



**WORKING PAPER**

**ITLS-WP-17-06**

**Assessment of the Hunter Valley  
Coal Export Supply Chain**

**By  
Natashia Boland<sup>1</sup>, Mohsen Reisi<sup>2</sup> and  
Martin Savelsbergh<sup>3</sup>**

<sup>1</sup> & <sup>3</sup> H. Milton Stewart School of Industrial and Systems  
Engineering, Georgia Tech, USA

<sup>2</sup> The University of Sydney, Australia

**March 2017**

**INSTITUTE of TRANSPORT and  
LOGISTICS STUDIES**

The Australian Key Centre in  
Transport and Logistics Management

The University of Sydney

*Established under the Australian Research Council's Key Centre Program.*



**NUMBER:** Working Paper ITLS-WP-17-06

**TITLE:** **Assessment of the Hunter Valley Coal Export Supply Chain**

**ABSTRACT:** We develop a decision support tool that assesses the throughput of a coal export supply chain for a given level of demand. The tool can be used to rapidly evaluate a number of infrastructures for several future demand scenarios in order to identify a few that should be investigated more thoroughly using a detailed simulation model. To make the natural model computationally tractable, we exploit problem structure to reduce the model size, and we employ aggregation as well as disaggregation to strengthen the structure of model. We use the tool in a computational study in which we analyze system performance for different levels of demand to identify potential bottlenecks.

**KEY WORDS:** *Supply Chain Optimisation \_ Integer Programming \_ Coal Export*

**AUTHORS:** **Boland, Reisi and Savelsbergh**

**CONTACT:** INSTITUTE OF TRANSPORT AND LOGISTICS STUDIES (H73)  
The Australian Key Centre in Transport and Logistics Management  
The University of Sydney NSW 2006 Australia  
Telephone: +612 9114 1824  
E-mail: [business.itlsinfo@sydney.edu.au](mailto:business.itlsinfo@sydney.edu.au)  
Internet: <http://sydney.edu.au/business/itls>

**DATE:** March 2017



## 1 Introduction

Since the global population is growing, and living standards are improving in developing countries, the demand for energy is increasing. Coal is still considered a necessary source of energy to satisfy future demand. Over the past 10 years, Australia's coal exports, to countries such as Japan, China, Korea, India and Taiwan, have increased by more than 50 percent. Demand for coal is still expected to increase dramatically over the next 10 years in China and India (<http://www.australiancoal.com.au>).

The Hunter Valley Coal Chain (HVCC) is the largest coal export operation in the world. It comprises the transport of coal from mines located in the Hunter Valley to the Port of Newcastle in New South Wales, Australia, and is a complex system, involving 11 producers operating 35 coal mines, 27 coal load points, 2 rail track owners, 4 above rail operators, 3 coal loading terminals with a total of 8 berths, and 9 vessel operators. Approximately 1700 coal vessels are loaded at the Port of Newcastle each year, and the throughput of the port has increased rapidly, from around 92 million tonnes in 2008 to 169 million tonnes in 2014.

After coal is mined, it is stored temporarily at a load point facility, before transport to one of the terminals at the Port of Newcastle, almost exclusively by train, although small quantities are carried by truck. Coal is unloaded at the terminal via a dump station, then transported by conveyor belt to a stacking machine, which places the coal on the ground so as to form stockpiles. Coal from different mines with different characteristics is mixed in a stockpile to form a coal blend that meets the specifications of a customer. Once a vessel arrives at a berth at the terminal, the stockpiles with coal for the vessel are reclaimed and loaded onto the vessel. The vessel then transports the coal to its destination.

The coal export supply chain is managed by the Hunter Valley Coal Chain Co-ordinator (HVCCC). One of the most important and far-reaching decision problems faced by HVCCC is the planning of long-term capacity. The demand for coal continues to grow and thus export through the Port of Newcastle is expected to increase in the future. Even the optimal use of the existing infrastructure will not be sufficient to accommodate anticipated future demand. Therefore, the infrastructure needs to be upgraded and the capacities in the system expanded. As upgrading infrastructure and expanding capacity is extremely expensive, a careful and thorough system analysis is crucial to ensure that investments are made in the right place and at the right time.

The HVCCC uses an elaborate and detailed discrete event simulation model of the HVCC, developed initially by ?, to analyze and assess the throughput of the system, to detect any bottlenecks in the system, and to investigate the benefits of infrastructure upgrades and expansions. As the simulation model is very detailed, it takes a considerable amount of time to run, and as a consequence, few scenarios can be analyzed. Given the number of possible infrastructure upgrades and expansion options (e.g., load point capacity, track capacity, dump station capacity, stacking capacity, stockyard capacity, reclaiming capacity, ship loading capacity, and berth capacity) and the number of future demand scenarios of interest, the long run times of the discrete event simulator pose a major challenge for the HVCCC.

Our research aims to complement the simulation tool with an integer programming based decision support tool that more quickly determines the throughput of the system for a given level of demand. It may thus be used to rapidly evaluate a number of infrastructures for several future demand scenarios in order to identify a few that may be investigated more thoroughly using the simulation model.

The paper focuses on two main topics: the design and implementation of an integer programming model for assessing the performance of the HVCC under a given infrastructure setting, and a computational study of its performance in response to variations in the nature and timing of the demand presented to it.

To make the natural integer programming model computationally tractable, we strengthen several classes of constraints, explore variable aggregation to eliminate

symmetry from the model, and develop several strategies to reduce the size of an instance. These measures combine to produce a speed-up in solution time of more than 40-fold, on average.

We use our integer programming model in a computational study in which we compare the results produced by the model for several different demand scenarios, using a variety of metrics, such as the percentage of the vessels served after their expected due time and the average daily stockyard utilization in each terminal, to identify and analyze bottlenecks in the system. The results show that the bottlenecks are in the outbound part of the system, with one terminal suffering a lack of capacity at the berths, and the other a lack of stockyard capacity.

The remainder of the paper is organized as follows. Section 2 describes the HVCC and the role of decision support tools in strategic planning for the system. Section 3 presents a review of relevant literature. Section 4 introduces an integer programming model for assessment of the performance of the HVCC, and develops solution acceleration strategies. Section 5 discusses the results of our computational study. Finally, Section 6 provides some concluding remarks.

## 2 Problem description

In this section, we describe the infrastructure components of the HVCC, as well as the nature of the demand on the system, in more detail. We also review relevant planning processes and tools. In so doing, we introduce terminology that will be used throughout the paper.

### 2.1 The Hunter Valley Coal Chain: An Overview

The HVCC refers to the inland part of the coal chain in the Hunter Valley. The coal transport system follows the Hunter River from mining areas in the Hunter Valley to the port at Newcastle. After coal is mined, it is stored either at a coal loading facility shared by several mines, or at a railway siding located at the mine, before transport to the terminals, primarily by rail, where the coal is unloaded at the terminal dump stations and stacked as stockpiles in the stockyard. Once the vessel for which the coal is destined arrives at the berth, the coal is loaded onto the vessel.

The trains used to transport coal have up to 6 locomotives and up to 148 wagons each, and a length of up to 2 kilometers. Each train carries about 8,500 tonnes of coal, on average. Each trip by a train consists of travel empty from a terminal to a load point, loading, at a single load point only, travel to a (single) terminal, and dumping of the entire load at that terminal. For efficiency, trains are nearly always loaded to their full capacity.

The Kooragang Coal Terminal, (KCT), on Kooragang Island, and the Carrington Coal Terminal, (CCT), in the suburb of Carrington, are two terminals at the port of Newcastle, located on either side of the Hunter river. The terminals are owned and managed by Port Waratah Coal Services Limited (PWCS), which is an unlisted public company providing equitable access to export terminal capacity. In each terminal, there are facilities for unloading and storing the coal on the stockyard and reclaiming and loading the coal onto the vessels. A new terminal,

(NCIG), owned and operated by a consortium of coal producers, the Newcastle Coal Infrastructure Group, was recently established to help accommodate export growth, providing access to additional terminal capacity for its members.

The HVCC comprises the mining companies, the track owners, the rail operators, PWCS, and the NCIG. Coordinating the activities of these organizations is necessary to the efficiency of the supply chain. The HVCCC is responsible for integrated planning and scheduling for the HVCC. Their job includes strategic modeling in support of infrastructure expansion, as well as operational scheduling.

The HVCCC plans and schedules all movements of coal in the system, but it is not responsible for the plan execution. The “inbound” plan created by the HVCCC covers the transportation of coal from the mines to the stockyards, including the activities of the load points, dump stations, and stacking machines, as well as train scheduling. The inbound plan covers 36 hours of operations and is ready 24 hours before its execution time. The HVCCC also creates “outbound” plans, which cover the movement of coal from the stockyard onto the vessels, including the reclaimers and ship loaders activities. The outbound plan also covers 36 hours of operation and is released 7 hours before it goes into effect. The “Live Run” team at PWCS, (together with planners at NCIG), is responsible for executing the plans that are provided by the HVCCC.


The HVCC is a *pull* system in which activities are driven by the customers’ demands, in contrast to a *push* system that first produces products, and then tries to find customers for them. In the HVCC, the demand is specified in the form of a *shipping stem*, which is a list of anticipated vessel arrivals. Each vessel arrival, referred to as a *trip*, is characterized by (i) an arrival time, that is the date and time at which the vessel will arrive at the port, (ii) the terminal at which the vessel is to be loaded, and (iii) a *cargo-profile*, which specifies the brands of coal and the tonnage of each that is to constitute the vessel’s cargo. The combination of a brand and the tonnes of that brand in a cargo-profile constitute a single *cargo*. For example, a trip may have a cargo-profile consisting of two cargoes, in which 100,000 tonnes of brand X is one cargo and 30,000 tonnes of brand Y is the other. Since coal is a blended product, each brand in a cargo-profile comes with a *brand-recipe*, which specifies a set of mines and their percentage contribution to the brand. For example, a recipe for brand X might be 25 percent from mine A and 75 percent from mine B. A *component* is the material in a cargo that comes from a particular mine, so a cargo consisting of 100,000 tonnes of brand X with the given brand-recipe will have two components, one consisting of 25,000 tonnes from mine A and the other 75,000 tonnes from mine B.

The purpose of building a stockpile is to store the coal until its vessel is ready to be loaded, and to mix the coal from different sources so that, at the time of reclaiming, a specific blend is achieved. Importantly, the HVCC operates largely using what is known as *cargo assembly*, meaning that each cargo for each vessel trip is assembled as a stockpile that is built specifically for that trip. For example, if a vessel trip requires 100 kt of brand X, with brand-recipe 25% from mine A and 75% from mine B, then (assuming average train tonnages), approximately 3 train loads from mine A and 9 train loads from mine B will be combined in a single stockpile, prior to the arrival of the vessel. The entire stockpile will then be loaded onto the vessel. Typically, each cargo corresponds to a single stockpile, but, occasionally, a cargo may be split into two or more smaller stockpiles so as

to fit better in the stockyard. Building a stockpile for a given vessel trip usually only starts quite a short time, (a matter of days, rather than weeks), before the vessel is expected to commence loading. We refer to the period of time between the unloading of the first train for a stockpile, and when the stockpile has completed loading onto its vessel and the stockyard space it has occupied cleared, as the *stockpile tenure*.

A complex infrastructure supports the coal export supply chain in the Hunter Valley. Coal moves from the load points at the mines along the rail tracks to the dump stations at the terminals, passing through a number of rail junctions on the way. From the dump stations, conveyors take the coal to stacker machines that assemble the stockpiles on pads in the stockyard. Blended coal is subsequently reclaimed from the pads in the stockyard by reclaiming machines and transported via conveyors to the ship loader machines at the berths to finally be loaded onto waiting vessels. Each train load is stacked onto a single stockpile using a single stacker machine. Each stockpile is reclaimed using a single reclaimer and ship loader combination.

The following list describes the infrastructure components and their characteristics in more detail. A schematic, showing how these are related, is given in Figure 1.



**Fig. 1** A schematic view of how different infrastructure elements combine in the HVCC.

- *Load point*: a facility at which trains are loaded with coal, typically collocated with, or near, a mine. In many cases, there is a one-to-one correspondence between a mine and a load point, but in some cases a load point serves more than one mine. Each load point has a specified maximum load rate, i.e. a maximum tonnes per hour at which it can load coal onto a train, and is usually assumed

to operate at this rate. For each load point, there is a set of possible train configurations that can be served by that load point, where a configuration is a combination of number of locomotives, number (and capacity of) wagons, and the rail operator. When planning, it is usually assumed that each load point has a single, preferred, train configuration. Since, on each train trip, the train is nearly always loaded to capacity, this implies that each train served by a load point has a specified tonnage, and takes a specified time to load.

- *Junction*: a junction connects two rail track segments and is used to model rail track capacity. Specifically, a junction capacity limits the number of trains that can pass through the junction in a day.
- *Stockyard*: the area at a coal terminal where coal is stacked and reclaimed, i.e., where coal blends are assembled.
- *Pad*: a specific area in the stockyard where stockpiles of coal are assembled. Each pad has a given length and width (in meters). Pads are typically arranged in parallel in the stockyard (see Figure 1). Each stockpile takes the entire width of a pad, and occupies a certain amount of its length, which can be derived from the tonnage of the stockpile and the width of the pad. Because of the stacking method generally used in the HVCC, a stockpile occupies the same length of the pad for its entire tenure.
- *Dump station*: a facility where trains are unloaded. Each dump station has specified maximum dumping rate (given in tonnes per hour). When a train is being unloaded at a dump station, the coal is immediately transported to the stockyard by a conveyor belt.
- *Stacker*: a machine that is used to stack coal on the pads in the stockyard of a terminal. Each stacker has a specified maximum stacking rate (given in tonnes per hour). Some stackers can reach only one pad, while others can reach two, and some stackers share a conveyor belt, and so cannot operate simultaneously, while others have a dedicated conveyor belt (see Figure 1).
- *Reclaimer*: a machine that is used to reclaim stockpiles from pads in the stockyard. Each reclaimer has a specified maximum reclaiming rate (given in tonnes per hour). As shown in Figure 1, each reclaimer can reach only specified pads, two each in the PWCS terminals, in which each also has a dedicated conveyor belt.
- *Ship loader*: a machine that loads coal onto a vessel. Each ship loader has a specified maximum loading rate (given in tonnes per hour).
- *Berth*: a place where a vessel can be loaded. Each terminal has a fixed, specified, number of berths.

In long-term planning, it is generally assumed that each dump station, stacker, reclaimer and ship loader, will operate at its maximum rate, and is available for a specified number of working hours per day. Both stackers and reclaimers need time to move between working on different stockpiles; we discuss modeling of this time in Section 4. In each PWCS terminal, there is a ship loader for each berth, these have identical rates, and any reclaimer can feed coal to any of the ship loaders. Thus the outbound loading capacity is typically represented as a single *reclaim stream* rate, with one rate given for each reclaimer.

If the cargo-profile of a vessel includes more than one cargo, and, as a result, multiple stockpiles have to be reclaimed and loaded onto the vessel, this has to be done in a pre-specified order to ensure the vessel stability. Furthermore, in

practice, the loading of a vessel will only commence when all of its stockpiles have been fully assembled so as to avoid the risk of an interruption in the loading of a vessel due to coal supply interruptions. The assembly of a stockpile takes between three to seven days depending on the coal components, (some mines are located several hundreds of miles away from the port), and therefore, the assembly of the stockpiles of a vessel starts about five to ten days before its arrival time.

PWCS terminals operate primarily what is known as a *turn of arrival* service, in which there is a preference to load vessels at each terminal in order of their arrival time at the port. However, this order is not strictly applied, especially when efficient operations of the system dictate otherwise; in planning and scheduling of operations, some flexibility is employed to reduce overall vessel delays. Indeed, the direct costs of vessel delays are incurred as *demurrage*: a vessel-specific cost per unit time is applied to any time from some stipulated due time to the actual completion time of vessel loading. The due time is usually taken to be the arrival time at the port, plus some length of time that is contractually specified, but which reflects what is reasonably necessary for loading the vessel. For example, 82 hours is a typical figure.

A primary goal of the HVCCC is to increase the throughput of the system, i.e., increasing the total number of tons that can be exported per year. However, these tons are not homogeneous, but are presented in the form of shipping stems, and so must be delivered in this form, in a timely fashion, without introducing significant demurrage costs. Thus, for modeling in support of strategic planning, and in operational scheduling, the HVCCC usually seeks to minimize total delays in the completion of vessel loading.

## 2.2 Planning and Scheduling in the Hunter Valley Coal Chain

One of the most important and far-reaching decision problems faced by the HVCCC is the planning of the long-term capacity. The Australian Rail Track Corporation Ltd has forecast the growth of the contracted export coal volumes in the Hunter Valley from 158 mtpa in 2013 to 206 mtpa in 2019 (<http://www.artc.com.au/library/2013%20HV%20Strategy%20-%20Final.pdf>). This is a downward revision from earlier forecasts, which reached 250 mtpa by 2018, (see, for example, discussion in Bayer et al. (2009), which also identifies some of the major infrastructure projects planned to achieve these volumes), but still represents significant growth over a relatively short period. Thus coal export through the Port of Newcastle is expected to continue to increase, apace, and even the optimal use of the existing infrastructure will not be sufficient to accommodate anticipated future demand. Therefore, the infrastructure needs to be upgraded, and the capacities in the system need to be expanded. As upgrading the infrastructure and expanding the capacity is extremely expensive, a careful and thorough system analysis is crucial to ensure that investments are made in the right place and at the right time. The possible upgrading and expanding options include load points improvements, new train control technology, track capacity upgrades, additional tracks, new passing loops, new overpasses, new dump stations, additional stockyard space, new stackers, new reclaimers, and new berths.

In order to manage a complex system, with many sub-systems, such as the HVCC, decision support tools are highly advantageous. In the following section, we discuss one of the decision support tools that is currently used by the HVCCC.

### *2.2.1 Simulation*

The HVCCC uses a discrete event simulation model of the HVCC to analyze and assess the performance of the system, to detect and identify any bottleneck in the system, and to investigate and explore the benefits of infrastructure upgrades and expansions. (? discuss the original development of the model and its first use in this way.) The assessment is performed using demand in form of a shipping stem. Creating a shipping stem that is a reasonably accurate representation of expected future demand is a significant challenge in its own right. ? discuss an integer-programming based hierarchical approach to generating one or more shipping stems for a given level of expected demand, (in the form of forecast mine production tonnages), that are realistic and reflect characteristics of historical shipping stems.

The simulation tool has two main modules:

- a master schedule generating module, which processes a shipping stem to provide a schedule for delivering coal from mines to the terminals, taking into account capacities in the system, including capacity reductions due to planned maintenance events, and
- a dynamic network operation module, which evaluates the performance of the master schedule under dynamic incidents, such as train delays.

The simulation model requires demand data, load point data, track data, train data, terminal data, and maintenance data. It generates key performance indicators, such as throughput, ship queue lengths, ship turn-around times, ship loading times, demurrage cost, train cycle times, cargo build times, stacker utilization, and ship loader utilization. The model is validated using historical data reflecting alternative anticipated operating scenarios. For the validation, annual export tonnage, delivery performance (in terms of late arrivals for each mine and overall late arrivals), average daily tonnage delivered, average number of trains per days, and train utilization are used.

As the simulation model is very detailed, it takes a considerable amount of time to run, and to analyze its output. Ideally, the HVCCC would like to consider a large number of possible infrastructure upgrade and expansion options, and, in order to be confident of the robustness of any recommended infrastructure plan, to evaluate these under a large number of alternative future demand scenarios. However, this would result in too many potential scenarios to be analyzed, in practice, with the simulation model.

### *2.2.2 Performance Assessment Tool*

Our research aims to complement the simulation tool with an integer programming based decision support tool that more quickly estimates the performance of the system for a given level of demand. Although it does not take into account all features that the simulation model does, and so may be less accurate, it is designed to

be sufficiently detailed to provide a meaningful discriminator between alternative infrastructure expansion options, and has the advantage of running fast enough to allow each option to be assessed under many more demand scenarios than is practical with the simulation model. Its purpose, for the HVCCC, is to identify a few infrastructure settings that may be investigated more thoroughly using the simulation model.

The integer programming (IP) model covers the entire system from the load points to the terminals. It represents load point facilities, rail tracks, trains and wagons, dump stations, stackers, pads, reclaimers, ship loaders, and berths. For the sake of efficiency, the IP model does not consider operational details, but includes a reasonable level of detail for tactical and strategic decision making. The balance between accuracy and efficiency is reflected in our choice of time granularity: the planning horizon is divided into days, and all decisions are at the daily level.

In order to measure the performance of the system, total delay, which is the sum of all vessels delay, the average delay per vessel, the daily production, which is daily stockpile tonnage and number of trains per day, and the number of vessels in the queue on each day, are reported by the tool, for analysis.

### 3 Literature review

Literature on bulk goods supply chains optimization models and, in particular, bulk handling terminals, appears to be relatively sparse. [Boland and Savelsbergh \(2010\)](#) present an overview of the HVCC, and discuss opportunities for the use of quantitative decision support tools to better plan and more effectively operate the coal chain. As the integrated system is very complex, optimization technology is identified as having the potential to be highly valuable to the HVCCC in achieving their goal to maximize the efficiency of the system. The opportunities identified for its use range from strategic infrastructure planning, to tactical stockyard management, to operational train scheduling, and to day-of-execution disruption handling.

[Singh et al. \(2012\)](#) develop a strategic infrastructure capacity planning model for the HVCC. It is a mixed integer programming model that simultaneously schedules coal chain operations and decides capacity expansions. The objective function trades off demurrage (or delay) costs and capacity expansion costs. The typical instance considers a year-long planning horizon with a daily time granularity, but with start times of activities, such as the stacking a train load at a terminal or the reclaiming of a stockpile, modeled at the minute level. Because solving such instances is computationally prohibitive, various heuristics, e.g., a genetic algorithm and a large neighborhood search algorithm, are developed to determine an initial feasible solution, which is then passed to a commercial solver (CPLEX) for possible improvement using the mixed integer programming model. An extensive computational study is conducted to validate the model, and the performance of different heuristics is compared, in terms of their computation time and objective function value. The solutions produced by the model are compared to historical data in terms of, among others, the daily vessel queue lengths, the number of train arrivals at a terminal per day, and the tonnage of coal loaded onto vessel per day.

The work in [Singh et al. \(2012\)](#), while an important precursor to that we present here, has some significant differences to our work. The key difference in

purpose – we seek to rapidly assess the coal chain performance, rather than find an optimal capacity expansion option – allows us to develop a model with several significant advantages over that in Singh et al. (2012). By decoupling capacity expansion decisions from the coal chain performance assessment model, we are able to obtain a more tractable model, and to invest more effort on investigating alternative models (e.g., models using different sets of decision variables), preprocessing techniques to reduce the number of variables, and valid inequalities that tighten the formulation and produce stronger linear programming bounds. It also allows us to include some real-life features not modeled by Singh et al. (2012). For example, Singh et al. (2012) enforce a particular vessel sequence at each the berth, which is determined *a priori* by a simulation, whereas we model the operational flexibilities permitted in the system, to alter the service sequence so as to reduce overall vessel delays. We reserve the entire space that a stockpile occupies on a pad as soon as the first train with coal for that stockpile arrives (which is what happens in reality), whereas Singh et al. (2012) reserve space on the pad per train basis. We explicitly take into account that the reclaiming of some stockpiles takes more than 24 hours. In such situations, our model divides the reclaim time over two consecutive days. This situation is not explicitly accounted for by Singh et al. (2012), and, as a result, it is possible that such stockpiles are reclaimed on a single day using two reclaimers, which would not happen in real-life.

In other work on the HVCC system, ? describe a method for scheduling trains on the Hunter Valley Coal Chain rail network. The operators of the rail network provide a number of paths, with specified arrival and departure times, that can be used for coal movement. A Lagrangian heuristic is presented, which is able to produce high quality solutions in a reasonable amount of time. Boland et al. (2012), ?, and ? develop approaches for stockyard management at the largest coal export terminal in the Port of Newcastle, focusing on stockpile placement and stockpile stack and reclaim timing decisions. The approaches differ in terms of the level of detail of the modeling and/or the methodologies used (from integer programming to purpose-built construction and local search heuristics to constraint programming).

Other than the research on the HVCC discussed above, Conradie et al. (2008) consider the end-to-end optimization of a coal chain from a mine to a factory, where the factory is producing liquid fuel products. The process includes transporting coal from mine bunkers to coal stacking locations where it is stacked into stockpiles, and then reclaiming coal from the stockpiles and conveying it to the coal processing facility. Conflicting objectives, such as minimizing excess coal extracted at the mines (due to the limited storage at the mines), maximizing stacked coal in the yard (to minimize the risk of interruptions in production), and minimizing the ashes and fines stacked on the yard, (to ensure high quality blends), need to be accommodated, and uncertainty in demand needs to be captured. In order to solve the resulting multi-objective, stochastic scheduling problem in a reasonable time, a simulated annealing approach is applied.

Hanoun et al. (2013) study the scheduling of a continuous coal handling problem with multiple conflicting objectives. An effective heuristic is described for planning stockpiles and scheduling resources to minimize the coal age in the stockyard and the delays in production. A model of stockyard operations within a coal mine is described and the problem is formulated as a biobjective optimization problem. Together, the model and the proposed heuristic act as a decision support system

for the stockyard planner to explore the effects of alternative decisions, such as balancing age and volume of stockpiles, and minimizing conflicts due to stacker and reclaimer movements.

The papers discussed above address operational scheduling problems, and well illustrate that such problems are very challenging in their own right. In the model presented here, we seek to avoid operational details as far as possible, and especially avoid the details associated with scheduling the precise movements of stockyard handling equipment. Instead, to best address the goals of strategic planning, we approximate stockyard capacity, with the exigencies of terminal handling equipment modeled via realistic choices for the specified operating rates of stacking and reclaiming machines, and for the specified stockyard capacity. However, for strategic planning in the HVCC, it is essential to model train scheduling and berth scheduling in some detail, otherwise crucial interactions will not be captured and the results will not be meaningful.

Although a large number of supply chain management papers have been written, few of them consider scheduling subproblems in detail. Two examples that do so, and that also point to the benefits of coordination between subproblems, are the papers of [Hall and Potts \(2003\)](#) and [Chen and Hall \(2007\)](#). Both consider a manufacturing supply chain, in which a manufacturer has to wait until all parts have been produced and delivered by the suppliers, so the effectiveness of the manufacturer depends on the scheduling decisions made by the suppliers. In the latter paper, the focus is on the manufacturer's ability to satisfy customer demand in a timely manner, and on practical mechanisms for obtaining cooperation between suppliers and manufacturers. [Hall and Potts \(2003\)](#) focus on cost minimization, and present algorithms that schedule the work of the supplier and the manufacturer separately, (where the latter algorithms take the supplier delivery schedule as input), and algorithms that schedule the work of the supplier and the manufacturer jointly. Computational experiments show that coordination causes a substantial reduction in the total cost.

Indeed, the HVCC operates as a coordinated system: the role of the HVCCC is to make coordinated decisions. Thus our model optimizes schedules for the whole system, in an integrated way. It provides a plan that is jointly optimized over the suppliers (load points), the transportation (trains and railroads), and the manufacturers (the terminals).

Since the HVCC can be viewed as a make-to-order assembly system, as the arrival of a vessel triggers the assembly of one or more stockpiles of particular coal blends, it is relevant to look at the literature of make-to-order production planning. One paper that addresses this area, and that also develops an integer programming model for such a system, is that of [Kolisch \(2000\)](#), which studies a make-to-order production planning problem with assembly and fabrication. Precedence constraints between jobs may exist and each job requires several parts, which are either fabricated locally or bought from suppliers. The goal is to provide a resource and time-feasible plan to satisfy customer orders before their due date. A mixed-integer programming model is presented, which includes scheduling based assembly, lot-sizing based fabrication, and coordination of assembly and fabrication. Although the supply chain modeled has some features in common with that we study here, the objective is very different, as is the solution methodology employed to solve the model. [Kolisch \(2000\)](#) seek to minimise the inventory holding

and setup costs, rather than delivery delays, and use a two-level heuristic to solve the model, rather than improving the IP model or solution approach.

[Bodon et al. \(2011\)](#) model a mineral export supply chain using a combination of optimization and discrete event simulation techniques to analyze capacity and evaluate expansion options. The discrete event simulation model is run over a one year time period, with the optimization model used within it to plan material movements on a more frequent basis, such as fortnightly, weekly, or a number of days in advance. This approach is broadly similar to the current simulation model used at the HVCCC, which also embeds an optimization procedure for short-term re-scheduling in response to simulated disruptions.

A key question in the research we present in this paper is: on what basis should different infrastructure options be compared? In supply chains in general, one might wish to compare how different supply chain infrastructures (or, indeed, operational practices) perform, and consider on what basis to make such a comparison. In the context of manufacturing, [Beamon \(1999\)](#) develops a framework for supply chain performance assessment. It is observed that using only one performance measure, such as cost, is attractive due to its simplicity, but it does not necessarily explain the system performance. It is suggested that three aspects of a supply chain need to be assessed: resources, output, and flexibility. These three aspects address the following three questions, respectively. (1) Are resources used efficiently, for example, what is the minimum level of resource, or what is the resource utilization, required to achieve the supply chain objectives? (2) How well does the supply chain serve its customers: what customer service levels are achieved? (3) How responsive is the supply chain to change? In the HVCC, the primary supply chain objectives are to meet the demand, represented by shipping stems, on time, in the sense that all vessels are served within a reasonable time of their arrival at the port, and to minimize delays in serving vessels. Thus the model we develop constrains service of all vessels to be completed within a fixed time after their arrival at port, given as a parameter of the model, that can be adjusted, and minimizes the total vessel service delay. In the analysis of the HVCC we carry out with this model, we address all three aspects of supply chain assessment suggested by [Beamon \(1999\)](#). We analyze the output of the model to determine the utilization of the supply chain resources, (such as berths, handling equipment, rail track and trains), and the delays in servicing vessels (reflecting the level of customer service provided). Our study measures the performance of the supply chain in response to shipping stems that have been manipulated in specific ways, for example, by keeping the demand quantities constant while increasing the variability of vessel interarrival times. This provides insights into the flexibility of the supply chain in response to different potential stressors on the system.

#### 4 HVCC Performance Assessment Models

In order to assess the performance of the Hunter Valley coal chain for a given infrastructure setting, we must model the schedule of all operations for critical pieces of infrastructure. Specifically, we need to model the time for which each vessel occupies a berth, the time each cargo for each vessel is reclaimed and loaded, the time and space occupied by each stockpile in a terminal stockyard, the time each train load is received at a terminal and stacked on its stockpile, the time

each piece of rail line is occupied by a train, the time each load point is required to load a train, and the time each piece of rolling stock is required for transport. This presents several key challenges.

First is the issue of time granularity in the model. Stacking a train load takes approximately two hours, reclaiming a stockpile can take anywhere from a few hours to almost two days, and a train round trip between the port and a mine can take from a few hours to almost 36 hours. However, stockpiles occupy stockyard space for anything from a few to 10 days, and, a planning horizon of at least a few months is needed for accurate assessment.

Second is the issue of operational variability. Even without major disruptions, train and stockyard handling equipment operations are, in practice, highly dynamic. Detailed train, stacking and reclaiming schedules are notional, at best. For example, the status of the area of the stockyard occupied by a stockpile, after it is loaded, is not known until after the vessel has finished loading; so-called stockpile remnants may need to be handled. Also, frequent train rescheduling occurs on the day of operations, since train arrivals dictate stacker operations, and reclaimers and stackers must be jointly scheduled to avoid clashes, (their operations may interfere with each other, depending on where in the stockyard each machine is working or traveling through), complex rules and human judgment are employed to juggle the precise sequence of train arrivals, stacker and reclaimer activities, with many choices only finalized on the day of operations. Within this dynamic setting, however, live run planners have sufficient flexibility so that they can usually achieve daily planning targets in terms of total trains unloaded, and vessels loaded, that are stipulated in advance plans, which are made taking into account rail, train and terminal handling equipment capacities at the daily level.

In creating the model, and addressing these issues, we need to balance accuracy and efficiency, because the ultimate goal is to use the model to quickly identify scenarios that should be investigated more thoroughly using the simulation model. The balance between accuracy and efficiency is reflected in our choice of time granularity: all our decisions, and our models of stockyard handling equipment and space capacity, are at the daily level. However, we take extra care in accounting for berth capacity, as there is much less operational flexibility in its use than there is for the rail and stockyard components of the infrastructure.

The driving decision variables in the model arise from the following observation. Coal from a specific mine is transported to a terminal using a train of a specific size, in terms of the tonnage it can haul, which depends on the load point characteristics and the above-rail operator serving the load point. As a consequence, each coal component tonnage of a vessel's cargo can be converted into a number of train loads, each of which is referred to as a *train-job*: each vessel has a specified set of train-jobs, each of which dictates that a train of a particular type (implying its tonnage) must be brought from a particular load point to a particular terminal, in order to assemble the cargo for that vessel. To effectively operate the coal export chain, it is crucial to time the train-jobs associated with a vessel's cargoes properly. Not surprisingly therefore, this is one of the key decision variables in the model discussed in this paper.

We next introduce a natural performance assessment model, driven by deciding the timing of the train-jobs associated with the cargoes of vessels at a daily level. After that, we explore various avenues to improve the solvability of the model.

#### 4.1 A Train-job Based Model

As already mentioned, each vessel cargo, which corresponds to a stockpile, can be converted to a set of (an integer number of) train-jobs. Specifically, for each vessel,  $v \in V$ , specified in the shipping stem, and each stockpile,  $s \in S(v)$ , associated with a cargo of vessel  $v$ , there is a set of components,  $C(s)$ . Each component,  $c$ , has a specified number of tons,  $Tonnage(c)$ , and dictates a load point from which coal for the component is to be sourced, which, in turn, dictates a train type, and hence a tonnage, for each train,  $TrainSize(c)$ , used to supply the component. The set of train-jobs,  $J(c)$ , for each component,  $c$ , is thus indexed by  $\{1, \dots, |J(c)|\}$ , where  $|J(c)|$  is calculated as  $|J(c)| = \lceil \frac{Tonnage(c)}{TrainSize(c)} \rceil$ , where  $\lceil \cdot \rceil$  is the nearest integer function, and each train-job,  $j \in J(c)$ , carries  $m_j = TrainSize(c)$  tons. Thus, for each stockpile  $s$ , the total number of trainloads to be delivered must be the total number of train-jobs for all components,  $|J(s)| = \sum_{c \in C(s)} |J(c)|$ . We note that for our data, we have that, averaged over all cargoes of all vessels in the shipping stem,  $\sum_{c \in C} |J(c)| \times TrainSize(c) = 0.9997 \sum_{c \in C} Tonnage(c)$ , which suggests that this approximation does not introduce a systematic error.

Our choice of a daily time granularity in the model has a few consequences. The reclaim time of a stockpile must be converted from hours to days. Because reclaiming of a stockpile can commence only after the stockpile is completely assembled, and some buffer time, typically of the order of 12 hours, is needed between stacking of the last train-job and reclaiming, (to obtain the results of quality tests on the assembled cargoes and to finalize the vessel paperwork before berthing), we enforce that the reclaiming of a stockpile can only start on the day after the day on which the last train with coal for the stockpile arrived at the terminal.

Also, we do not distinguish dumping and stacking and we do not distinguish reclaiming and ship loading. When coal is unloaded at the dump station of a terminal it is dumped on a conveyor belt and immediately transported to a stacker. In the model, we use their combined processing capacity on a day, which is calculated as the maximum amount of coal that can be dumped and stacked at a terminal, per day. Similarly, when coal is reclaimed by a reclaimer it is dumped on a conveyor belt and immediately transported to a ship loader. We use their combined processing capacity on a day, which is calculated as the maximum amount of coal that can be reclaimed and loaded at a terminal, per day. It takes some time to get a reclaimer ready to reclaim the next stockpile, about 30 minutes in practice. This is included in the reclaim time of each stockpile.

Table 1: Model parameters

$H$ :	set of time periods
$V$ :	set of vessels
$S$ :	set of stockpiles
$C$ :	set of components
$J$ :	set of train-jobs
$L$ :	set of load points
$U$ :	set of junctions
$E$ :	set of terminals
$W$ :	set of wagon types

---

$P$ :	set of pads
$P(e)$ :	set of pads at terminal $e \in E$
$V(e)$ :	set of vessels loaded at terminal $e \in E$
$S(v)$ :	set of stockpiles associated with vessel $v \in V$ ( $S(v) = \{s_1^v, s_2^v, \dots, s_{n(v)}^v\}$ )
$S(e)$ :	set of stockpiles assembled at terminal $e \in E$
$C(s)$ :	set of components associated with stockpile $s \in S$
$J(e)$ :	set of train-jobs delivering coal at terminal $e \in E$
$J(v)$ :	set of train-jobs delivering coal for vessel $v \in V$
$J(s)$ :	set of train-jobs delivering coal for stockpile $s \in S$ .
$J(c)$ :	set of train-jobs associated with component $c \in C$
$J(l)$ :	set of train-jobs picking up coal at load point $l \in L$
$J(u)$ :	set of train-jobs passing through junction $u \in U$
$J(w)$ :	set of train-jobs using wagon type $w \in W$

---

$b_v$ :	day of arrival of vessel $v$
$m_s$ :	tonnage of stockpile $s$
$l_s$ :	length of stockpile $s$ (meters)
$\tilde{d}_s$ :	number of hours required to reclaim stockpile $s$
$d_s$ :	number of days on which stockpile $s$ may be reclaimed ( $d_s = \lceil \tilde{d}_s / 24 \rceil$ )
$m_j$ :	tonnage of train-job $j$
$n_j^w$ :	number of wagons of type $w$ used in train-job $j$
$d_j$ :	cycle time of train-job $j$ (hours from the port to its load point and back to the port)
$T^e$ :	number of hours preparation required for dumping and stacking each train-job at terminal $e$
$S_1^e$ :	number of dumping and stacking hours available at terminal $e$ per day
$\rho_1^e$ :	minimum of dumping and stacking rates for terminal $e$ (tons per hour)
$S_2^e$ :	number of reclaiming and ship loading hours available at terminal $e$ per day
$\rho_2^e$ :	minimum of reclaiming and ship loading rates for terminal $e$ (tons per hour)
$e_v$ :	day vessel $v$ is due to depart
$L_p$ :	length of pad $p$ (meters)
$B_e$ :	number of berths at terminal $e$
$C_l$ :	loading capacity of load point $l$ (tons per day)
$C_u$ :	capacity of junction $u$ (number of trains per day)
$C_w$ :	number of wagons of type $w$
$\Delta^-$ :	number of available days before arrival of a vessel for railing and stacking of its stockpiles
$\Delta^+$ :	number of available days after arrival of a vessel for railing, stacking, and reclaiming of its stockpiles

---

The cycle time,  $d_j$ , of a train-job,  $j \in J$ , represents the time it takes a train to go from the port to the load point associated with the train-job and back. Since the decision variables are at a daily level, we provide feasible solutions (considering all capacity limitations) in which each train-job is done in only one day (within the midnight-to-midnight range). In reality, a train could depart at 6pm and arrive at 6pm the next day, for example, and so could span more than one day. Our assumption to have a train-job within a day is more restricted than the reality, and may result in solutions with higher delays. However, our numerical experiments

demonstrate that the raiing operation is not the bottleneck of the supply chain, which supports our use of this assumption.

As already mentioned, each train-job has an associated load point, which dictates the train type used. The train type specifies the number of wagons, where each wagon must be one of a certain type in the set of all wagon types,  $W$ . Thus for each  $c \in C$  and each  $j \in J(c)$ , we have  $n_j^w$ , defined to be the number of wagons of type  $w \in W$  used for train-job  $j$ . The cycle time,  $d_j$ , is used to calculate the number of wagon-hours that are required for train-job  $j$ , for wagons of each type. These are constrained to ensure that the fleet size limit,  $C_w$ , for each wagon type,  $w \in W$ , is respected.

Two model parameters,  $\Delta^-$  and  $\Delta^+$ , indicate the number of days, before and after the arrival date of each vessel, respectively, in which activities associated with the vessel can be scheduled. These parameters are primarily used to reflect operational practices, and to embody the supply chain objective to serve every vessel within a reasonable time. Their choice also has a significant effect on the size of the formulations.  $\Delta^-$  is restricted by the terminal managers, who do not want coal for a vessel to arrive too long before the arrival of the vessel, because it may unnecessarily take up space in the stockyard. (Coal age in a stockyard can also have safety implications.)  $\Delta^+$  should be at least the time needed to load a vessel, otherwise no feasible solution exists. Choosing  $\Delta^+$  too small takes away flexibility, as it limits the maximum delay that a vessel is allowed to incur.

The complete list of infrastructure and demand parameters used in the model can be found in Table 1.

There are three sets of decision variables in the model:

$$\begin{aligned} x_j^t &= \begin{cases} 1 & \text{if train-job } j \text{ is done on day } t, \\ 0 & \text{otherwise,} \end{cases} \\ y_{sp}^t &= \begin{cases} 1 & \text{if stacking of stockpile } s \text{ on pad } p \text{ starts on day } t, \\ 0 & \text{otherwise, and} \end{cases} \\ z_{sp}^t &= \begin{cases} 1 & \text{if reclaiming of stockpile } s \text{ from pad } p \text{ starts on day } t, \\ 0 & \text{otherwise,} \end{cases} \end{aligned}$$

with domains given by

$$\begin{aligned} t &\in \{b_v - \Delta^-, b_v - \Delta^- + 1, \dots, b_v + \Delta^+ - 1\}, \quad j \in J(v), \quad v \in V, \\ t &\in \{b_v - \Delta^-, b_v - \Delta^- + 1, \dots, b_v + \Delta^+ - 1\}, \quad p \in P, \quad s \in S(v), \quad v \in V, \text{ and} \\ t &\in \{b_v, b_v + 1, \dots, b_v + \Delta^+\}, \quad p \in P, \quad s \in S(v), \quad v \in V, \end{aligned}$$

respectively, where  $b_v$  is the arrival day of vessel  $v$ .

The following constraints ensure that the stockpiles for all vessels are assembled and loaded in the required sequence. For readability, time index sets are not given explicitly, but are assumed to be those given by the domains for the referenced variable.

- Each train-job must be completed, and so

$$\sum_t x_j^t = 1, \quad \forall j \in J. \quad (1)$$

- All stockpiles must be stacked on some pad:

$$\sum_{p \in P} \sum_t y_{sp}^t = 1, \quad \forall s \in S. \quad (2)$$

- A stockpile must be stacked and reclaimed on the same pad:

$$\sum_t y_{sp}^t = \sum_t z_{sp}^t, \quad \forall s \in S, \forall p \in P. \quad (3)$$

Constraints (2) and (3) also ensure that all stockpiles are reclaimed.

- All train-jobs for a stockpile have to arrive on or after the start day of the stacking:

$$\sum_{j \in J(s)} \sum_{t' \geq t} x_j^{t'} \geq |J(s)| \sum_{p \in P} y_{sp}^t, \quad \forall s \in S, \forall t. \quad (4)$$

- Reclaiming of the stockpiles of a vessel cannot start until all stockpiles have been completely assembled, and the buffer time elapsed, and thus cannot start until the day after all train-jobs for the vessel have been completed:

$$\sum_{j \in J(v)} \sum_{t' < t} x_j^{t'} \geq |J(v)| \sum_{p \in P} z_{s_p^v}^t, \quad \forall v \in V, \forall t. \quad (5)$$

- The stockpiles of each vessel  $v$ ,  $s_1^v, s_2^v, \dots, s_{n(v)}^v$ , have to be reclaimed in the specified order. We model this with the constraint

$$\sum_{p \in P} \sum_{t' \geq t} z_{s_p^v}^{t'} \leq \sum_{p \in P} \sum_{t' \geq t + \lfloor \sum_{k=i}^{j-1} \tilde{d}_{s_k^v} / 24 \rfloor} z_{s_p^v}^{t'}, \quad \forall i, j \in \{1, \dots, n(v)\}, j > i, \forall v \in V, \forall t. \quad (6)$$

This constraint also models the total duration of vessel loading, optimistically, as we illustrate with the following example. Consider a vessel  $v$ , with three stockpiles  $s_1^v$ ,  $s_2^v$ , and  $s_3^v$ , having associated reclaim times  $\tilde{d}_{s_1^v} = 12$ ,  $\tilde{d}_{s_2^v} = 8$ , and  $\tilde{d}_{s_3^v} = 9$  hours. Stockpiles  $s_1^v$  and  $s_2^v$  can be reclaimed on the same day, but it is not possible to reclaim all three stockpiles in one day, because the total reclaim time is 29 hours and exceeds a day. Therefore, the reclaiming of the third stockpile has to start at least  $\lfloor (12+8+9)/24 \rfloor$  days after the reclaiming of the first stockpile commences. In general, the reclaiming of stockpile  $s_j^v$  has to start at least  $\lfloor \sum_{k=i}^{j-1} \tilde{d}_{s_k^v} / 24 \rfloor$  days after the reclaiming of  $s_i^v$  commences, where  $j > i$ . Constraint (6) says that if the reclaiming of  $s_i^v$  starts at or after day  $t$ , then the reclaiming of  $s_j^v$  can only commence on or after day  $t + \lfloor \sum_{k=i}^{j-1} \tilde{d}_{s_k^v} / 24 \rfloor$ . Note that there is no need for this constraint to be a very accurate model of the time needed to reclaim all stockpiles, as this is captured in the reclaiming capacity and berth capacity constraints, described later.

- Load point capacity must be respected:

$$\sum_{j \in J(l)} m_j x_j^t \leq C_l, \quad \forall l \in L, \forall t. \quad (7)$$

- Junction capacity must be respected:

$$\sum_{j \in J(u)} x_j^t \leq C_u, \quad \forall u \in U, \forall t. \quad (8)$$

- The fleet size for the wagons of each type must be respected, i.e., the number of wagon-hours traveled by wagons of type  $w$  in each day cannot exceed the

number of wagon-hours that can be provided by the fleet in a day, which is given by the fleet size multiplied by the number of hours per day:

$$\sum_{j \in J(w)} \frac{d_j}{24} n_j^w x_j^t \leq C_w, \quad \forall w \in W, \forall t. \quad (9)$$

- Dumping and stacking capacity at each terminal, on each day, must be respected:

$$\sum_{j \in J(e)} (T^e + \frac{m_j}{\rho_1^e}) x_j^t \leq S_1^e, \quad \forall t, \quad \forall e \in E, \quad (10)$$

where  $T^e + \frac{m_j}{\rho_1^e}$  gives the number of hours required to unload and stack train-job  $j$  at terminal  $e$ .

- Pad capacity, for each pad, on each day, must be respected:

$$\sum_{s \in S} l_s (\sum_{t' \leq t} y_{sp}^{t'} - \sum_{t' \leq t-d_s} z_{sp}^{t'}) \leq L_p, \quad \forall p \in P, \forall t. \quad (11)$$

When the stacking of stockpile  $s$  on pad  $p$  starts on day  $t$ , then  $l_s$  of the length of the pad will be occupied until  $d_s$  days after the start of the reclaiming of the stockpile. For a pad  $p$  and a day  $t$ , the summation  $\sum_{t' \leq t} y_{sp}^{t'} - \sum_{t' \leq t-d_s} z_{sp}^{t'}$  indicates whether stockpile  $s$  is on the pad on that day or not. (The summation is zero if stacking has not started, one if the stockpile is on the pad, and zero if the stockpile has been on the pad, but has already been reclaimed.)

- Reclaiming and ship loading capacity at each terminal, on each day, must be respected:

$$\sum_{s \in S(e)} \sum_{p \in P(e)} \sum_{t'=t-d_s+1}^t (\frac{m_s}{d_s \rho_2^e}) z_{sp}^{t'} \leq S_2^e, \quad \forall t, \quad \forall e \in E. \quad (12)$$

Although similar to the dumping and stacking capacity constraints, the constraints are slightly more involved because the number of hours required to reclaim a stockpile ( $\frac{m_s}{\rho_2^e}$ ) can be as high as 48 hours. When more than 24 hours of reclaiming time are required, a fraction of the reclaiming capacity needs to be allocated to the first day of reclaiming and the remaining fraction of the required capacity needs to be allocated to the second day of reclaiming. We have chosen to allocate equal parts to both days ( $\frac{m_s}{d_s \rho_2^e} = \frac{m_s}{2 \rho_2^e}$  in this case). Note, too, that stockpile  $s$  consumes reclaiming and ship loading capacity on day  $t$  if its reclaiming starts in  $[t - d_s + 1, t]$ , and therefore the summation  $\sum_{t'=t-d_s+1}^t z_{sp}^{t'}$  occurs in constraints (12).

- Berth capacity, at each terminal, on each day, must be respected:

$$\sum_{v \in V(e)} \left( \sum_{p \in P} \sum_{t' \leq t} z_{s_1 p}^{t'} - \sum_{p \in P} \sum_{t' \leq t-d_{s_{n(v)}}} z_{s_{n(v)} p}^{t'} \right) \leq B_e, \quad \forall t, \quad \forall e \in E. \quad (13)$$

A vessel  $v$  occupies a berth from the day on which the reclaiming of its first stockpile,  $s_1^v$ , starts to the day on which the reclaiming of its last stockpile,  $s_{n(v)}^v$ , ends. For a terminal  $e$  and a day  $t$ , the summation  $\sum_{p \in P} \sum_{t' \leq t} z_{s_1 p}^{t'} - \sum_{p \in P} \sum_{t' \leq t-d_{s_{n(v)}}} z_{s_{n(v)} p}^{t'}$  indicates whether vessel  $v$  is at a berth on that day or not.

In modeling space occupancy (pad and berth capacity) restrictions at the daily granularity, we use the parameter,  $d_s$ , for the number of days needed to load stockpile  $s$ , which is taken to be the number of hours required to load the stockpile, rounded *up*. However, this approximation does not yield a systematically pessimistic model, since the loading may start at any time of the day. For some stockpiles, loading may, in reality, start early in the day, in which case the model over-estimates the real occupancy requirement of the stockpile, but in other cases, loading may start late in the day, in which case the model under-estimates the real occupancy requirement. On balance, the model provides a realistic approximation.

As discussed earlier, the performance of the HVCC is assessed in our model by its ability to meet demand, in a reasonably timely fashion, while minimizing total vessel delays. Although the HVCCC is also interested in understanding the maximum throughput that the system is capable of, demand arrives as vessel-loads, each of which requires resources in different parts of the system, at different times: the nature and timing of these demand arrivals affects the ability of the system to deliver. Thus our model is based on demand in the form of shipping stem data, which stipulates the nature and timing of demand arrivals, and seeks to meet all demand in the given stem within a reasonable time (specified by the parameter  $\Delta^+$ ), with minimum total trip delay. The trip delay for a particular vessel is the difference between the trip departure time and the trip due time if the trip departure time is greater than the trip due time (otherwise there is no delay). This is calculated as

$$\min \sum_{v \in V} \sum_{p \in P} \sum_{t + d_{s_n^v} > e_v} (t + d_{s_n^v} - e_v) z_{s_n^v p}^t, \quad (14)$$

where  $e_v$  is the due time, i.e., the day on which vessel  $v$  is due to be loaded. This due time is taken to be the earliest possible departure time under ideal circumstances, determined as follows. First, we assume that a berth is available for the vessel at its arrival time and that the stockpiles corresponding to its cargoes are ready to be reclaimed at its arrival time. Then the due time is calculated as the vessel's arrival time plus the sum of its stockpile reclaiming times.

An understanding of throughput can be obtained from a model with the above objective by presenting it with shipping stems that represent alternative realisations of different aggregate demand tonnage scenarios. For example, the model may be run with shipping stems representing 100 Mtpa, 120 Mtpa, 140 Mtpa, 160 Mtpa, and so on. One expects the model will produce small delays for smaller aggregate demand tonnage scenarios, and that vessel queue lengths will not grow uncontrollably; as aggregate demand tonnage increases, delays will increase, and the queue of vessels observed in the output will become unstable, indicating that the system is not capable of delivering the corresponding throughput.

#### 4.2 A Component-job Based Model

One of the challenges encountered when solving an instance of the train-job based model is the symmetry introduced by the train-jobs. The issue arises when the solution to the linear programming relaxation has  $\sum_{j \in J(c)} x_j^t = f$  with  $f$  fractional for some stockpile  $s \in S$ , some component  $c \in C(s)$ , and some day  $t$ . In this

case, because  $f$  is fractional, there is at least one train-job variable with a fractional value. This fractionality will be eliminated by branching. In other words, the variable is set to zero on one branch and set to one on the other branch. Unfortunately, in the former case, in the solution to the linear programming relaxation at the child nodes, the other train-job variables for the same component on the same day will simply be adjusted to essentially return the same solution (the values for the variables will have been permuted). This is because, for a given component  $c \in C$ , all the model parameters associated with  $x_j^t$  for each  $j \in J(c)$  are identical. For a given  $c \in C$ ,  $m_j$  is identical for all  $j \in J(c)$ , (recall it equals  $Trainsize(c)$ ). Since each component is associated with a specific load point, and a specific terminal where the component's vessel is to be served, there is a unique path in the rail network for every train-job associated with the component, so the cycle time,  $d_j$ , is identical for all  $j \in J(c)$ , and, furthermore, for all junctions  $u \in U$ , either  $J(c) \cap J(u) = \emptyset$  or  $J(c) \subseteq J(u)$ . Since each component dictates a specific train type,  $n_j^w$  is also identical for all  $j \in J(c)$ . Finally, all train-jobs,  $j \in J(c)$ , are also (only) in  $J(s)$  for the unique  $s \in S$  with  $c \in C(s)$ , and (only) in  $J(v)$  for the unique vessel  $v$  with  $s \in S(v)$ . Thus for a given component  $c$ , the coefficients of  $x_j^t$  and  $x_{j'}^t$ , for any  $j, j' \in J(c)$ , are, in any constraint of the model, identical; the variables  $\{x_j^t : j \in J(c)\}$  are symmetric. This type of symmetry can lead to very large search trees, but can be eliminated by reformulation: by using component based model where the variables represent the number of train-jobs of a component that are carried out in one day. This change has the additional benefit that the number of variables is reduced.

To convert the train-job based model to the component based model, we need to introduce a few new parameters, replace a decision variable, and modify a number of constraints. The new infrastructure and demand parameters can be found in Table 2.

Table 2: New model parameters

$C$ :	set of components
$C(e)$ :	set of components associated with terminal $e \in E$
$C(v)$ :	set of components associated with vessel $v \in V$
$C(s)$ :	set of components associated with stockpile $s \in S$
$C(l)$ :	set of components associated with loading point $l \in L$
$C(u)$ :	set of components associated with junction point $u \in U$
$C(w)$ :	set of components using wagon type $w \in W$
$m_c$ :	tonnage of coal for each train-job of component $c \in C$
$d_c$ :	cycle time for each train-job of component $c$
$n_c^w$ :	number of wagons of type $w$ required by each train-job of component $c$

We note that  $J(c) = J(s) \cap J(l)$  for the component  $c$  of stockpile  $s$  that is loaded at load point  $l$ .

The binary train-job variables,  $x_j^t$ , for  $j \in J(c)$ , are replaced by the general integer component-job variable,  $x_c^t$ , which represents the number of train-jobs of component  $c$  carried out on day  $t$ , with domain  $t \in \{b_v - \Delta^-, b_v - \Delta^- + 1, \dots, b_v + \Delta^+ - 1\}$  when  $c \in C(v)$  and  $v \in V$ .

The component-job model has the same constraints as does the train-job model, after modifications, as follows.

- Each component-job must be completed (replacing constraints (1)):

$$\sum_t x_c^t = |J(c)|, \quad \forall c \in C. \quad (15)$$

- All component-jobs of a stockpile must arrive at or after the start day of stacking (replacing constraints (4)):

$$\sum_{c \in C(s)} \sum_{t' \geq t} x_c^{t'} \geq |J(s)| \sum_{p \in P} y_{sp}^t, \quad \forall s \in S, \forall t. \quad (16)$$

- Reclaiming of the stockpiles of a vessel cannot start until the day after all stockpiles have been assembled, and thus until all train-jobs for the vessel have been completed (replacing constraints (5)):

$$\sum_{c \in C(v)} \sum_{t' < t} x_c^{t'} \geq |J(v)| \sum_{p \in P} z_{s_p^t}^t, \quad \forall v \in V, \forall t. \quad (17)$$

- Load point capacity must be respected (replacing constraints (7)):

$$\sum_{c \in C(l)} m_c x_c^t \leq C_l^t, \quad \forall l \in L, \forall t. \quad (18)$$

- Junction capacity must be respected (replacing constraints (8)):

$$\sum_{c \in C(u)} x_c^t \leq C_u^t, \quad \forall u \in U, \forall t. \quad (19)$$

- Fleet size for each wagon type must be respected (replacing constraints (9)):

$$\sum_{c \in C(w)} \frac{d_c}{24} n_c^w x_c^t \leq C_w, \quad \forall w \in W, \forall t. \quad (20)$$

- Dumping and stacking capacity at a terminal must be respected (replacing constraints (10)):

$$\sum_{c \in C(e)} (T_1^e + \frac{m_c}{\rho_1^e}) x_c^t \leq S_1^e, \quad \forall t, \forall e \in E. \quad (21)$$

Another source of symmetry is the stock pads. Since stacking and reclaiming capacity are modeled at an aggregate level (at the terminal level), the pad that a stockpile is assigned to is immaterial in terms of stacking and reclaiming capacity usage. The assignment of a stockpile to a pad only impacts the pad capacity usage, that is the length of the pad that is occupied by the stockpile. But that means that if the pads are of equal length, the assignments of stockpiles to pads can be interchanged without affecting the feasibility and quality of the solution. This symmetry can be avoided by also aggregating pad capacity at a terminal: the pad space can be assumed to consist of a single pad with a length equal to the sum of the lengths of the actual pads at the terminal. All our computational experiments are performed with aggregated pad capacity.

### 4.3 Strategies to Increase Efficiency

In this section, we discuss various strategies to solve instances of the component-job based model more efficiently. These strategies include pre-processing, to reduce the number of variables, and strengthening of some of the constraints in the model.

#### 4.3.1 Pre-processing

Because the reclaiming of the stockpiles of a vessel has to be done in a pre-specified order, we may be able to reduce the domain of some variables. Observe that because the total number of days required to reclaim stockpiles  $s_{i+1}^v, \dots, s_{n(v)}^v$  is  $\sum_{s=s_{i+1}^v}^{s_n^v} \tilde{d}_s/24$ , the reclaiming of  $s_i^v$  has to start at or before  $b_v + \Delta^+ - \lfloor (\sum_{s=s_{i+1}^v}^{s_n^v} \tilde{d}_s)/24 \rfloor$  and the stacking of stockpile  $s_i^v$ , and thus the transport of coal for stockpile  $s_i^v$ , has to be completed at or before  $b_v + \Delta^+ - \lfloor (\sum_{s=s_i^v}^{s_n^v} \tilde{d}_s)/24 \rfloor - 1$ . We update the domains of the variables accordingly.

Because it may not be possible to perform all the train-jobs of a stockpile in a single day, due to various capacity limits, (particularly if coal for the stockpile is required from highly constrained load points or areas of the rail network), the domain of certain start stacking variables may be reduced. To obtain a lower bound on the minimum number of days required to transport all the coal for a stockpile  $s$  of vessel  $v$ , we solve an auxiliary optimization problem in which the number of days required to transport the coal for stockpile  $s$  is minimized, subject to constraints (18) - (21), where only stockpiles of vessel  $v$  are considered in the constraints. We denote the minimum value of this optimization problem by  $\mu_s$ . Then the stacking of stockpile  $s$  has to be started at or before  $b_v + \Delta^+ - \mu_s - \lfloor \sum_{s \in S(v)} \tilde{d}_s/24 \rfloor - 1$ .

In our computational study, we will refer to the basic component-job based model as CM, and we refer to the component-job base model with pre-processing as CM-P.

#### 4.3.2 Disaggregated and Strengthened Constraints

Next, we provide some alternative constraints, which provide a stronger linear programming relaxation.

Constraint (17) forces all component-jobs of a vessel to be completed before the reclaiming of the first stockpile of the vessel can commence. A similar constraint can be formulated for each individual component:

$$\sum_{t' < t} x_c^{t'} \geq |J(c)| \sum_{p \in P} z_{s_1^v p}^t \quad \forall c \in C(s) \quad \forall s \in S(v) \quad \forall v \in V \quad \forall t. \quad (22)$$

This is a disaggregated form of constraint (17) and so may be stronger.

Constraint (16) forces all component-jobs of a stockpile to arrive at or after the start day of stacking, and can be disaggregated in a similar way:

$$\sum_{t' \geq t} x_c^{t'} \geq |J(c)| \sum_{p \in P} y_{sp}^t \quad \forall c \in C(s) \quad \forall s \in S(v) \quad \forall v \in V \quad \forall t. \quad (23)$$

In our computational study, we will refer to the component-job based model with disaggregated constraints (22) and (23) by CM-D.

Constraint (23) can be further strengthened by replacing  $y_{sp}^t$  with  $\sum_{t' \geq t} y_{sp}^{t'}$ . The strengthening is valid because when  $\sum_{t' \geq t} y_{sp}^{t'} = 1$ , the stacking of stockpile  $s$  starts at or after  $t$ , which means that the component-jobs  $c$  for stockpile  $s$  have to be carried out at or after day  $t$ . (Note, too, that  $\sum_{t' \geq t} y_{sp}^{t'} \leq 1$  because of (2).) This leads to

$$\sum_{t' \geq t} x_c^{t'} \geq |J(c)| \sum_{p \in P} \sum_{t' \geq t} y_{sp}^t \quad \forall c \in C(s) \forall s \in S(v) \forall v \in V \forall t. \quad (24)$$

The same idea can be applied to constraint (22) leading to

$$\sum_{t' < t} x_c^{t'} \geq |J(c)| \sum_{p \in P} \sum_{t' \leq t} z_{s_1 p}^{t'} \quad \forall c \in C(s) \forall s \in S(v) \forall v \in V \forall t. \quad (25)$$

In our computational study, we will refer to the component-job based model with strengthened disaggregated constraints by CM-DS. The component-job based model in which the original constraints are strengthened in the same way is referred to as CM-S.

Using constraint (15) and (2), we can complement constraint (24), i.e.,

$$\begin{aligned} |J(c)| - \sum_{t' < t} x_c^{t'} &\geq |J(c)| \left(1 - \sum_{p \in P} \sum_{t' < t} y_{sp}^t\right) \\ \Rightarrow \sum_{t' < t} x_c^{t'} &\leq |J(c)| \sum_{p \in P} \sum_{t' < t} y_{sp}^t. \end{aligned} \quad (26)$$

Similarly, we can complement constraint (25), i.e.,

$$\sum_{t' \geq t} x_c^{t'} \leq |J(c)| \sum_{p \in P} \sum_{t' > t} z_{s_1 p}^{t'} \quad \forall c \in C(s) \forall s \in S(v) \forall v \in V \forall t. \quad (27)$$

Computational experiments revealed that the techniques embedded in CPLEX, the integer programming solver used for our computational experiments, are more effective when provided with the complemented constraints and result in faster solution times.

In our computational study, we will refer to the component-job based model with complemented strengthened disaggregated constraints by CM-DSC.

## 5 A Computational Study

## 6 Final Remarks

## References

- Bayer, A., Rademacher, M., and Rutherford, A. (2009). Development and perspectives of the Australian coal supply chain and implications for the export market. *Zeitschrift für Energiewirtschaft*, 33:pp. 255–267.
- Beamon, B. M. (1999). Measuring supply chain performance. *International Journal of Operations and Production Management*, 33:pp. 275–292.

- Bodon, P., Fricke, C., Sandeman, T., and Stanford, C. (2011). Modeling the mining supply chain from mine to port: A combined optimization and simulation approach. *Journal of Mining Science*, 47(2):202–211.
- Boland, N., Gulczynski, D., and Savelsbergh, M. (2012). A stockyard planning problem. *EURO Journal on Transportation and Logistics*, pages pp. 1–40.
- Boland, N. L. and Savelsbergh, M. (2010). Optimization the hunter valley coal chain. *Managing Supply Distributions (ed: H Gurnani, A Mehrotra and S Ray) (Springer-Verlag: London)*.
- Chen, Z.-L. and Hall, N. G. (2007). Supply chain scheduling: Conflict and cooperation in assembly systems. *Operations Research*, 55(6):pp. 1072–1089.
- Conradie, D., Morison, L., and Joubert, J. (2008). Scheduling at coal handling facilities using simulated annealing. *Mathematical Methods of Operations Research*, 68:pp. 277–293.
- Hall, N. G. and Potts, C. N. (2003). Supply chain scheduling: Batching and delivery. *Operations Research*, 51(4):pp. 566–584.
- Hanoun, S., Khan, B., Johnstone, M., Nahavandi, S., and Creighton, D. (2013). An effective heuristic for stockyard planning and machinery scheduling at a coal handling facility. In *Industrial Informatics (INDIN), 2013 11th IEEE International Conference on*, pages 206–211.
- Kolisch, R. (2000). Integration of assembly and fabrication for make-to-order production. *International Journal of Production Economics*, 68(3):pp. 287 – 306.
- Singh, G., Sier, D., Ernst, A. T., Gavrilouk, O., Oyston, R., Giles, T., and Welgama, T. (2012). A mixed integer programming model for long term capacity expansion planning: A case study from the hunter valley coal chain. *European Journal of Operational Research*, 220:pp. 210 – 224.