limit. The material is then placed into the bulk stockpile until this stockpile reaches its capacity, and is then put again in the live stockpile until this reaches the yard limit. For mines,

$$S^{\max}_{mpt} + \alpha^{\uparrow}_{mpt} \le YLM_{mpt} \quad \forall\, m \in M,\; p \in P_m,\; t \in T, \tag{A16}$$

$$y^{+}_{mpt} \ge \alpha^{\uparrow}_{mpt} \quad \forall\, m \in M,\; p \in P_m,\; t \in T, \tag{A17}$$

where YLM_{mpt} is the yard limit. Similar constraints apply at ports.

4. At the mines, the amount of product in the live stockpiles at the beginning of a period equals the existing material in the stockpile, plus the produced material, plus the material transferred from the corresponding bulk stockpile, minus the material transferred to the corresponding bulk stockpile, minus the material railed to the ports:

$$s_{mp,\,t+1} = s_{mpt} + IOP_{mpt} + y^{-}_{mpt} - y^{+}_{mpt} - TS_{mp} \sum_{d\in D_{mrp}} \sum_{s\in S_{mp}} x_{mpdst} \quad \forall\, m \in M,\; p \in P_m,\; t \in T, \tag{A18}$$

where TS_{mp} is the capacity of a train in tonnes transporting product p from mine m, and IOP_{mpt} is the amount of product p produced at mine m in period t. Similar bookkeeping constraints apply to bulk at the mines and bulk and live material at the ports.

5. At the ports, as at the mines, the amount of product in the live stockpiles to be shipped at the beginning of a period equals the existing material in the stockpile, minus the material shipped plus the material transferred from the corresponding bulk stockpile, minus the material transferred to the corresponding bulk stockpile, plus the material railed from the mines.

6. The number of trains that service a region cannot exceed the total allowed number of trains in that region:

$$\sum_{m\in M_g} \sum_{p\in P_m} \sum_{d\in D_{mrp}} \sum_{s\in S} x_{mpdst} \le MT_{gt} \quad \forall\, g \in G,\; t \in T, \tag{A19}$$

where MT_{gt} is the maximum number of allowed trains in region g at time t. Similar capacity constraints apply to train fleets, car dumpers, inloaders at the mines and at the ports, and outloaders at the mines and at the ports.

7. The total number of train and train hours for each fleet should not exceed the available pooled fleets and fleet hours in a given period,

$$\sum_{m\in M_f} \sum_{p\in P_m} \sum_{d\in D_{mrp}} \sum_{s\in S} x_{mpdst} \le MFT_{ft} + \gamma_{ft} \quad \forall\, f \in F,\; t \in T, \tag{A20}$$

$$\sum_{m\in M_f} \sum_{p\in P_m} CT_{mpt} \sum_{d\in D_{mrp}} \sum_{s\in S} x_{mpdst} \le PFT_{ft} + \mu_{ft} \quad \forall\, f \in F,\; t \in T, \tag{A21}$$

where CT_{mpt} is the cycle time of a train carrying mined product p from mine m at period t, MFT_{ft} are the available train numbers, and PFT_{ft} is the available pooled hours of fleet f at period t. The γ_{ft} are the additional trains required, and μ_{ft} is the additional cycle time needed for fleet f at period t.

Other constraints (not shown) address special requirements, for example, penalties for not delivering a specific product to a particular car dumper.

Appendix B. Iron Ore Grades

We calculate live stockpile grades at mines by finding the fraction of the total component that is left at the live stockpile of the mine after loading trains and moving material in or out of the bulk stockpiles. More precisely, the live stockpile grade of component c of product p mined at m in period t is

$$LM_{mpct} = \frac{L\_tns_{mpct} + IOP\_tns_{mpct} + B\_tns_{mpct} - R\_tns_{mpct}}{s_{mpt}}, \quad \text{(B1)}$$

where

$$L\_tns_{mpct} = LM_{mpc,\,t-1}(s_{mpc,\,t-1} - y^{+}_{mp,\,t-1}), \quad \text{(B2)}$$

$$IOP\_tns_{mpct} = IOP\_grade_{mpc,\,t-1} IOP_{mp,\,t-1}, \quad \text{(B3)}$$

$$B\_tns_{mpct} = BM_{mpc,\,t-1} y^{-}_{mp,\,t-1}, \quad \text{(B4)}$$

$$R\_tns_{mpct} = RG_{mpc,\,t-1}\left(\sum_{d\in D_{mrp}} \sum_{s\in S_{mp}} TS_{mp} x_{mpds,\,t-1}\right). \quad \text{(B5)}$$

The term L_tns_{mpct} is the fraction of component c left in the stockpile at the end of the previous period, IOP_tns_{mpct} is the amount of c in the mined ore as measured by the company, B_tns_{mpct} is the amount of c transferred from the bulk stockpile, and R_tns_{mpct} is the amount of c transported by rail to the ports. As the equation shows, the estimation of LM_{mpct} requires $LM_{mpc,\,t-1}$, which is also a variable, thus making the equation nonlinear. Bulk stockpile grades at mines are calculated in a similar fashion. For grades of material transported by rail, the model considers material-handling schemes in the mines as one of LIFO or FIFO, as explained in the *Problem Description* section. If the scheme of the mine is LIFO, then

$$RG_{mpct} = \frac{\sum_{d\in D_{mrp}} \sum_{s\in S_{mp}} (IOP\_tns_{mpct} + L\_tns_{mpct} + B\_tns_{mpct})}{\sum_{d\in D_{mrp}} \sum_{s\in S_{mp}} TS_{mp} x_{mpds,\,t-1}}, \quad \text{(B6)}$$

where

$$IOP\_tns_{mpct} = \min\{TS_{mp} x_{mpds,\,t-1}, IOP_{mp,\,t-1}\}\, IOP\_grade_{mpc,\,t-1}, \quad \text{(B7)}$$

$$L\_tns_{mpct} = \min\{s_{mp,\,t-1}, \max\{0, TS_{mp} x_{mpds,\,t-1} - IOP_{mp,\,t-1}\}\}\, LM_{mpc,\,t-1}, \quad \text{(B8)}$$

$$B\_tns_{mpct} = \max\{0, TS_{mp} x_{mpds,\,t-1} - IOP_{mp,\,t-1} - s_{mp,\,t-1}\}\, BM_{mpc,\,t-1}. \quad \text{(B9)}$$

Equations (B7)–(B9) imply that if the mine's scheme is LIFO, the produced quantity will first be loaded onto the trains leaving the mine, and the remaining amount will be loaded from the live stockpiles. If this amount is not sufficient, the remaining amount will come from the bulk stockpile. Accordingly, the respective grades are multiplied to accurately calculate the grade of material transported by rail from the mine. However, if the mine's scheme is FIFO, then

$$IOP\_tns_{mpct} = \min\{IOP_{mp,\,t-1}, \max\{0, TS_{mp} x_{mpds,\,t-1} - s_{mp,\,t-1}\}\}\, IOP\_grade_{mpc,\,t-1},$$

$$L\_tns_{mpct} = \min\{\mathrm{TS}_{mp} x_{mpds,\,t-1}, s_{mp,\,t-1}\}\, \mathrm{LM}_{mpc,\,t-1},$$

$$B\_tns_{mpct} = \max\{0, \mathrm{TS}_{mp} x_{mpds,\,t-1} - IOP_{mp,\,t-1} - s_{mp,\,t-1}\} BM_{mpc,\,t-1}.$$

In this case, trains will be loaded first using live stockpile material, next from the mine production, and finally, if required, from the bulk stockpile.

To obtain linear approximations of the nonlinearities introduced by the above equations into the model, we use Algorithm 1 (Méndez et al. 2006). In the first step, we solve the relaxed problem without grade constraints and without restricting the number of trains to be integer. In the following steps, we use the previous iteration's solution values for stockpile levels and numbers of trains in the grade equations as estimates for the denominators. We repeat this procedure for a fixed number of iterations.

Algorithm 1 (Grade Preservation)
Require: $N =$ Number of iterations for which we should run the algorithm. $\Phi^o =$ Rail planning model to optimize without integer restrictions on numbers of trains.
Ensure: $\Pi^o =$ A rail plan with small deviations from the target compositions.
  $\Pi^T \leftarrow$ Solution of Φ^o without grade constraints.
  for $i = 1$ to N **do**
    $\Phi^T \leftarrow$ A model like Φ^o, but where stockpile levels and numbers of trains from Π^T are used as estimates of the denominators for grade calculation in blending equations.
    $\Pi^T \leftarrow$ Solution of Φ^T.
  end for
  $\Pi^o \leftarrow \Pi^T$.

Appendix C. Heuristics

We implemented two heuristic methods to provide RTIO with good-quality solutions in a short execution time. The heuristics decrease the memory requirements and the load on the solver by dividing the time horizon into subintervals and solving the corresponding subproblems. The underlying rationale is that the farther back in time a schedule was realized, the less effect it should have on present plans. We also implemented these heuristics to assess and compare their effectiveness relative to the optimal method.

Heuristic 1

Heuristic 1, which we denote by H1(I), uses a sequence of iterations over intervals of I periods as follows:

1. In the first iteration, we solve the problem for the full-time horizon by taking the number of trains to be integers for only the first I periods, and as real numbers for the rest of the horizon.
2. In the second iteration, we solve the problem for the complete horizon by giving a tolerance of two trains to the solution from the previous iteration $[0, I]$ to ensure feasibility, and taking the number of trains in periods $[I, 2I]$ to be integers.
3. In general, in the $(j+1)$th iteration, we solve the model over the complete horizon, giving a tolerance of two trains to the solution from the previous iteration $[0, jI]$, and taking the number of trains in periods $[jI, (j+1)I]$ to be integers. That is, we allow the $(i+1)$th iteration to violate the solution of the ith iteration by two trains to ensure feasibility.
4. The process continues until we reach the last period.

Heuristic 2

Heuristic 2, which we denote by H2(I', i'), uses a sequence of iterations over intervals of I' periods with overlap i' as follows:

1. In the first iteration, we solve the model for a limited horizon $[0, I']$.
2. In the second iteration, we solve the model over periods $[0, 2I' - i']$, with the solution from the first iteration used to fix (within a tolerance of two trains to ensure feasibility, as in Heuristic 1) the values of x_{mpdst} for $t \in [0, I' - i']$.
3. In general, in the $(j+1)$th iteration, we solve the model over periods $[0, I' + j(I' - i')]$, with the solution from the jth iteration used to fix (within a tolerance of two trains to ensure feasibility) the values of x_{mpdst} for $t \in [0, j(I' - i')]$.
4. The process continues until we reach the last period.

For the problem presented, a tolerance of two trains is sufficient to ensure feasibility for both heuristics; however, this value is adjustable.

Appendix D. Optimizer Parameters and Flags

For both the CPLEX and Gurobi solvers, we invoked the relaxation-induced neighborhood search (RINS) every 50th node in the tree and set the search parameter to 100 nodes in the corresponding subMIP tree. We set the maximum number of Gomory cut passes to 1,000 and the number of passes of the feasibility pump heuristic (Fischetti et al. 2005) to 100. We also set the termination criterion to a certain percentage gap (different for different scenarios), the maximum CPU time to 600 minutes, and configure the optimizers' flags to aggressively generate mixed-integer rounding cuts and cover cuts. We initially designed the tool to use only the CPLEX solver. However, we realized from our experiments that we found better upper and lower bounds by using the Gurobi solver; therefore, we also added an interface to this solver.

Acknowledgments

We wish to acknowledge the contributions of Meng Zhang and Matthew Baxter.

References

Audet C, Brimberg J, Hansen P, Le Digabel S, Mladenović N (2004) Pooling problem: Alternate formulations and solution methods. *Management Sci.* 50(6):761–776.

Bilgen B, Ozkarahan I (2007) A mixed-integer linear programming model for bulk grain blending and shipping. *Internat. J. Production Econom.* 107(2):555–571.

C++ Standards Committee (2012) JTC1/SC22/WG21—The C++ Standards Committee. Accessed August 28, 2012, http://www.open-std.org/jtc1/sc22/wg21/.

Caris A, Macharis C, Janssens GK (2008) Planning problems in intermodal freight transport: Accomplishments and prospects. *Transportation Planning Tech.* 31(3):277–302.

Cordeau JF, Toth P, Vigo D (1998) A survey of optimization models for train routing and scheduling. *Transportation Sci.* 32(4):380–404.

Cullenbine C, Wood RK, Newman A (2011) A sliding time window heuristic for open pit mine block sequencing. *Optim. Lett.* 5(3):365–377.

Everett JE (2001) Iron ore production scheduling to improve product quality. *Eur. J. Oper. Res.* 129(2):355–361.

Everett JE (2007) Computer aids for production systems management in iron ore mining. *Internat. J. Production Econom.* 110(1–2):213–223.

Fischetti M, Glover F, Lodi A (2005) The feasibility pump. *Math. Programming* 104(1):91–104.

Fröling M, Schwaderer F, Bartusch H, Rentz O (2010) Integrated planning of transportation and recycling for multiple plants based on process simulation. *Eur. J. Oper. Res.* 207(10):958–970.

García-Flores R, Singh G, Ernst A, Welgama P (2011) Medium-term rail planning at Rio Tinto Iron Ore. Chan F, Marinova D, eds. *Proc. 19th Internat. Congress Model. Simulation* (Modelling and Simulation Society of Australia and New Zealand, Perth, Australia), 311–317.

Gurobi Optimization (2012) Gurobi Optimization home page. Accessed January 18, 2012, http://www.gurobi.com/.

IBM (2012) IBM ILOG CPLEX optimizer. Accessed January 18, 2012, http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/.

Kumral M (2011) Incorporating geo-metallurgical information into mine production scheduling. *J. Oper. Res. Soc.* 6(1):60–68.

Kuo A, Miller-Hooks E, Mahmassani HS (2010) Freight train scheduling with elastic demand. *Transportation Res. Part E—Logist. Transportation Rev.* 46(6):1057–1070.

Liu CM, Sherali HD (2000) A coal shipping and blending problem for an electric utility company. *Omega–Internat. J. Management Sci.* 28(4):433–444.

Macharis C, Bontekoning YM (2004) Opportunities for OR in intermodal freight transport research: A review. *Eur. J. Oper. Res.* 153(2):400–416.

Méndez CA, Grossmann IE, Harjunkoski I, Kaboré P (2006) A simultaneous optimization approach for off-line blending and scheduling of oil-refinery operations. *Comput. Chemical Engrg.* 30(4):614–634.

Microsoft Corporation (2012) Visual basic. Accessed August 28, 2012, http://msdn.microsoft.com/en-us/vstudio/hh388568.aspx.

Newman A, Yano CA (2000) Centralized and decentralized train scheduling for intermodal operations. *IIE Trans.* 32(8):743–754.

Rehman SU, Asad MWA (2010) A mixed-integer linear program MILP model for short-range production scheduling of cement quarry operations. *Asia-Pacific J. Oper. Res.* 27(3):315–333.

Sandeman T, Fricke C, Bodon P, Stanford C (2010) Integrating optimization and simulation—A comparison of two case studies in mine planning. Johansson B, Jain S, Montoya-Torres J, Hugan J, Yücesan E, eds. *Proc. 2010 Winter Simulation Conf.*, IEEE, 1898–1910.

Shih LH (1997) Planning of fuel coal imports using a mixed integer programming method. *Internat. J. Production Econom.* 51(3):243–249.

Ulstein NL, Nygreen B, Sagli JR (2007) Tactical planning of offshore petroleum production. *Eur. J. Oper. Res.* 176(1):550–564.

Winston WL (1994) *Operations Research Applications and Algorithms*, 3rd ed. (Duxbury Press, Belmont, CA), 164–169.

Zanakis SH, Evans JR (1981) Heuristic "optimization": Why, when, and how to use it. *Interfaces* 11(5):84–91.

Gaurav Singh is a senior research scientist in the Mathematics, Informatics, and Statistics Division of CSIRO, the Australian federal government's premier research agency. He holds a PhD in scheduling theory and joined CSIRO in 2004 as an OR specialist. He heads the supply chain research team in CSIRO and has successfully led several consulting and optimization systems development projects with major mining supply chains across Australia. He is published widely.

Rodolfo García-Flores has 10 years experience in OR, particularly in mathematical programming, data mining, simulation, and multiagent systems. As a research scientist at CSIRO, he has participated in different commercial projects, for example, in the development of software for minimization of waste in dairy production, the design of beef supply chains, and the analysis of environmental sensor network data. Before joining CSIRO, he was an associate professor in Mexico and a data analyst for a chemical company in Pudsey, England.

Andreas Ernst is the leader of the Operations Research Group at CSIRO's Mathematics, Informatics, and Statistics Division. He has worked for more than a decade on OR problems in bulk material supply chains. His research interests include hub location problems, parallel hybrid metaheuristics, and integer programming.

Palitha Welgama has more than 18 years experience in simulation modeling of logistics systems and in optimization. He has consulted and developed decision support systems for the mining, manufacturing, and service industries. He is manager of systems modeling and optimization with Rio Tinto and is published widely.

Meimei Zhang is senior analyst for optimization and system modeling in integrated planning for Rio Tinto Iron Ore. She earned a PhD in manufacturing systems engineering. She has about 10 years of academic research experience in OR in manufacturing, aerospace, and mining. She was published widely prior to joining Rio Tinto in 2011.

Kerry Munday is a qualified surveyor with more than 25 years experience in the mining industry in planning and scheduling. He has seen his focus shift from a single site to the entire supply chain, with its increases in complexity and volume over the past 10 years and expectations for more complexity in the near future.