Proceedings of the Twenty-Third International Conference on Automated Planning and Scheduling

Integrated Operations (Re-)Scheduling from Mine to Ship

S Kameshwaran and **Alfiya Tezabwala**
IBM Research - India, Bangalore
{kameshwaran.s, alfiya.tezabwala}@in.ibm.com

Alain Chabrier
IBM Software Group, Spain
achabrier@es.ibm.com

Julian Payne and **Fabio Tiozzo**
IBM Software Group, France
{julianpayne, tiozzo}@fr.ibm.com

Abstract

Mining companies have complex supply chains that start from the mining location and stretch thousands of kilometers to the end customer in a different country and continent. The logistics of moving the materials from mines to ship is composed of series of optimization problems like *berth allocation*, *ship scheduling*, *stockyard scheduling*, and *rail scheduling*, which are individually NP-hard. In this paper, we present a scheduling application, called as *IBM Optimization: Mine to Ship*, for end-to-end *integrated operations* scheduling. The application is built on IBM ILOG ODM Enterprise with advanced features like rescheduling under deviations and disturbances, and maintenance scheduling. The modeling and computational complexity of integrated scheduling optimization is tamed using hybrid optimization technique that leverages mathematical programming and constraint programming. The application will benefit the mining companies with increased resource usage, higher throughput, reduced cost of operations, and higher revenue.

Introduction

The Iron Ore Mining Industry

Around 98% of iron ore is used to make steel and goes directly to primary steel plants (MII 2006). It has been argued that iron ore is *more integral to the global economy than any other commodity, except perhaps oil* (FT 2009a). The trade in iron ore makes it the second-largest commodity market by value after crude oil. Unlike oil, iron is abundant and makes up 5% of the earth's crust. The difficulty is finding it in sufficient concentrations and then shifting millions of tons of dirt to where it is needed (TE 2012). The iron ore supply chain is about this *shifting*. The supply chain starts from the mine where the ore is extracted and extends till delivery at the steel plants. Most of the iron ore trade are global in nature, starting from Brazil and Australia, with most of the deliveries destined to China via sea (the so called seaborne trade).

Iron ore mining is a high volume low margin business, as the value of iron is significantly lower than base metals (FT 2009b). As most of the costs of an iron ore company are processing and logistics costs, managing the supply chain efficiently is one of the key drivers for the profitability of a mining company. Integration of the supply chain stages in the hands of a single company has been the key approach followed to increase supply chain efficiency. Given the high volume, low margin, and high investment nature of the business, the expected return on investments (ROIs) mandates a cost-efficient operations management of end-to-end supply chain.

Integrated Supply Chain and its Challenges

Following goals directly influence the ROIs that the stakeholders expect: *a)* Optimal balancing of capacity versus demand via a long/medium term supply chain planning, and *b)* high resource utilization and customer service via a proper short term scheduling and execution.

The mining companies generally focus on the long and medium term planning. The detailed short term scheduling are left to individual stages, with the mining company just handling the coordination across the stages. Lately, a couple of new issues are forcing companies to give more importance to short term scheduling.

The first issue is the growing fragmentation of the steel market. There were strong long term commercial relationships between the mining companies and few big steel companies. These long term contracts guaranteed steady supply chain flows that could be planned even years in advance. Now the rules of the game have changed: mining companies are forced to work with a lot of small spot orders coming from large number of clients sending their requests on short notice.

The second issue is that, after years of exploitation, the quality of the extracted material is deteriorating in many mines. This requires optimization of blending and processing activities to fulfill the target final quality and also reactive scheduling capabilities to quickly manage the quality unpredictability of the mines output.

In this paper, we focus on the short term scheduling, also called as *supply chain operations scheduling*. There are two main challenges to the integrated operations scheduling: *strong coupling across the stages* and *system wide percolation of deviations and disturbances*.

Strong Coupling The first issue is the strong coupling and cohesion of the different supply chain stages due to *short*

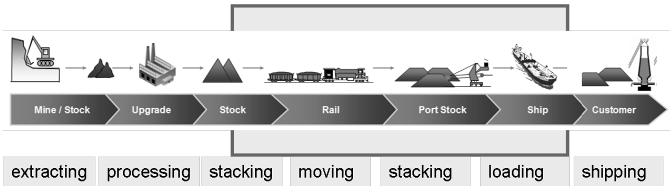

Figure 1: Mining Supply Chain

lead time and *limited inventory*. The lead time for moving the extracted material from mine to port is short and can be less than 24 hours. The intermediary storage is limited to two or three days of demand coverage.

Generally, the operations scheduling of the each stage of the supply chain is performed by a specific team, focusing on the respective stage. Coordination between the different teams ensure consistent and seamless movement of the material across the supply chain, but is not sufficient to avoid low equipment utilization rate due to poor synchronization of strongly connected upstream and downstream activities. In recent past, many mining companies have setup *remote control center* infrastructure for centralized coordination and control of supply chain activities from a single, remote location. New IT solutions for monitoring, planning, and scheduling that complement these centralized and integrated processes need to be used.

The market has responded with planning and scheduling solutions for mining supply chains which claim to schedule the end-to-end supply chain. However, so far the authors do no have a clear evidence of advanced end-to-end scheduling solutions that covers end-to-end and also can schedule activities to the resolution of individual resources.

Disturbances and Deviations The second issue is the critical reliability of equipment and infrastructure due to the tough operating conditions and intensive utilization. The equipments are capital intensive and there are no redundancies in the supply chain. Hence, whenever there are deviations (equipment under performance) and disturbances (equipment breakdowns), the schedule that is released for execution can be rendered inefficient or even infeasible. This gap between scheduling and execution is a major drawback in many commercial or custom-made applications that results in a rigid scheduling which is not agile enough to respond swiftly to disturbances in the supply chain.

In this paper, we present a scheduling system *IBM Optimization: Mine to Ship*[1] (MSS) developed at IBM for iron ore supply chain operations scheduling. We use the terms *solution*, *asset*, *system* interchangeably to refer to the MSS.

Integrated Operations Scheduling

Mine-to-Ship Supply Chain

A typical iron ore mining supply chain is shown in Fig. 1. The supply chain extends from the mines and ends at the customer points-of-delivery. The main component blocks of an iron ore supply chain are: *a)* the mines, where the iron ore is extracted, processed and stocked before to be loaded into the trains; *b)* the rail system, which transports the ore from the mines to the port; and *c)* the port where the material can be stocked and then loaded into the ships.

The scope of the MSS is highlighted in the figure, starting from the stockpiles in the mines till loading to the ships. The *supply schedule* (scheduled availability of materials at the mines) and the *demand schedule* (scheduled arrival of ships with demand for materials) are given. The problem that is addressed by MSS is the operations scheduling of moving materials from the mines (supply) to the ships at the port (demand) through intermediary stages of storage and transportation.

Planning, Execution, and Maintenance

There are three important business processes that govern the material movement, namely, *planning*, *execution* and *maintenance*. Planning is the process of *planning and scheduling* that comes up with the *schedule* to be executed. It is done either manually or through an automated system like MSS. The *schedule* is then communicated to the individual operations teams of respective stages for execution. The maintenance process deals with the monitoring the health of the equipments and attending to repair and maintenance. These three processes usually work in silos, leading to gaps and inefficiencies in the overall operations.

- The *schedule* given by the planning and scheduling system would not often be completely implementable, as the status of the operations in the field would be different from that of the time when the schedule was created.
- The maintenance is often reactive, and hence a equipment failure leads to break-down of original schedule, requiring manual intervention.

The scheduling systems generally focus on the *planning and scheduling* process that automates the derivation of schedules using AI, mathematical, and algorithmic techniques. In MSS, in addition to automated scheduling using math programming and constraint programming, we reduce the gaps between the above three business processes.

Deviations, Disturbances, and Rescheduling

Deviations and disturbances during execution can render the original schedule inefficient or infeasible. Due to the tight coupling between the stages, a deviation or disturbance in a specific stage, if not contained, can easily percolate to the upstream and downstream stages. This would often require *end-to-end* rescheduling that takes into account all the current activities across the supply chain and finds a new global schedule in response to the disturbance or deviation. This *reactive* or *dynamic* scheduling requires a *frozen horizon*, where no new activities are started and only the on-going activities are allowed to continue execution. The new schedule should be computed, communicated, and validated during this period. As the resources are under utilized during the frozen horizon, one expects short response time to minimize losses. Frozen horizon is a user-defined paramerter, but a

[1] http://www-01.ibm.com/software/data/information-agenda/catalog/profiles/Opt_Mind_to_Ship_IND_12.html

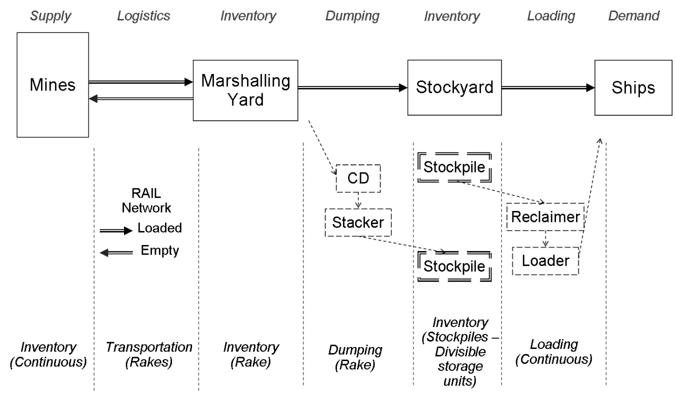

Figure 2: Material Flow

minimum/reasonable frozen horizon could be computed automatically.

Aims and Goals

The MSS asset was aimed at tackling the above inefficiencies, by providing following features: *a)* Integrated operations scheduling from mine to ship, covering the end-to-end operations; *b)* Minimize the *planning - execution* gap by creating plans that can be implemented with minimal manual intervention; *c)* Integrate maintenance with planning and execution; and *d)* Fast solution time that results in shorter *frozen horizon* and response time for (re-)scheduling.

Problem Description

Material Flow and Constraints

The schematic of the material flow supply chain (with the refined scope of MSS) is shown in Fig. 2. The system constraints are listed in the Table 1. We explain briefly the different stages and constraints in the supply chain, as shown in Fig. 2.

Mines (Supply) The inventory at the mines are the supply points with scheduled availability of materials. The *mine loaders* load the material from the mine on to the empty *rakes* for transportation.

Rail (Logistics) The rail logistics handle the transportation of the loaded material in rakes from the mines to the port.

Marshalling yard (Inventory) The rakes (loaded with material) that arrive at the port are stored in the marshalling yard.

Dumping The dumping activity empties the material from rakes and dumps it to a stockpile in the stock yard. The emptied rakes return to the mines. Material from a rake is emptied using a *car dumper* which is stacked on to a stockpile using an equipment called *stacker*.

Stockyard (Inventory) The stockyard consists of stockpiles of materials, where a stockpile has a designated material and a maximum capacity. It is the main storage unit at the port, where the materials arriving from the trains are stored in large heaps for loading on to the ships. The stockpiles are usually arranged in a grid, with rails operating between them for the movement of stackers and reclaimers.

Loading The material from a stockpiles are collected using *reclaimer* which is then transfered to the *ship loader* for loading on to the ships. We model one ship loader per pier. There are many reclaimers at stockyard, with one feeding to a ship loader at a time.

Input

The input consists of *static* data that remains unchanged for different scheduling horizons, *dynamic* data that are specific to a problem instance, *maintenance* data, and *WIP* data corresponding to work-in-progress (on-going activities).

Static Data The static data corresponds to supply chain configuration and capacities of constituent resources: number of mines; available products at mine; number and specifications of various equipments like mine loader, car dumper, stacker, reclaimer, and ship loader; rolling stock of engines and rakes; allowable train configurations (number of engines and rakes per train); description of the rail network; capacity of marshalling yard; stockyard configuration in terms of number of stockpiles and its material assignments; the number and capacity of piers.

Dynamic Data The scope of the dynamic data is defined by the *scheduling horizon*, which could be from hours to few weeks (as it is for short term operations scheduling). The schedule for periodic replenishment of the material stock at mines forms the supply input. Demand data consists of ships and orders: scheduled arrival of ships with information of expected time of arrival (ETA), demurrage date, end lay date (ELD), and orders per ship. The demurrage date is is an agreed upon date between the port and the customer ship lying between ETA and ELD . A ship when loaded and sailed out before the demurrage date earns the bonus for the company as against penalty which needs to be paid if ship sails out after the same. The rates of bonus and penalty are part of the input and total bonus earned has an upper bound. An order consists of material and demand quantity for a specific customer. Another important dynamic data is the high-tide schedule, which is port-specific and is made available in advance by the ports. Loaded ships can leave the port only in hightide to compensate for the loaded weight.

Maintenance Data Maintenance schedule describes the *planned* downtimes and expected performance degradations of various equipments during the scheduling horizon.

WIP Data WIP or *work-in-progress* captures on-going activities from which the status of the equipments and inventories are calculated. Schedule created without WIP would often not be implementable.

Decisions and Activities

Given the input, the problem is to determine the schedule of various activities in the supply chain during the scheduling horizon. The decisions and activities to be scheduled at each

Table 1: Constraints in various stages of supply chain

MINES: • Scheduled replenishment of stock at mine; specifies date, material and quantity • Scheduled maintenance of mine loaders • Each route from mine to port specifies admissible train configuration which gives number and capacity of rakes and locos to form a train
RAIL: • Headway constraints give time offset between two successive trains on the same segment or route • Rail capacity constraint gives max. number of trains simultaneously on a segment or route
MARSHALLING YARD: • Yard capacity in terms of max. number of loaded rakes, empty rakes and locos • Scheduled maintenance for car dumpers
STOCKYARD: • A stacker can access only certain stockpiles given by admissible stacker-stockpile configurations. Likewise, admissible stockpile-reclaimer configurations are defined • Max. capacity of of stockpiles and fixed assignment of material to stockpile • Stacker and reclaimer usually share a single rail to move between stockpiles, hence have movement restrictions for crossing • Min and max levels of stockpile for triggering refill and reclaim, respectively • Scheduled maintenance of stacker/reclaimer
PORT: • Given the max. capacity of pier and size of the ship, certain ship-pier allocations are admissible • Fixed number of sail in/out channels with max. number of ships simultaneously sailing in/out • Ship must be loaded and sailed out before its demurrage date to earn bonus, else incurs the demurrage penalty • End lay date sets the hard deadline before which a ship MUST sail out • Ship load plans (provided by each ship) specify the sequence of loading of orders to maintain ship balance • A loaded ship can sail-out only in a high-tide window given as high-tide time +/- X hours, where X is a function of ship capacity • Scheduled maintenance for ship loaders

Table 2: Decisions and Activities

MINES: • material assignment to rakes • scheduling rake loading activity
RAIL: • scheduling of loaded trains from mine to port • forming trains from empty rakes and locos at port • scheduling empty trains from port to mine
MARSHALLLING YARD: • scheduling rake dumping activity
STOCKYARD: • scheduling ship loading activity
PORT: • selection of customer demands • ship-pier (berth) allocation • scheduling ship sail-in, berthing and sail-out activities

of the stage is shown in Table 2. It is worth noting the mutual dependence of the decisions and activities due to the high cohesion between successive stages.

Objectives

The various decisions and activities are optimized with respect to a number of business objectives which are indicatives of the performance of the supply chain: *a)* Minimize penalty for rejected orders; *b)* Minimize demurrage penalty; *c)* Maximize demurrage bonus; and *d)* Minimize difference between achieved and target inventory levels.

Modeling and Solution Approach

Multistage Scheduling

The integrated operations scheduling is a series of interdependent decision and scheduling subproblems. We briefly discuss the subproblems and their interdependency below.

Demand Selection Given the schedule of ships and the customer orders therein, the demand selection is to select the orders that will be fulfilled. This is indeed dependent on the availability from the upstream nodes, arrival schedule of the ships, and the capacity and availability of the piers. The penalty for rejecting an order is user specified.

Ship Berthing and Scheduling The pier capacity constrains the ships that can be berthed and problem of allocating ships to piers is the NP-hard *berth allocation problem* (Imai, Nishimura, and Papadimitriou 2001; Umang, Bierlaire, and Vacca 2011). In our case, the berthing schedule is also contingent on the total order loading time (based on the customer orders accepted for that ship), sail-in channel capacity, and high-tide window for sail-out. Math programming is the predominant solution approach for ship berthing and scheduling (Buhrkal et al. 2011), while constraint programming has also been showed to be successful (Kelareva et al. 2012).

Stockyard Scheduling The stockyard operations consists of two main activities: *dumping* and *loading*. The material from rakes in marshalling yard dumped to a stockpile using a car dumper and a stacker. The material is loaded on to the ships from a stockpile using a reclaimer and ship loader. The two activities are interdependent due to the common inventory (stockpile) and stacker-reclaimer dependency on movement and operations. The stacker-reclaimer scheduling is NP-hard in nature (Hu and Yao 2010), even without other dependencies.

Rail Logistics In MSS, we address macro level rail logistics scheduling without considering the micro details of sidings, junctions, and control. The multi-level approach (Caimi et al. 2011) allows to work with aggregated information at the macroscopic level, which can be used for further scheduling with finer details. As the micro level information are specific to individual rail systems, the current implementation focuses only on macro level scheduling with provisions for customizing micro details later. At the macro level we focus on allowable train configurations between a mine - port route. A train configuration is the number of engine (of specific type/capacity) and number of rakes (of specific capacity). There are constraints on the rail network route like bound on the number of trains on a segment and headway constraints (time off-set between consecutive trains). The rail logistics also handles the reverse logistics of empty rakes from port to the mines. The reverse logistics is dependent on the dumping activities in stockyard operations as dumping results in empty rakes, which in turn is tied to forward logistics due to the rolling stocks of engines and rakes.

Mine Cluster Operations The MSS models supply as cluster of mines where the mines could be geographically separated. The supply schedule of materials at each mine is given as input. The decision problem at the mines is the allocation of material to empty rakes (supply to the port) and scheduling of loading the materials on the rakes.

The above constituent subproblems are NP-hard in general without any mutual dependencies. The benefits of jointly scheduling adjacent stages like rail and stockyard (Abdekhodaee et al. 2004) have been addressed to a limited extent but we are not aware of end-to-end integrated operations scheduling. A monolithic optimization or scheduling model to tackle the end-to-end integrated operations scheduling model is prohibitive in terms of modeling and solving. We decompose the problem to tame the modeling and computational complexity.

Decomposition

Decomposition is the standard technique in AI and operations research to handle complex interdependent problems. The two common decomposition techniques are *stage-wise decomposition* and *hierarchical decomposition*.

Stage-wise Decomposition Stage-wise decomposition is a *vertical* decomposition technique where the decisions and activities for a stage is completed before the adjacent stage. The decomposition can be from supply to demand (left to right), demand to supply (right to left), or hybrid (start with supply and demand, and make it consistent mid-way).

Hierarchical Decomposition Hierarchical decomposition is *horizontal* in nature, with higher level planning followed by lower level scheduling. The higher level decisions are taken for all the stages by *aggregating* the finer scheduling constraints (Pinedo 2009, Ch. 8). These decisions are then input to the lower level where the problem can split into smaller sub-systems, each dealing with respective constraints.

Hybrid Decomposition In MSS, we implemented a *hybrid* of the above techniques which consists of two *higher levels* and two lower level *sub-systems*. The two higher levels follow a stage-wise decomposition in order to fix the higher level planning decisions. The lower lever sub-systems use the output from the higher levels to create finer schedules. For example, the rail logistics is dealt at two levels. At the higher (planning) level, it is combined with supply and dumping operations. The decisions are number of trains per planning horizon per route and the material they carry. At the lower (scheduling) level, the decisions are formation of trains from the rolling stock of engines and rakes, and then scheduling them. The resulting sub-models have additional objectives and constraints in order to make the overall final solution consistent. Each of the levels and sub-systems are modeled as either a *linear program* or a *constraint program* depending on the structure of the problem. The optimization models are coded in OPL (optimization programming language) (Hentenryck 1999) and solved using the ILOG optimization engines. ILOG constraint programming optimizer provides a new scheduling language supported by a robust and efficient automatic search (Laborie 2009). All the scheduling subproblems in MSS are modeled as constraint programs.

Scheduling Modes

MSS supports *rolling horizon* scheduling by taking into account WIP and real time status of inventories and equipments from the execution and monitoring system. The user can do *dynamic rescheduling* when an unplanned disruption or deviation renders the current schedule inefficient or infeasible. The user can define a new scheduling horizon and frozen horizon, and MSS automatically processes the relevant data for end-to-end rescheduling. Proactive maintenance scheduling of equipments is also supported in MSS where the user can specify the maintenance duration and the time horizon within which it needs to be scheduled. The above features are enabled by the advanced data manipulation and visualization capabilities of ILOG ODME.

Implementation in ILOG ODME

IBM ILOG Optimization product line is dedicated to prescriptive analytics. This is the part of business analytics using operations research techniques to help decision makers take the right decision to optimally use their resources and

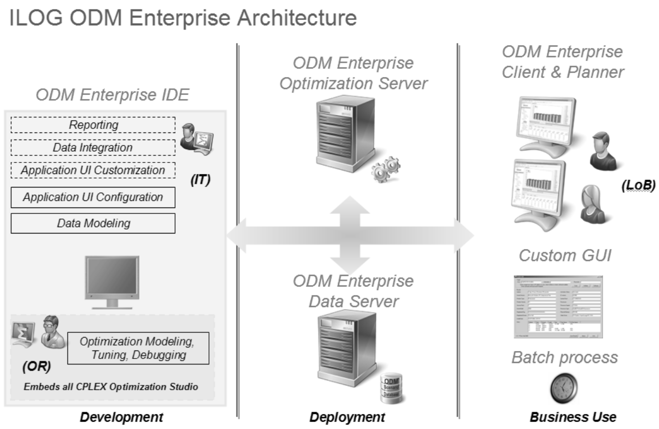

Figure 3: Architecture and application development with IBM ILOG ODME

reach their objectives in future. The IBM ILOG Optimization product line is organized in three areas:

- IBM ILOG CPLEX Optimization Studio is dedicated to operations research experts. It includes all the optimization technology required to develop, debug, tune, deploy and run optimization models. It contains engines covering linear, mixed integer, and constraint programming techniques.
- IBM ILOG ODM Enterprise is a platform dedicated to the development and deployment of complete decision support applications based on the optimization technology. In addition to all functionalities from the CPLEX Optimization Studio, it also facilitates data integration and data management of the application, creation of user interfaces, and scalable deployment of the application.
- Industry assets, developed on top of the platform, provide a significant amount of the required business functionality (from optimization model to business data model through user interfaces) for a specific set of industries.

The MSS is implemented as an application in ILOG ODM Enterprise (ODME). The ODME architecture and constituent components used in application development is shown in Fig. 3. ODME provides infrastructure for creating a business application from a purely mathematical operations research model. An ODME application is created by starting the definition of an *application data model* (ADM). The ADM includes data wider than what is strictly required for the optimization model as the application may allow other decision support mechanisms than pure optimization (e.g. manual planning). Also ADM integrates with the legacy systems of the users and hence the data format can be same as that of the users, which can be preprocessed within ODME to suit the requirements of the optimization model.

The decomposition approach described in the previous section comprises of few constraint and mathematical programming models with complex control and data flow between them. The individual models are coded in OPL with the control and data flow between the models managed by ILOG script. The optimization logic forms the core of the application, with appropriate pre-processing and post-processing of data from and to ADM.

In ODME parlance, a business centric data set (a problem instance from optimization perspective) is referred as a *scenario*. The ODME *data server* manages different scenarios. User interface, data integration, and optimization are mapped to ADM. When the MSS user requests a scenario (problem instance) to be scheduled, the current scenario data is mapped to the optimization data through preprocessing. The preprocessing in MSS not just transforms scenario data that the OPL models can understand. It also includes complex operations that identifies WIP data; flags dependent activities that should follow WIP activities; and computes equipment status and inventories during the frozen horizon.

Similarly, after optimization, the results are post-processed to be stored in the ADM in a business oriented way. Optimization can be run locally or on an ODM Enterprise Optimization Server. Optimization servers can be set up in clusters so that the applications can scale with the needs of rescheduling scenarios and/or offer fail-over functionalities for high-availability. The integration of the solution into the broader infrastructure can be done through simple data integration from/to the data server. The user interface is directly connected to the ADM, allowing possible scenario manipulation without interference with optimization.

ODME provides powerful scenario management capabilities to do various what-if analysis on the schedule given by the optimizer. For example, analyzing the effect of increased downtime of an equipment on number of rejected orders can be done visually. Several output schedules can be visually compared against each other, making it easier for the end user to select one of the several available activity schedules.

Visualization

Visualization in ODM Enterprise

As part of the platform, ODME provides some standard visualization components as well as a framework for building custom graphical views. These custom graphical views can be built with any Java based graphical components including IBM ILOG JViews that is bundled as part of ODME.

Visualizing the Supply Chain

In order to effectively manage the supply chain it is important that organizations can visualize the supply chain to gain better insights on the structure of the supply chain and the flow of materials in the supply chain. Visualization allows users to better understand solutions that are generated by optimization-based tools thereby increasing their confidence in the solution as well as providing simple guided mechanisms to make changes to the generated plans. Visualizing the supply chain requires visualizing the flow of material, in this case from the mines, to the trains then to the stockyards and finally on to the ships. Given the potential size of the supply chain it is also important that different filtering mechanisms are provided that allow an user to focus on the key business factors that interest them, this may include

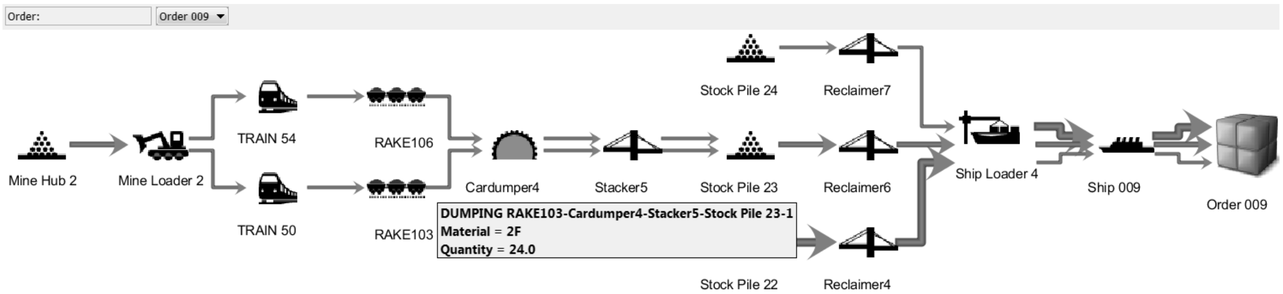

Figure 4: Material flow from mines to ship

looking at particular customer orders, looking at orders for a given ship, looking at orders related to a given mine or material and so on.

For example, an user that is interested in visualizing the material flow for a given customer order can use a custom diagram to visualize just the information related to this order. In addition an user can get information through *tool tips* and other mechanisms at each stage of the supply chain, this is illustrated in Fig. 4. This chart enables the user to have end-to-end visualization of material flow across all the stages, specifying activities that took place to satisfy one specific order.

Users will also need to be able to visualize the availability of material at each stage of the supply chain in order to verify availability for new orders as well as making decisions in the case of disturbances in the supply chain. Fig. 5 shows an example of a custom view showing the inventory levels of various products at the different stages of the supply chain with the possibility to filter based on a location, a material and a sub location.

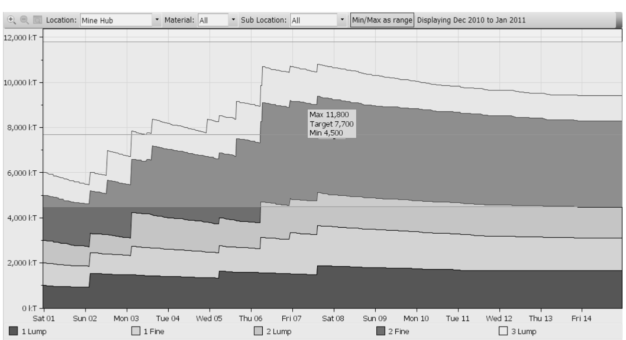

Figure 5: Material availability at different stages of the supply chain

Gantt Chart

A common visualization for a scheduling solution is a Gantt chart that allows the user to see each operation, the start and end time of the operation, the equipment used by an operation as well as other information related to the solution, in this case for example the related customer order or the type and quantity of material. An example of such a view is shown in Fig. 6, this diagram shows the different operations that are being performed in the supply chain. Each operation is colored based on a criteria specified by the user, in this case operations are colored by order number. Finally pegging logic is used to calculate and display the relationship between operations for a given customer order.

Visualizing Disturbances in the Supply Chain

As has been already mentioned, an important part of MSS is the ability to understand and manage disturbances in the supply chain. In order to quickly respond to the disturbances, the users need to be able to visualize the disturbance event as well as the potential impact of the event. Once the impact has been understood then an user may want to make manual modifications to the system or reschedule a solution in order to take into account the disturbance. In both cases the user needs visualization tools that allow easy identification of the disturbance as well as real-time guidance when making changes to a plan.

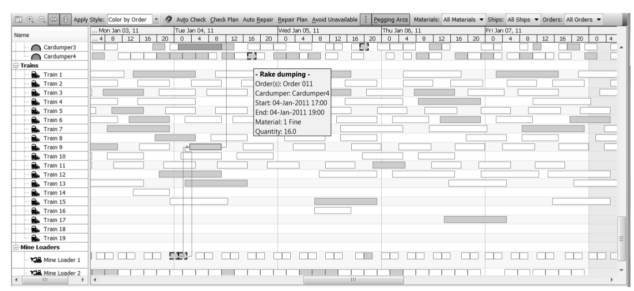

Figure 6: Visualizing material movements in a Gantt chart

In order to visualize disturbances the user needs to be able to compare the initial plan with the actual plan as well as disturbances. In this case it is possible to load additional information into the Gantt chart related to the work in progress as well as disturbances. In this example the disturbances are unplanned maintenance tasks. Fig. 7 shows an example where the tasks that have been finished are in gray shaded region, tasks that are in progress are hashed and unplanned maintenance tasks are displayed with a gray dotted pattern.

In some cases an user may want to make manual modifications to the plan and the solution needs to be able to guide the user in real time as changes are made. Fig. 8 shows an example of an user making a manual modification to a car

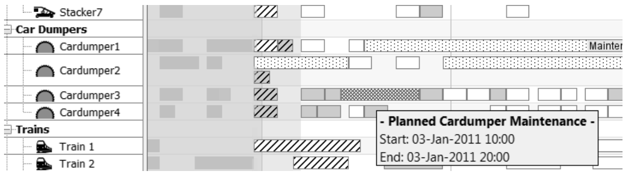

Figure 7: Visualizing work in progress and un-planned maintenance periods

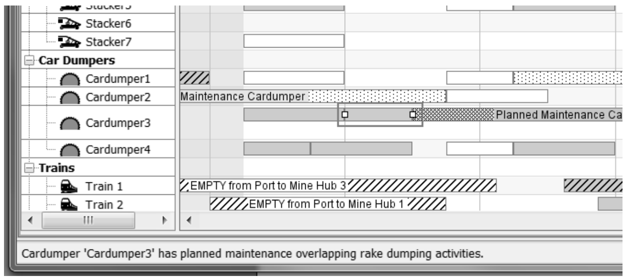

Figure 8: Managing manual modifications

dumping activity and a real time status message that tells the user why a move is not possible in the case that the move would violate a constraint in the system. This optional mode allows an user to make a manual modification whilst avoiding any conflicts with the current plan, if necessary an user can override this feedback and make the modification as required.

Computational Experiments

In this section, we share the results of computational experiments conducted with a synthetic data set that closely resembles a typical real world iron ore supply chain. The supply chain configuration from the synthetic data set is shown in Table 3. The decomposition has planning models and scheduling models that can work with different time granularities. We used one minute granularity for scheduling models and the planning models worked with the higher aggregated time granularity of six hours (the planning horizon is divided into planning periods of six hours each). Two weeks of scheduling horizon was used with 30 ships of different capacities, each with one customer order. The total number of binary, real, and interval (for constraint programming) variables is given in Table 4.

The testing of optimality of the decomposed optimization model is not straight-forward due to the following reasons. Firstly, the decomposition trades-off optimality for the reduced complexity in modeling and the quantification of optimality gap is difficult. Secondly, we use constraint programming, which takes very long time to prove optimality and hence it is mainly used with limit on the solution time. We used indirect measurement of optimality by comparing the maximum utilization of all the stages in the supply chain with different initial conditions. By varying the inventory levels (stockyard and marshalling yard) and distribution of rakes and engines at the port and mines, we created various initial conditions. The matrix of computational experiments showed that a time limit of 180 seconds can ensure a good solution for any of the initial conditions. The experiments were conducted with ILOG ODM Enterprise 3.6, running on Windows 7, with Intel i7 and 4 GB RAM. All the experiments had no WIP data to ensure uniformity. The rescheduling test cases also showed that end-to-end rescheduling can be done within 180 seconds thus providing a quick response for dynamic rescheduling.

Table 3: Synthetic Data Set

Mines	Products	Mine Loaders	Rakes
3	5	3	44
Locos	Train Config.	Car Dumpers	Stockpiles
23	3	4	30
Stackers	Reclaimers	Piers	Ship Loaders
7	7	4	4

Table 4: Number of Variables and Constraints

# Variables			# Constraints
Binary	Real	Interval	
1834	3090	19336	42616

Conclusions

In this paper, we present a scheduling application *IBM Optimization: Mine to Ship*, built for integrated operations scheduling of iron ore mining supply chain. The asset is implemented using IBM ILOG ODM Enterprise, a platform to develop industry solutions by providing complex data manipulation, advanced scenario management and visualization functionalities, in addition to core optimization capabilities.

Currently the application is referred as an industry (service) asset, rather than as a product to differentiate the off-the-shelf commercial products. The features of commercial products are nearly hard-coded that is uniform for all users. On the other hand, an industry asset allows customization for specific client needs, which is indeed the requirement for a complex business case like integrated operations scheduling of mining supply chains.

The optimization models in OPL are easily extensible to accommodate additional constraints and objectives that are specific to individual mining companies. The decomposition approach is modular and one can easily modify the asset to companies that require subset of functionalities like only ship berthing for a port or a mining company with third party rail logistics that does not require the rail scheduling part. Further, MSS is also available, under special agreement, with source code for customization to suit client specific requirements. The asset is expected to benefit the delivery of advanced customized solutions to mining companies, solutions that are expected to deliver increased resource efficiency, higher throughput, reduced cost of operations, and higher revenue.

Acknowledgements

We would like to acknowledge the contributions of Krishna K. Nagarajan (Client Technical Architect, IBM Singapore), Kizhanatham Srikanth (Mining Solution Specialist, IBM Australia), Amit Merchant (CTO, India Design and Tools, IBM India), Kalpana Doraisamy (Solution Representative, IBM India Software Lab), and Prasanna S. Venkataraman (Global Business Services, IBM India).

References

Abdekhodaee, A.; Dunstall, S.; Ernst, A. T.; and Lam, L. 2004. Integration of stockyard and rail network: A scheduling case study. In *Proceedings of the Fifth Asia Pacific Industrial Engineering and Management Systems Conference*, 25.5.1 – 25.5.17.

Buhrkal, K.; Zuglian, S.; Ropke, S.; Larsen, J.; and Lusby, R. 2011. Models for the discrete berth allocation problem: A computational comparison. *Transportation Research Part E: Logistics and Transportation Review* 47(4):461 – 473.

Caimi, G.; Fuchsberger, M.; Laumanns, M.; and Schupbach, K. 2011. A multi-level framework for generating train schedules in highly utilised networks. *Public Transport* 3(1):3 – 24.

FT. 2009a. Iron ore pricing emerges from stone age. *Financial Times*.

FT. 2009b. Iron ore pricing war. *Financial Times*.

Hentenryck, P. V. 1999. *The OPL Optimization Programming Language*. Cambridge, Massachusetts: The MIT Press.

Hu, D., and Yao, Z. 2010. Stacker-reclaimer scheduling for raw material yard operation. In *International Workshop on Advanced Computational Intelligence*, 432 – 436. IEEE.

Imai, A.; Nishimura, E.; and Papadimitriou, S. 2001. The dynamic berth allocation problem for a container port. *Transportation Research Part B: Methodological* 35(4):401 – 417.

Kelareva, E.; Brand, S.; Kilby, P.; Thiebaux, S.; and Wallace, M. 2012. CP and MIP methods for ship scheduling with time-varying draft. In McCluskey, L.; Williams, B.; Silva, J. R.; and Bonet, B., eds., *ICAPS*. AAAI.

Laborie, P. 2009. IBM ILOG CP Optimizer for detailed scheduling illustrated on three problems. In *Proceedings of the 6th International Conference on Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, CPAIOR '09, 148–162. Berlin, Heidelberg: Springer-Verlag.

MII. 2006. Iron ore - hematite, magnetite & taconite. *Mineral Information Institute*.

Pinedo, M. L. 2009. *Planning and Scheduling in Manufacturing and Services*. Springer.

TE. 2012. The lore of ore: The most important commodity after oil deserves more attention than it gets. *The Economist*.

Umang, N.; Bierlaire, M.; and Vacca, I. 2011. The Berth Allocation Problem in Bulk Ports. Swiss Transport Research Conference, STRC, Monte Verita / Ascona, Switzerland, May 13, 2011.