Consideration for Multi-objective Metaheuristic Optimisation of Large Iron Ore and Coal Supply Chains, from Resource to Market

J. Balzary and A. Mohais

Abstract Dynamic market and operating conditions coupled with an environment in which multiple objectives and trade-offs are common, pose major challenges for planners and schedulers working in any mining entity. Many mining companies recognise the need to shift from a siloed mining-focused push model to an integrated value chain, demand-driven approach but there are still fundamental barriers in business process and the supporting technology preventing a consideration of end-to-end optimality. This paper presents some elements of experiences working with companies to adopt such advanced approaches. In addition to algorithmic elements, an approach to phased and gradual deployment of progressively more sophisticated optimisation models is described. From a practical software adoption perspective, it is believed that this last concern is also of primary importance. Next generation approaches to the optimisation of complex bulk commodity demand chains; namely iron ore and coal are presented, with case studies in the world's largest integrated operations in Western Australia and Queensland from the raw material mined through to market. Utilising accurate simulation models supported by metaheuristic optimisation techniques, a range of ways to engineer a dynamic decision support framework that can adapt and change with the inevitable changes in commodity markets is explored. Objectives such as total revenue, margin, cost, NPV, throughput, asset utilisation, contractual penalties and bonuses, and energy consumption can be managed simultaneously across the mine, plant, logistics network, port operation, shipping and sales domains. AQ1

J. Balzary (✉) · A. Mohais
Supply and Demand Optimisation, Schneider Electric, Level 1,
99 Frome Street, Adelaide 5000, Australia
e-mail: james.balzary@schneider-electric.com

A. Mohais
e-mail: arvind.mohais@schneider-electric.com

R. Dimitrakopoulos (ed.), *Advances in Applied Strategic Mine Planning*,
https://doi.org/10.1007/978-3-319-69320-0_20

Introduction

The Resource-to-Market mining supply or demand chain can be represented most broadly from the pre-extracted, in situ resource through to the point at which an organisation can invoice upon sale. Much effort is placed on systems facilitating mathematical and computational improvements in decision making at various points in this value chain. Technologies exist to assist in optimal mining operations sequencing, and also in subsequent material handling and logistics processes down the chain. According to current mathematical knowledge, for the class of problem represented by the full Resource-to-Market supply chain and all of its complexities, there is no known method of solution that would give an absolute and irrefutable optimal planning or scheduling outcome. Whilst this is a mathematical reality that businesses must come to terms with, from an opportunistic point-of-view, it presents stakeholders with the ever-present possibility that they can continually improve decision support and modelling technologies and do better.

With this opportunity as a motivator and using of practical and real-world learnings as atomic components, this paper presents a next-generation optimisation framework that would deliver further benefit and profit to mining organisations globally. Included is a brief overview of the nature of bulk mining supply chains, conceptualised from a software point of view—from available raw material, through beneficiation, transport, storage and onto vessels. Within this supply chain, we will identify a number of important component segments that can be treated as silos, or preferentially, should be treated as integrated parts of a larger global operation. Standardised key performance indicators are described, with targets set for each as reward-based fitness measures. Methodologies for the utilisation of advanced science solutions involving modern heuristic optimisers and metaheuristic algorithms are described guiding the search efforts of the lower-level searches.

Supply Chain Objectives

Companies seeking to optimise the planning and scheduling of their Resource-to-Market supply chains express their view of an optimal solution in terms of certain objectives that they would like to achieve. These objectives are often either to maximise or minimise a particular measure of performance in the supply chain, or sometimes to keep a measure confined to a targeted band of values. Although we can identify typical objectives that are common across mining entities, different companies often would attribute different degrees of importance to these objectives, in effect weighting their contribution to the overall evaluation of a plan or schedule. Typical objectives encountered in the mining industry are:

1. Increase margin
2. Maximise cash inflow
3. Minimise cash outflow
4. Maximisation of asset utilisation, fixed and mobile plant
5. Maximisation of sequenced activities—e.g. vessels berthed per tide
6. Maximisation of efficiency—e.g. direct train to vessel loading
7. Minimisation of variability—e.g. quality through processing
8. Minimisation of penalty—e.g. demurrage
9. Achieve target tonnage value e.g. rail and shipped.

Conflicting Objectives

Typically in mining supply chains, and in fact every business, objectives conflict with each other. These conflicts involve the interrelationships of complex business rules, processes, constraints and performance measures. Take for example the desire to maximise fixed asset utilisation in a port operation. Maximising asset utilisation is in direct conflict with a common objective to maximise direct train to vessel loading. A model of these activities seeking to keep car dumpers, stackers and reclaimers in continual use would like come up with a sequence that schedules and dumps trains as soon as there is an available time slot on any of these pieces of equipment. The alternate view is to delay arrival of a train so that it coincides with the berthing of a vessel therefore allowing for direct loading, but possibly at the cost of keeping the aforementioned pieces of equipment idle.

There are many other such examples of conflicting objectives. It is inherent in any company's expression of their optimisation wishes.

Handling Multiple Objectives

In the literature a common theoretical construct that is proffered for managing situations with multiple conflicting objectives is to use the notion of Pareto fronts (citation), but limited application of this exists in decision support in production environments (citation). Under this approach, instead of seeking a single optimised solution, a collection of solutions is retained, forming a so-called Pareto front of non-dominated solutions. Any solution in this set has the characteristic of being better than all the others on at least one of the objectives.

The philosophy behind this approach is that the value of the work done by the optimisation software should be retained in the form of the Pareto front, and then this set of high-quality solutions should then be passed to a human expert for final analysis and evaluation, and the human expert would make the final decision on which should be selected. In some situations, this approach is feasible. However, we believe that in the Resource-to-Market context, the number of possible conflicts is reasonably high, and the presentation of a Pareto front to a human expert by the software would be of limited value because of the number of potential solutions that would be expected in such a set, and the work required to make a final decision would not at all be straightforward.

Another factor that currently weighs against purely multi-objective algorithms in the Resource-to-Market space is that when the magnitude of the supply chain, coupled with the number of data elements and the time horizon are taken into consideration, the potential running time of such an algorithm is infeasible for the decision making timeframe. A more appropriate methodology given these considerations is the relative weighting of each objective, and their combination into a single unified evaluation measure. A downside to this approach is the fact that the scales on which different objectives are measured could vary dramatically Therefore trying to combine them using appropriate weights would often require re-tuning to get the right values. A possible approach to eliminate this variability is to allow an authorised end-user to state their weighting preferences using a unit-free normalised scale, and to let the science and software experts to devise and tune appropriate multiplicative factors to compensate for the inherent scale differences of the different objectives.

Objective Function

Although each mining operation would have idiosyncrasies necessitating modifications to the objective function used by an optimisation algorithm, it is still possible to provide in closed mathematical form, an expression of that function, using typical and common terms. This is illustrated below for the case of a scheduling problem:

Let X be an element of the unconstrained solution space. Then X can be expressed as a collection of scheduled activities, across train loading, railing, car dumping, stacking, reclaiming, conveyance, ship loading and berthing. Hence X can be expressed as the union of disjoint subsets of activities in each of these areas as follows:

$$X = X_{TLO} \cup X_R \cup X_{CD} \cup X_S \cup X_{RE} \cup X_C \cup X_{SL} \cup X_B$$

where X_{TLO} is a set of train loading activities, X_R is a set of railing activities, X_{CD} is a set of car dumping activities, X_S is a set of stacking activities, X_{RE} is a set of reclaiming activities, X_C is a set of conveyance activities, X_{SL} is a set of ship loading activities, and X_B is a set of berthing activities. These discrete activities are the required steps move excavated material from the pit onto a ship at berth at the port.

The elements contributing to the objective function would be:

Revenue	$\$R = \sum_{x \in X_{SL}} saleprice(x)$
Costs	$\$C = \sum_{x \in X} cost(x)$
Resource utilisation (fraction)	$RU = \sum_{Y \in \{X_{TLO}, X_R, X_{CD}, X_S, X_{RE}, X_C, X_{SL}, X_B\}} \frac{constrained\ capacity - \sum_{x \in Y} duration(x)}{constrained\ capacity}$
Demurrage costs	$\$D = \sum_{x \in X_B} demurrage\_penalty(x)$
Silo constraint violations	$CV = \sum_{Y \in \{X_{TLO}, X_R, X_{CD}, X_S, X_{RE}, X_C, X_{SL}, X_B\}} \sum_{x \in Y} constraint\_violation\_severity(x)$
Target shipped tonnes penalty	$TSTP = shipped\ tonnage\ target - \sum_{x \in X_{SL}} tonnage(x)$
Target railed tonnes penalty	$TRTP = railed\ tonnage\ target - \sum_{x \in X_R} tonnage(x)$

Using these contributing elements as representative, the objective function can then be expressed as:

$$f(X) = w_1\$R - (w_2\$C + w_3\$D) - (w_4RU + w_5CV + w_6TSTP + w_7TRTP)$$

The coefficients $w_1, \ldots, w_7$ are weights that are configurable by users with the right access privileges. The first three terms of the function are intuitive dollar values which are readily justified. The last four terms are penalties due to violations of constrains and operating rules, or for un-achieved targets. To justify the a simply numeric difference with the dollar values requires human input into the weightings so that their relative importance is correctly judged in relation to the hard dollar values.

Literature Survey

Although Supply Chain Modelling and Supply Chain Management are heavily researched areas, the published literature addressing resource-to-market optimisation in the mining context is relatively small. Bodon et al. (2017, in this volume), describe the challenges of using a discrete event simulation language to model the complexities of a pit to port coal supply chain, and propose a de-coupling of the simulation aspect of the model from optimisation aspects. They used a general linear program as an optimiser, and couple this with discrete event simulation, and presented their results on scenarios from a real-world coal mining operation in Indonesia. Further related work is available in Bodon et al. (2011). Peng et al. (2009) provide an analysis of an integrated coal supply chain, and apply the model to the Xuzhou coal mine in China. They present results showing that not only optimal profit is obtained, but that a level of customer satisfaction is achieved. Their results enabled recommendations to be made for the mine operations, and assisted in decision making.

Montiel and Dimitrakopoulos (2013) look at a coper mining supply chain, with emphasis on global optimisation from the point of view of taking into account the output of multiples mines and products in a given mining complex. Their work also focuses on the variability of orebody models, and deals with their stochastic nature by producing stochastic mine production schedules. Montiel and Dimitrakopoulos' work was based on using Simulated Annealing, a metaheuristic approach, for producing mine schedules. Their results showed that a stochastic schedule produced expected deviations from mill and waste production targets smaller than 5%, versus that of conventionally generated schedules which was 20%. Although their work did not focus on the Resource-to-Market supply chain as it has been outlined in this paper, their model nevertheless considers a large subset of the mining supply chain, particularly around the details of excavation, waste haulage, milling, and further value-adding preparation and handling of the product.

Singh et al. (2012) provide a detailed elaboration of a mathematical model constructed to represent the operations of the Hunter Valley coal chain in eastern Australia. Their goal was to create a model that could find a supply chain schedule/plan that would meet a given demand profile, whilst concurrently suggesting any capacity increases or new equipment that would be required to support that solution. Singh et al's model was not built into an end-user enterprise application, and their results potentially could take up to several hours to compute, which would make it challenging for the kind of software implementations that Schneider Electric's SDO is interested in. However, their work is remarkable to us because of the level of detail that was built into the model in certain places, and because of the hybrid nature and multi-phase approach to their solution.

Their model was developed around assumptions for a demand-driven, cargo-assembly type operation. Historical demand profiles were used to drive the model and optimisation process. The main goal of the optimisation model was to minimise the cost of running the terminal for that demand profile. The cargo-assembly approach required that all products required for loading a vessel be delivered and already stacked at the port before loading begins. Hence, direct loading was not considered in their model. Rail was modelled around the key factors of a limited number of consists (potential trains) per day, and a limited number of paths through rail junctions at the mine and at the port.

They explored using Genetic Algorithms, and Squeaky Wheel heuristics to generate individuals with representation components involving job sequences, and capacity/equipment increments. The solutions produced by these algorithms are then passed to a CPLEX algorithm to generate a final solution. Singh et al. concluded that it was a challenging problem that could not be easily solved by then-currently available MILP (mixed integer linear programming) commercial software or straight application of general metaheuristics like genetic algorithms. Some of their other approaches (squeaky wheel and another called large neighbourhood search) produced somewhat better results, but they acknowledged room for improvement, possibly by exploring alternative or more closely coupled hybridisations between the MILP approach and heuristic search methods.

Live Software Implementation Experience with Mining Companies

Enterprise level software solutions in the Resource-to-Market domain for iron ore and coal mining have been deployed into live use for major Australian mining companies over the last three years. Some experiences, modelling and algorithmic details from these implementations are provided utilising two scenario sections of the paper. In each case, a future extension of the approach is described, which seeks to apply meta-level optimisation in an effort to further improve on the results that have been previously achieved in practice.

Two scenarios are presented to illustrate different time horizons, one of which necessitates a finer-grained "scheduling" approach, and the other a more coarse-grained "planning" approach. There are substantive differences between the approaches, and the algorithms used must be tailored accordingly. As an added benefit, the cases described were chosen so as to reflect both iron ore mining and coal mining.

Scenario #1—Scheduling System for Iron Ore

Fortescue Metals Group (FMG) is Australia's third largest iron ore producer operating 3 mines, a dedicated rail line and port in Western Australia. In this scenario the model manages several silos from post-beneficiation to vessel. Focus is placed on important elemental aspects of the algorithms used in the software with an outline of a future-state meta-level algorithm proposed. This progression in algorithmic complexity follows a prescribed staged approach where initial deployment of optimisation technology is managed in a step by step fashion, beginning with simplified acceptable techniques and migrating to more advanced, automated decision support paradigms.

The deployed decision support model focuses on the modelling of trains and the rail network between the various train load-out (TLOs) and the port. The system is configured with a fixed number of rakes or consists (a collection of wagons assembled to carry an iron ore product) that need to be scheduled in order to meet demand at the port. Queuing of rakes at the TLOs is an important factor in the local scheduling decisions considered when looking at the rail silo. For the iron ore scheduling problem, two elements of the deployed algorithm that are particularly important include (i) the demand-driven nature of the algorithm, and (ii) the technique of disruption propagation. These are both used as baseline elements of the current and future version of the algorithm.

Layout: T1_Standard Book ID: 457301_3_En Book ISBN: 978-3-319-69319-4
Chapter No.: 20 Date: 14-10-2017 Time: 4:45 pm Page: 8/20

Components of a Scheduling Solution

In the scheduling (versus planning) domain, the emphasis is on very detailed and comprehensively-specified activities scheduled with a start and an end time. Many details of each activity in question, such as the equipment utilised and inventory produced must be modelled and calculated. The computational effort is often prohibitive and the respective granularity and accuracy of data diminishes over a long time horizon, making long term decisions on highly detailed models infeasible. The need to manage the level of detail and the importance of these finite elements in the scheduling horizon naturally focuses attention in the short-term (hours, shift, days).

In a typical mining Resource-to-Market requirement for a scheduling purpose, the following activities are specified (as examples):

Train (rake/consist) service	
• Rake/Consist ID • Train destination (mine) • Port depot departure time • Selected loader at mine (TLO) • Product to be loaded (type, tonnage, quality)	• Queuing time at mine • Loading duration • Journey time • Queuing time at port • Selected unloader at port • Optional periodic maintenance at port
Train loading activity	
• TLO ID • Product type • Product tonnage	• Product quality • Loading start time • Loading end time
Car dumping activity	
• Car Dumper ID • Rake ID • Product type • Product tonnage • Product quality • Conveyor route ID	• Stockpile destination ID (if applicable for stacking) • Shiploader ID (if applicable for direct loading) • Dumping start time • Dumping end time
Stacking activity	
• Car Dumper ID • Stacker ID • Stockpile ID • Conveyor route ID • Product ID	• Product tonnage • Product quality • Stacking start time • Stacking end time
Reclaiming activity	
• Reclaimer ID • Stockpile ID • Shiploader ID • Conveyor route ID • Product ID	• Product tonnage • Product quality • Reclaiming start time • Reclaiming end time

(continued)

(continued)

Ship berthing/de-berthing activity	
• Vessel ID • Berth ID • "Pilot on board" (POB) time • "First line" time	• "All fast" time • "Ready to load" time • Depart berth time
Ship loading activity	
• Shiploader ID • Berth ID • Conveyor route ID • Product ID	• Product tonnage • Product quality • Loading start time • Loading end time

Demand-Driven Solution Generation

This iron ore case study uses a demand-centric perspective to drive the optimised solution generation with primary demand based on vessel nominations and the associated attributes for contractual fulfilment.

The market factor is very important in this model, as it is the primary determinant in the schedule produced. The client organisation provides data on future sales for the time horizon under consideration. This consists of firm orders, which are ones which have already been confirmed by the end buyers, as well as tentative orders, which are indications of intention to buy. This data is provided to the scheduling software by means of direct data integration. The scheduling software has a data exchange interface with other software systems used by the client organisation, and the latest variations are always available for use in generating new and updated schedules. The data is provided in the form of Vessel Nominations, which are contracts for the sale of iron ore commodities to be loaded at a designated port by a particular vessel. The data contains the Estimated Time of Arrival (ETA) of the vessel at the anchor point associated with a port. This date and time is used by the scheduling algorithm to determine possible choices for a time of berthing for that vessel.

During simulation, when a particular instant in time is being considered, a vessel that has been tentatively selected for berthing at that time is examined to see what its nomination is, i.e. which products, their respective volume and quality are required for loading once it is berthed at the port. This demand triggers a backward-looking analysis agent that retraces the steps along the supply chain that are needed for the required amount of the right products to be available at the port at the time of the vessel's arrival. This then creates precursor demands within upstream silos in the supply chain, which must be optimised concurrently with other scheduled activities in those silos.

Disruption Propagation

When a certain magnitude of change occurs to a scheduled activity, for example the berthing of a vessel an hour earlier than planned, it is possible to locally propagate the effect of those changes and quickly get the overall schedule back into a correct feasible state without having to undergo a computationally expensive re-building of the entire schedule. An understanding of the implications of this propagation without optimisation is managed via constraint handling which references the available capacity and buffer between each related activity and determines if violations have occurred that are infeasible.

Assuming stockpiles that were intended to be used to load the vessel were already at their required inventory levels several hours before, the fact that the vessel is early does not cause any problem with loading. Alternatively, it may be that a sequence of trains that were scheduled to arrive throughout the duration of the vessel being at berth are now out of sync with this shift in time, and the problem could be fixed by shifting the train schedules all by one hour earlier. We would have to check how this triggers knock-on effect higher up the supply chain, and also potentially look at effects like congestion or conflict on the rail network, if the supply chain is being modelled to that extent.

The key thing to note is that it is possible for small disruptions in one silo to be relatively easily absorbed by adjacent silos in the supply chain, and it is a wise tactic to attempt to use this propagation opportunity to quickly absorb these changes as opposed to attempting brute force re-optimisation.

Realities of Global Optimisation

It is important to note that within the confines of the decision making timeframe, it would be impractical to create a problem representation that encompasses the entire supply chain, and then use a population-based algorithm that simply treats the individuals as candidate solutions to this massive problem. In practice the computational power required to process that magnitude of scope and complexity would be infeasible, and so would the required computing time under current hardware constraints. Furthermore it would be naïve to expect that simple operators (such as intra-silo mutations, or crossovers across silos, or even across the global representation) working on a massive representation would be able to effectively or efficiently find the truly high-quality solutions that human experts are seeking.

It is important to recognise that although the global context must be considered, and the desired solution would have less-than-optimal sub-solutions within silos, we should nevertheless respect the local logic and intelligence that exists within the silos (human or modelled). It is through judicious use of this intelligence that we can arrive at a solution that can be considered globally optimised.

Hybrid Global Optimisation

We propose that a hybrid approach is needed, which acknowledges the fact that a truly optimised solution must take into account the global, multi-silo nature of the problem, but which also intelligently operates on the representation so that infeasible solutions are avoided, and also that natural heuristic corrections to adjacent silos are carried out in response to an evolutionary disruption in a target silo.

Consider for example a change in the vessel-loading activities at the port, wherein a particular ship loader requires more material than is currently scheduled, and thus draws upon a stockpile to an extent surpassing its current stock. (This is not possible in physical reality, but can certainly be considered as part of an individual representation). This shortfall in inventory at that stockpile is a natural impetus for the adjacent rail module to undergo an amount of re-optimisation, whether it be a small or a large change remains to be determined). This principle gives rise to a multi-silo algorithm which can be called “Disruption Dampening and Transmission”. Changes in one silo may cause nudging on an adjacent silo, which may be accommodated by slight movement, i.e. a dampening of the disruption, or it may be necessary to completely re-adjust the neighbour to try to align its endpoints with the disruption, i.e. a full transmission of the disruption.

Figure 1 illustrates the concept of disruption dampening and transmission, using the analogy of sitting on a bench. To get a better understanding of how this proposed algorithm would be implemented, Fig. 2 provides a pseudo-code outline. The key idea of this algorithm is to choose a most influential silo (or weight them in importance and choose probabilistically), and run a full optimisation routine on it, but after each iteration, as individuals are modified, the effect of their modifications either get dampened by virtue of adjacent silos being able to absorb the impact of the change with small-scale modifications, or get transmitted with a more disruptive effect into the adjacent silo, triggering a full re-optimisation of the current state within that neighbouring silo.

If this approach is contrasted with a more straightforward approach to global optimisation, one could imagine that a change in a silo would be followed immediately by an evaluation of the overall individual. The resulting individual is likely to contain multiple constraint violations and task misalignments. These could be handled by penalty components in the fitness evaluation of the individual, but the likelihood of this being able to successfully guide the algorithm is very low.

The above proposed approach could be likened somewhat to repair algorithms from evolutionary computation. What is substantially different however, is the

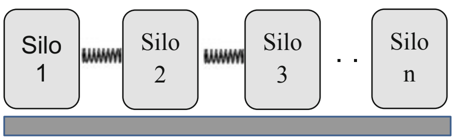

Fig. 1 Disruption dampening and transmission between silos

Layout: T1_Standard | Book ID: 457301_3_En | Book ISBN: 978-3-319-69319-4
Chapter No.: 20 | Date: 14-10-2017 Time: 4:45 pm | Page: 12/20


```
heuristically initialise N candidate solutions

while (termination condition not satisfied)
{
  probabilistically select silo for disruption initiation
  heuristically effect local beneficial change in selected silo
  initialize disruption-queue with adjacent silos
  while (disruption-queue is not empty)
  {
   remove and process silo at front of queue
    {
      run local propagation algorithm to repair silo

      if (repair successful){
        for each adjacent silo no longer aligned, add silo to disruption
        queue
      }
      else
      {
        run heavy-weight full optimisation on silo
      }
    }
  }
}
```

Fig. 2 Disruption dampening and transmission algorithm outline

possibility of complete local re-optimisation of certain silos, and also phased propagation of the disruption of a change throughout the supply chain.

Case Study #2—Planning System for Coal

Glencore (previously Xstrata) Coal is a major global energy materials producer. This example includes a multi-mine operation centred on raw coal management, coal handling and preparation through the plant, rail logistics to vessel loading using two berths at the Abbott Point Coal Terminal in Queensland, Australia.

In this model, attention was paid to the maximisation of the potential total revenue by not only relying on the supplied data on contracts for Month 1–3 years, but also considering the more detailed addition of place-holder vessels in order to make enable recommendations to the Sales department, highlighting where additional product is available to be sold. The importance of shipping data for capacity assessment is elaborated by Boland et al. (2011).

The fact that the Australian coal industry often fails to meet demand due to inadequate planning, infrastructure deficiencies and other reasons is outlined by Bayer et al. (2009) and represents a primary driver for organisations to look at exploiting latent value associated through improved planning and optimisation. As before with the iron ore case study, the coal case study is described from an abstract point of view, all the way to a vision of future algorithmic approach. For the coal planning problem, a few elements of the current production-implemented model and associated optimisation algorithm that are particularly important include (i) a stock accumulation-based representation for vessel loading, (ii) quality up-building and

Layout: T1_Standard Book ID: 457301_3_En Book ISBN: 978-3-319-69319-4
Chapter No.: 20 Date: 14-10-2017 Time: 4:45 pm Page: 13/20

down-building heuristics, and (iii) upstream heuristic plan completion based on the vessel-loading driver. These elements form the baseline optimisation improvements upon which enhanced future state optimisation is considered. The main realisations that are used when proposing the future-state algorithm is that heuristic construction of seed individuals is important for a modern heuristic algorithm, therefore a simplified heuristic approach as a foundation element to optimised plan generation is a feasible. There is a careful balance that is needed between those kinds of individuals and more randomly generated ones in order to find an appropriate balance between biased and free range exploration of the search space. A meta-level algorithm is part of the proposal to find this appropriate balance.

Solution Representation for a Planning Context

In contrast to the level of model complexity in scheduling (minutes, hours, shifts, days), planning systems (days, weeks, months) are orientated towards a higher-level summary view of what can be achieved, and safely planned for, in a long-term time horizon—for example 1 month to several years. Similar to heavily constrained activity scheduling relationships, the planning requirement must take into consideration numerous parameters and hard and soft constraints, in order to ensure that the results are valid. Whereas a schedule would consist of a number of discrete activities assigned to different resources, planning models are generally defined around summarised aggregated activities or capacities within a fairly large time bucket, in this case monthly. Plans are created starting at month 3 (from the current time), out to 3.5 years. The initial period of 3 months is not covered because this is considered to be within the scheduling period, not planning. For each month of the planning horizon, the following information must be generated:

Haulage Plan:

- Aggregated tonne-hours for movement between ROM and stockyards, inter-stockyard, and stockyard to CHPP. Individual journeys are not modelled.

Field Stockyards Plan:

- The total tonnes of each product type on field stockpiles for that month.

ROM Stockyard Plan:

- The total tonnes of each product type on a ROM stockpiles for that month.

CHPP Operation Plan:

- Tonnes of each coal type sent to each CHPP module, and bypass.
- Output tonnes for each coal type for each CHPP module, as well as new ash %, and reject tonnes.

CHPP Clean Coal Plan:

- Tonnes of each clean coal product added to each stockpile.
- Stockpile tonnes and % capacity.
- Quality attributes of each blended stockpile.

Rail Plan:

- Train-hours—Tonnes of each coal type transported by train from each mine

Port Stockyard Plan:

- Tonnes of each Brand
- Quality attributes for coal assigned to a Brand

Shipping Plan:

a. The number of satisfied 'TBC' (to be confirmed) shipments, including tonnage and product quality attributes. TBC shipments are derived using a typical vessel size and accounting for customer contracts (i.e. tonnage and quality requirements)
b. Number of non-contracted proposed shipments of coal brands, including tonnage and product quality attributes. This is coal that is not linked to a contract. This provides a view of how much additional coal product is produced by the mines and needs to be sold by Marketing.

Optimisation Algorithm

Part of the local optimisation heuristic for a vessel stock accumulation plan is the tuning of the selected components for blending to achieve a shippable product type, i.e., one which is within the target quality specification bandwidths. The representation of an individual in this algorithm consists of a set of ordered pairs, where each pair consists of a viable stockpile id, and a desired tonnage to be reclaimed from that stockpile (Fig. 3).

Low-Grade Up-Build Blending Sub-algorithm

The approach of the Low-Grade Up-Build blending algorithm is to initialise an individual in a deficient sub-space of the coal blending selection space. Such individuals would be of a low grade, and a search algorithm would need to be structured so that overarching directional vectors of the search tend towards sub-spaces that are richer in terms of coal quality. It is important to control the velocity of movement so that there is an increased likelihood of discovering suitable blends within the required quality tolerances at an early stage within the search process, without exploring too deeply within the high quality areas of the search space.

High-Grade Down-Build Blending Sub-algorithm

The High-Grade Down-Build blending algorithm uses a converse approach, which is to initialise an individual in an adequate or rich sub-space of the coal blending space. Such individuals would be of a high grade, and a search algorithm would need to be structured so that the direction of the search moves at low velocity towards lower quality regions. The objective being to have a high likelihood of settling in a region that maximises the use of lower grades, whilst still remaining within the required quality tolerances.

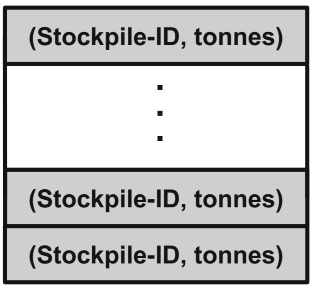

Fig. 3 Individual representation for stock accumulation

For both blending algorithms, evolutionary operators are used to gradually bring an individual to within tolerance, the main operator being a swap of a small tonnage of ore, exchanging low with high grade, or vice versa.

Upstream Scheduling Building Based on Ship Loading Profile

Once a candidate vessel berthing sequence has been determined, a heuristically-built individual representing a schedule for the entire supply chain can be constructed by working backwards, upstream in the supply chain, to create nominal activities to match the requirements of the vessel at berth within a given period of time. This gives rise to a so-called heuristically built individual that contains elements of a good-quality solution, but has not yet been optimised.

The Evolution to Metaheuristic Optimisation

Evolutionary algorithms often can be made to produce excellent results on problems in a particular domain, but one of the issues that arises is that there are often several algorithmic parameters involved, and these parameters need to be correctly tuned in order to achieve positive results. In certain situations, it is the case that a meta-algorithm, or metaheuristic, can be engineered to run at a higher level and perform the tuning of the lower-level evolutionary algorithm. Thus, any human manual intervention in the finding of high-quality solutions is minimised, and the work can be relegated mostly to the computational machinery and software.

Since we are considering primarily Population-Based Modern Heuristic (PBMH) optimisation algorithms as the key tools for optimising the supply chain, we will describe the concept of a metaheuristic optimiser in this context. For the Coal Planning problem under consideration, in order for the PBMH to operate effectively, it is critical that it be seeded with candidate solutions that have already been placed into reasonably feasible sub-spaces of the search space (Fig. 4). This is accomplished by using a percentage of the seeded individuals that are passed through a local optimisation routine to achieve some moderate level of fitness before entering into the PBMH algorithm. Furthermore, there is another percentage (typically quite small) of individuals that are heuristically built, but which have not been subjected to the local optimisation. There are also introduced into the seed population for the purpose of maintaining genetic diversity.

Details of how the Initial Seeding component of the algorithm would work are illustrated in Fig. 5. The V1 and V2 local searches referenced in that diagram refer to the Low-Grade Up-Build and High-Grade Down-Build heuristics defined earlier. The initial population of the main population-based modern heuristic (PBMH)

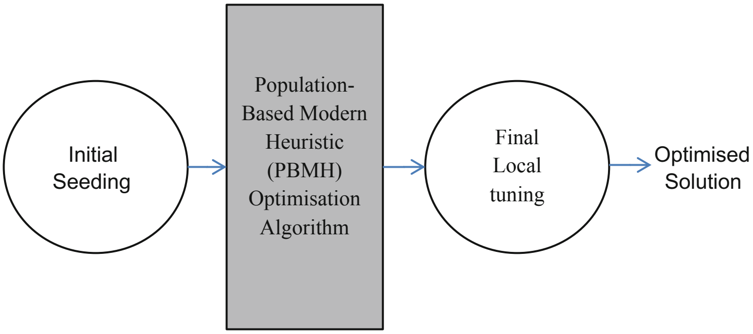

Fig. 4 Pre and post processing needed for a population-based modern heuristic optimisation algorithm

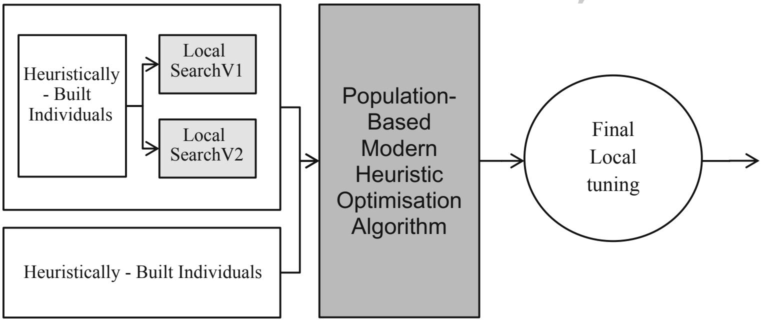

Fig. 5 Use of local search algorithm for coal blending, pre PBMH

optimisation algorithm is divided into 3 parts: (i) individuals that have gone through V1 local optimisation, (ii) individuals that have gone through V2 local optimisation and (iii) individuals that have not gone through any local optimisation. Each part can be thought of as being of a particular percentage.

For the PBMH operating at a meta-level, the goal is to find an optimised combination of these percentages (two would suffice), such that when used to seed the lower-level PBMH, the best possible planning solution results (Fig. 6). To get a better understanding of how this proposed algorithm would be implemented, Fig. 7 provides a pseudo-code outline.

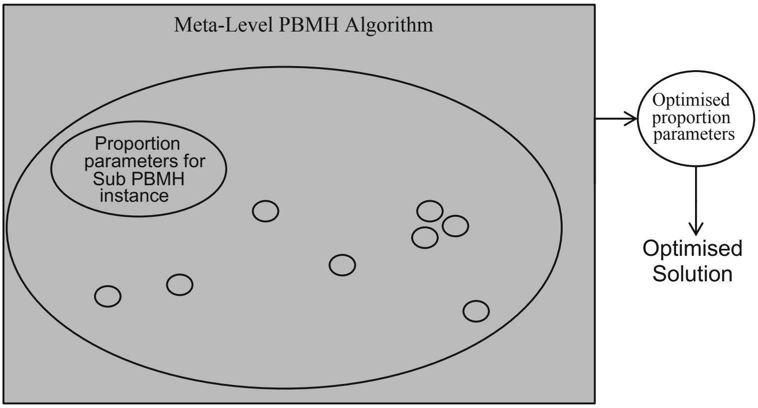

Fig. 6 Each individual of the meta-level algorithm is a set of parameters for a base-level PBMH

```
heuristically initialise  N candidate solutions consisting of
 1.  Percentage of locally optimised seed candidates using V1 algorithm.
 2.  Percentage of locally optimised seed candidates using V2 algorithm.

while (termination condition not satisfied)
{
  evaluate all candidate solutions by running lower-level PBMH
  optimisation

  prune population to standard size N

  probabilistically select individuals for percentage-shift mutation
  probabilistically select individuals for vector midpoint sampling

  run percentage-shift mutation
  run vector midpoint sampling
}
```

Fig. 7 Meta-level PBMH algorithm for parameter tuning

Conclusions

Any mining value chain scheduling problem involves the assignment of a high number of variable activities to a set of resources. From a computational complexity perspective, this problem is known to be NP-Complete, effectively indicating the Resource-to-Market scheduling problem is currently amongst the most challenging problems known. Furthermore, many of the constraints that exist within that domain are non-linear in nature. Due to these complexity characteristics it is expected that population-based modern heuristic methods are highly appropriate for finding high-quality solutions, as opposed to methods premised on linear constraints and linear models as they can inherently manage more complex business rules and non-linear constraints.

The planning problem appears to be less complex in nature than detailed scheduling since jobs are not being assigned to resources, but rather aggregated capacity is being consumed against those resources in less-granular time buckets. Nevertheless, this apparent reduction in complexity is usually offset by the practice of considering much longer time horizons—many months or years into the future.

The Resource-to-Market problem is currently managed in the real world production environment by predominantly talented human experts who, together with various rudimentary tools, for example spreadsheet models, and very limited-scope and narrowly-focused software applications such as discreet event simulators, are coping with the task of keeping businesses running by finding suitable, though arguably sub-optimal solutions to the problem. The mining business community has a strong appetite for advanced software solutions using novel and innovative mathematics, science and technology to improve in this area.

Considerable care must be taken when embarking upon the journey of making major changes to how scheduling and planning tasks are carried out by all mining organisations. The deployment of software that instantaneously and dramatically shifts the scheduling/planning paradigm in place, even if this does hold the potential for much higher-quality results, more often than not is a sure recipe for immediate reticence, incomprehension, doubt, overall inertia, and eventual rejection of the new system. Despite the potential of advanced scientific software solutions, it is important to recognise and respect that the process of adoption of such systems is in no small part a human activity. It is important to carry out such an endeavour as a staged process, using a roadmap of checkpoints that guides the organisation and its experts in an incremental fashion. At each step, clearly-understood solutions must be produced by the software in a manner that the human expert would feel comfortable signing-off on. Especially in the early stages of the roadmap, it is critical that the actions of the software be explainable and comprehendible.

Modern heuristic algorithms have been discussed and applied at length in the research community for more than 30 years. In the last ten years however, there has been a noticeable emergence of commercial-grade enterprise level software that incorporates these kinds of algorithms though arguably their uptake has been limited in production environments.

Currently implemented elements in existing clients for Schneider Electric are presented as components of a framework for meta-level optimisation. These baseline elements are designed to be expanded, scaled and enhanced as the understanding and acceptance of their output is trusted. The benefit that would be achieved from any optimisation technique needs to be carefully weighed against the increase in runtime that would ensue. The continual increase in power and capability of computer hardware, including the ability to leverage parallel computation diminishes the impact of this downside.

It is expected that these approaches will yield higher quality solutions than the perceived state-of-the-art production models in use today, whilst remaining amenable to implementation in enterprise software designed for mining supply chain experts who are not necessarily mathematical modelling and optimisation specialists.

References

Bayer A, Rademacher M, Rutherford A (2009) Development and perspectives of the Australian coal supply chain and implications for the export market. Zeitschrift fur EnergiewirtschaftI 33 (3):255–267 (Springer)

Bodon P, Fricke C, Sandeman T, Stanford C (2011) Modelling the mining supply chain from mine to port: a combined optimization and simulation approach. J Min Sci 47(2):202–211 (SP MAIK Nauka/Interperiodica)

Bodon P, Fricke C, Sandeman T, Stanford C (2017) Combining optimisation and simulation to model a supply chain from pit to port

Boland N, Engineer F, Reisi M, Savelsbergh M, Waterer H (2011) Data generation in the Hunter Valley coal chain: a case study in capacity assessment. In: Proceedings of the 35th Application of Computers and Operations Research in the Minerals Industry Symposium. The Australasian Institute of Mining and Metallurgy, Wollongong, NSW, Australia, pp 795–806

Bolan N, McGowan B, Mendes A, Rigterink F (2013) Modelling the capacity of the Hunter Valley coal chain to support capacity alignment of maintenance activities. In: 20th International Congress on Modelling and Simulation. Adelaide, Australia, pp 3302–3308 AQ2

Montiel L, Dimitrakopoulos R (2013) Stochastic mine production scheduling with multiple processes: application at Escondida Norte, Chile. J Min Sci 49(4):583–597 (Pleiades Publishing)

Peng H, Zhou M, Liu M, Zhang Y, Huang Y (2009) A dynamic optimization model of an integrated coal supply chain system and its application. Min Sci Technol (China) 19(6):842–846 (Elsevier)

Singh G, Sier D, Ernst A, Gavriliouk O, Oyston R, Giles T, Welgama P (2012) A mixed integer programming model for long term capacity expansion planning: a case study from the Hunter Valley coal chain. Eur J Oper Res 220(1):210–224 (Elsevier)

Author Query Form

Book ID :	**457301_3_En**
Chapter No :	**20**

Please ensure you fill out your response to the queries raised below and return this form along with your corrections

Dear Author
During the process of typesetting your chapter, the following queries have arisen. Please check your typeset proof carefully against the queries listed below and mark the necessary changes either directly on the proof/online grid or in the 'Author's response' area provided below

Query Refs.	Details Required	Author's Response
AQ1	Please confirm if the corresponding author is correctly identified. Amend if necessary.	
AQ2	References 'Bolan et al. (2013)' is given in the list but not cited in the text. Please cite it in text or delete them from the list.	