19th International Congress on Modelling and Simulation, Perth, Australia, 12–16 December 2011
http://mssanz.org.au/modsim2011

# Improved stockyard management strategies for coal export terminals at Newcastle

**N. Boland** [a], **D. Gulczynski** [a], **M.P. Jackson** [a], **M.W.P. Savelsbergh** [b] **and M.K. Tam** [a]

[a] *School of Mathematical and Physical Sciences, University of Newcastle*
[b] *CSIRO Mathematics, Informatics, and Statistics, Australia*
*Email: natashia.boland@newcastle.edu.au*

**Abstract:** Coal is Australia's leading export valued at close to A$50 billion and representing more than 20% of Australia's commodity exports in 2008 (Australian Department of Foreign Affairs and Trade).

The Port of Newcastle is home to the world's largest coal export operation. In 2008, it achieved a throughput of around 92 million tonnes, or more than 10 per cent of the world's total trade in coal for that year. That throughput is expected to double in the next decade. Several coal terminals operate at the Port of Newcastle. Crucial to achieving a high coal throughput is effective management of the stockyards at these terminals. The stockyard, where cargoes of (typically blended) coal product are assembled in stockpiles using stacking machines, and then reclaimed using bucket wheel reclaimers, is a pivotal component of a coal export chain.

In this paper, a model of stockyard operations within a coal export supply chain, with make-to-order cargo assembly, is described. Stockpiles are assembled from coal delivered by trains from load points at mines. For a given stockpile, it is known in advance how much coal must be delivered from each load point to completely assemble the stockpile. Trains arrive at a dump station and dump the coal onto a conveyor which transports it to a stacker which in turn assembles the stockpile on the yard. Once the stockpile is completely assembled (a process that usually takes several days), it is removed from the stockyard via a reclaimer and loaded onto a vessel (a process that usually takes around half a day to complete). A stockpile often remains in the stockyard for some time (several days) before its intended vessel arrives at the berth and it can be loaded. The system is constrained in a number of ways. There are limited berths available for vessels to be loaded. The departure of larger vessels may be restricted to high tide. Load point capacity at the mines, in terms of the number of trains as well as the volume that can be handled per day, is limited. Stockyard space is often at a premium, and dumping, stacking, and reclaiming capacity per day is limited too. All these constraining factors need to be taken into account when managing the sequencing and loading of vessels, and the management of the stockyard.

The model developed represents decisions and constraints typically applied at a planning stage of about 4 to 6 weeks in advance. The key decisions are where on the stockyard to place each stockpile for a vessel, when to start building the stockpile, when to bring the vessel to berth, and when to start reclaiming and loading each stockpile for the vessel. An approximate railing plan for transporting coal from mine load points to the stockyard is also required, largely as a check on load point and rail capacity limits. The model considers the stockyard itself and the outbound handling in some detail, with timing at the hourly level, but approximates in-bound capacity constraints more coarsely, at the daily level. We give a solution approach, simulating stockpile placement and scheduling decisions in a greedy fashion, with the goal of minimizing vessel sailing delays, and maximizing throughput of the system.

Since a greedy approach is unlikely to yield the most efficient schedules, a variant of the algorithm is developed in which the vessel scheduling order is randomized, and the resulting performance analysed computationally. The resulting model and algorithm together can be viewed as a prototype decision support system for the stockyard planner. However, it can also be used as a simulation tool, to explore the effects of alternative stockyard management strategies. Such strategies, may, for example, reserve some areas of the stockyard for fast-moving, and some for slow-moving, cargo, with the aim of balancing load across the stacking and reclaiming equipment. One such strategy is implemented in the model, and compared with the simple greedy approach. Computational results of this study are report. Our research shows that coal throughput can be substantially impacted by the stockyard management strategy employed.

***Keywords:*** *Coal export, supply chain management, logistics optimisation, stockyard management, decision support system*

## 1. INTRODUCTION

Coal is Australia's leading export valued at close to A$50 billion and representing more than 20% of Australia's commodity exports in 2008 (Australian Department of Foreign Affairs and Trade).

The Port of Newcastle is home to the world's largest coal export operation. In 2008, it achieved a throughput of around 92 million tonnes, or more than 10 per cent of the world's total trade in coal for that year. That throughput is expected to double in the next decade. Several coal terminals operate at the Port of Newcastle. Crucial to achieving a high coal throughput is greater efficiency in planning and operations. A key step in ensuring effective planning processes for the Hunter Valley Coal Chain (HVCCC) was the formation of the Hunter Valley Coal Chain Coordinator P/L (HVCCC). Founded in 2003 as a trial between Port Waratah Coal Services (PWCS) and rail operator Pacific National to initiate centralised planning activities, it quickly expanded to include other service providers via a memorandum of understanding signed in 2005. As a cooperative organisation responsible for planning all coal exports for the HVCC, it was a landmark for collaborative logistics. Its success led to incorporation in 2009 as the HVCCC, with a mission including

- planning and scheduling the movement of coal through the HVCC in accordance with the agreed collective needs and contractual obligations of coal producers and service providers; and
- minimising total logistics cost and maximising throughput via the provision of appropriate analysis and advice on capacity constraints affecting the efficient operation; and

A key component of the HVCCC's operational planning activities is vessel scheduling and stockyard planning. As information about coal vessel arrivals at the port is obtained, all aspects of the operation required to service the vessels must be planned. The first stage in this planning process typically occurs 4 to 6 weeks before the vessel arrival time. At this time, a vessel nominates its estimated arrival time (ETA) at port and gives its cargo specifications. A vessel typically carries between one and six cargoes (the majority of vessels are single-cargo, and very few carry more than two), each of which blends coal from one or more mines in specified proportions. A significant part of the HVCC operates as make-to-order cargo assembly, in which stockpiles are built for specific cargoes, usually between 3-10 days prior to loading onto the vessel. In such a cargo assembly system, the stockyard contains numerous stockpiles, each with its own distinct identity, in various stages of being built or destroyed (reclaimed). In such a system, effective management of the stockyard is clearly critical to the efficiency of the operation. Thus in the first stage of planning, in which the time at which the vessel should berth for loading is decided, the stockyard resources required to build, store and reclaim its cargo must also be considered; vessel scheduling and stockyard management are tightly integrated decisions.

The HVCCC planners responsible for this planning activity have the assistance of an IT system (the commissioning of which was one of its major early successes) providing a common view of the data across the HVCC's logistics system, graphical user interfaces, and some limited forms of decision support (for example, visual flagging of over-capacity points in the schedule). However the key choices of where in the stockyard to place the stockpiles for a vessel, when to start building them, and when to start loading the vessel, are in the hands of the human planner; it is manual planning process. The research reported in this paper is a first attempt to prototype a more automated decision support system. However it was quickly realized that the model and solution method developed could serve well as a simulation of the system, in which the stockyard operations are modeled in some detail. In particular, it could be used to assess alternative stockyard management strategies.

This paper describes the planning process, the model developed, and a greedy solution approach. Randomization in the vessel scheduling sequence is introduced, and the resulting algorithm performance compared with the original pure greedy method. An alternative strategy which reserves some areas of the stockyard for fast-moving, and some for slow-moving, cargo, with the aim of balancing load across the stacking and reclaiming equipment is also tested, and compared on historical data with the greedy approach. A summary of numerical results is presented. We observe that the scheduling method can have a big impact on the system performance in terms of vessel delays and throughput.

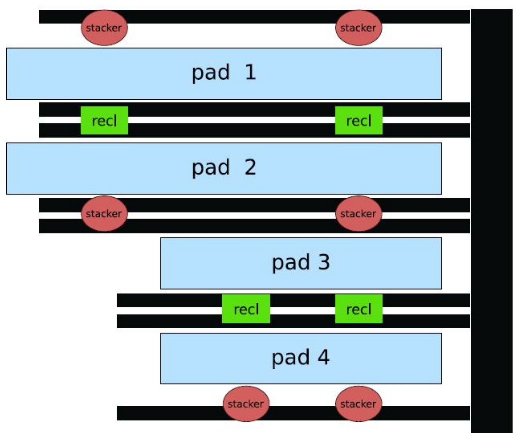


**Figure 1.** Sample stockyard with four pads, six stackers and four reclaimers.

## 2. LITERATURE REVIEW

A recent paper, (Conradie et al., 2008), considers end-to-end optimization of a coal chain from mine-to-factory, rather than mine-to-port, where the factory is producing liquid fuel products. The models of rail and stockyard capacity used in this paper are very simple, as the focus is on the robustness of schedules to uncertainty in factory demand. At the mining end of supply chain, there has been both a long history and intense recent activity in the development of decision support, and in particular optimization tools for mine scheduling (Newman et al, 2010). At the port end of the supply chain, the planning and scheduling of container ports has received increasing interest in the international literature. For example, the review of (Steenken et al., 2004) highlights the complexity of such operations, and describes it as "unthinkable" that such activities could be carried out efficiently without the use of appropriate mathematical optimization methods. Recently, interesting work (Wong and Kozan, 2010), explores the relationships between location of containers in the yard with terminal handling equipment schedules, developing promising heuristic and meta-heuristic approaches.

By comparison with mine extraction scheduling and container port operations, the research for bulk goods supply chains and in particular bulk handling terminals is relatively sparse. Aside from the paper of (Conradie et al., 2008), we can find no optimization-based work on bulk goods supply chains. There are a number of simulation studies of such systems, for example, (Glassrock and Hoare,1995) describe a discrete-event simulation system for an iron-ore supply chain. More recently (Welgama and Oyston, 2003) develop a simulator for the HVCC, spanning the system from mine load points to ship-loaders and berths. These studies are primarily used for answering strategic planning questions, such as the impact of new infrastructure. In the context of precast concrete products, (Marasani et al., 2001) develop a simulation model embedding artificial intelligence elements to help design the stockyard layout and manage the stockyard space. Unlike the cargo assembly stockyard studied here, the concrete products considered in (Marasani et al., 2001) are made-to-stock, rather than made-to-order. Analytic techniques such as queuing theory are used in (Binkowski and McCarragher, 1999) and (Ayu and Cardew-Hall, 2002) to answer questions such as the optimal number and size of stockpiles to maintain in a yard, and the train arrival rate, so as to minimize vessel delays, in a blended mineral resource context. This work is motivated by iron ore operations, again for make-to-stock systems, and the analysis requires somewhat idealized assumptions. (Abdekhodaee et al., 2005) does address a similar system to that considered in this paper. In fact, the HVCC itself is modeled. However, the main focus is the rail system, which is modeled in more detail than the stockyard. (Abdekhodaee et al., 2005) decompose their overall problem into several subproblems, each solved greedily. We can find no other published work algorithms or decision support for operational planning of bulk goods stockyards, or on the efficient scheduling of operations for stacker, reclaimer or ship-loading machinery in a bulk handling terminal.

In seeking to assign stockpiles to locations in the stockyard, the problem presented by the HVCC cargo assembly terminal clearly has some elements in common with 2D packing. Each stockpile consumes the entire width of a pad, occupying a length of the pad dependent on the stockpile volume. Thus for a single pad in the stockyard, its length constitutes one dimension into which the stockpiles must be packed. However two stockpiles cannot occupy the same space at the same time, so their dimension in time must be considered, as seen, for example, in Figure 2. There is a substantial literature on 2D packing in general (Lodi et al., 2002), for which it seems that heuristics provide the most practical solution approach, e.g. (Hopper and Turton, 2001). However the cargo assembly stockyard situation is substantially more complex than any of the 2D packing problems considered in the literature: the time dimension of the stockpile rectangles is unknown, and depends on (i) the train schedule for delivery of coal to the stockpile, (ii) the time at which all stockpiles for the same vessel are completed, and (iii) the reclaimer schedule for loading the vessel; the pad itself must be selected, introducing a third (discrete) dimension; and the objective is not simply to minimize waste in the encompassing rectangle, but to minimize lateness of vessels, which depends on the completion times of their associated stockpiles.

One paper which does consider a packing problem with complex constraints having some similarity to those encountered in the HVCC cargo assembly terminal is that of (Bay et al., 2010). This work uses a guided local search heuristic to schedule and locate shipbuilding activities needed for the manufacture of large ships. Ship "blocks" are assembled in a dedicated hall, called an assembly hall. The hall is divided into equal-sized rectangular areas, with blocks to be positioned in each area for the duration of their assembly. Thus we see that the hall resembles the HVCC stockyard, an area resembles a pad, and a block has the role of a stockpile in the HVCC system. There are clearly some key differences. Blocks require 2D packing in the assembly hall areas, leading to a 3D packing problem. However the time dimension of blocks is fixed, and known, as opposed to dependent on other constrained scheduling decisions. Despite these differences, the work of (Bay et al., 2010) may provide helpful insights.

## 3. A VESSEL SCHEDULING AND STOCKYARD PLANNING MODEL

### 3.1. Overview of the planning process

At about 4 to 6 weeks prior to ETA, a vessel communicates its ETA and cargo specification to the HVCCC. The planner then looks for stockyard space in which to place the cargo. At some terminals, a single cargo can be split into one or more stockpiles. Here only the single stockpile per cargo case is considered, as it is the dominant mode of operation in the HVCC.

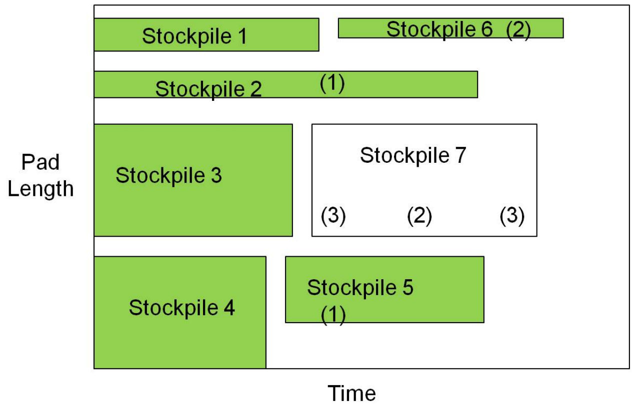


**Figure 2.** A stockyard pad over time, showing the positions and durations of planned stockpiles. Numbers in brackets indicate daily train requirements from a load point.

For each cargo on the vessel, the planner seeks a location in the stockyard in which to place the corresponding stockpile. The stockyard consists of a several pads, as shown in Figure 1. Each pad is served by one or more stackers, for loading coal from incoming trains onto the stockpiles, and by one or more reclaimers, for loading coal from the stockpile to the ship-loaders and hence onto the ship. (Note that both stackers and reclaimers can service pads on either side of them, and can move the length of a pad, provided they do not get too close to any neighbouring machinery.) Stockpiles occupy the width of a pad, with length determined by their volume. Obviously any new stockpile must be placed so that it does not overlap with (and leaves sufficient buffer between) other stockpiles that are on the ground in the pad at the same time.

Thus the planner must also determine the stockpile start time, and its duration. The latter is the length of time for which the stockpile remains on the ground in the yard. It depends on the schedule of trains delivering coal to it. Since a vessel typically does not start loading until after all stockpiles for it have been fully built (plus some buffer time), a stockpile's duration also depends on the start time and duration of any other stockpile for the same vessel. Of course the reclaimer schedule must also be considered, since even if all stockpiles are built, if the reclaimer is busy, it cannot start loading stockpiles for the vessel.

The planner typically "sketches in" a rail delivery plan for each stockpile as it is placed, indicating the number of trains that will be delivered to that stockpile from each load point on each day, and the stacker that will unload them onto the stockpile. In doing so, it is ensured that capacity limits on the load points are respected, and that the total load on each stacker, in terms of tons per day stacked, is below a specified limit. This is illustrated in Figure 2, which shows a new stockpile, number 7, provisionally scheduled. The number of trainloads by day from a load point having a maximum daily capacity of 5 is indicated against each stockpile sourcing coal from that load point. On each day, that maximum is respected. From the train plan, the planner can deduce the earliest time that the vessel can start loading (the earliest time after all stockpiles are completed, plus some buffer, that is after the vessel's ETA). After checking the berth schedule, and the ship-loader commitments to currently planned vessels, to determine when both are free for the vessel to use, the planner will determine its time to start loading, and advise the vessel accordingly.

The planner does not currently explicitly create a schedule for reclaimers, but will use estimated load times based on standard load rates for the equipment, together with the load rates specified for the vessel, to estimate when ship-loaders and berths will become free.

In deciding placement of stockpiles and allocating train/load points capacity to stockpiles, the planner will use experience and rules of thumb to ensure smooth operation of the system. However, operational variability means that all plans are revised on a daily basis, and the HVCCC reports that a great deal of planner time is spent in "re-work" with little time to look ahead or consider the longer term implications of their scheduling decisions. Thus an automated tool to assist planners in this activity is believed to be of great value.

### 3.2. The planning model

The model developed largely imitates the planning process described above, with the key exception that reclaimer schedules are modeled in detail.

A list of vessels V, ordered by ETA, is taken as input, together with a set of stockpiles S(v) to be built for vessel each $v \in V$. Each stockpile $s \in S(v)$ requires $num_{ls}$ trainloads of coal from load point l for each l in the set L(s), and uses up $len_s$ metres of pad space. Stockpiles take several days to build, so the inbound capacities (railing and stacking) are modeled at the daily level. For each stockpile s, a build start date, $t_s$, and end date, $u_s$, with $t_s < u_s$ must be decided, as well as $n_{lst}$, the number of trainloads from load point $l \in L(s)$ to convey to stockpile s on day t, for each $t=t_s,\ldots,u_s$. Obviously the sum of the $n_{lst}$ over t must give the required number, $num_{ls}$. Also the sum of the $n_{lst}$ over s cannot exceed the limit for load point l on day t. Each stockpile is also assigned a pad, $p_s$, and a location on the pad, $h_s$; the stockpile occupies the interval $[h_s, h_s+len_s]$ of the pad.

Loading a vessel can range from 12-45 hours, so the outbound capacities are modeled more precisely, at the hourly level. Otherwise there is a great risk that berth and ship-loading capacities are not accurately represented. The time a vessel v starts loading must be decided: it is recorded to the nearest hour, and must be at least $u_s$ for all $s \in S(v)$, plus a buffer. In fact, the stockpiles on a vessel are typically loaded in a given sequence, designed to avoid excessive stress on the vessel. Reclaiming a stockpile is carried out without interruption, thus the reclaim time $rec_s$ is fixed and known. A reclaiming start time $r_s$ must be set for each so that $r_s + rec_s < r_{\hat{s}}$, (with a margin), where $s, \hat{s} \in S(v)$ and $\hat{s}$ is the next stockpile in the vessel sequence after s.

Clearly for all pairs of stockpiles s and $\hat{s}$, if they are assigned to the same pad, and overlap in time, i.e. $p_s = p_{\hat{s}}$, and $[t_s, r_s] \cap [t_{\hat{s}}, r_{\hat{s}}] \neq \emptyset$, then their spatial extent cannot overlap, i.e. either $h_s > h_{\hat{s}} + len_{\hat{s}}$ or $h_{\hat{s}} > h_s + len_s$, where the inequality must be over-satisfied by some buffer value. There are a number of equipment constraints that must be applied. Stacker capacities are approximated by partitioning the stackers into sets that all service the same set of pads (e.g. in Figure 1, the sets are the two stackers along the top serving pad 1, the two stackers between pads 2 and 3 serving those pads, and the two stackers at the bottom serving pad 4). For each corresponding set of pads, the total coal to be stacked on those pads on each day, derived from the sum over l of the $n_{lst}$ values for s with $p_s$ in the set of pads, cannot exceed the total stacking capacity of the set of stackers. Each stockpile is assigned a reclaimer, and its reclaim time must be set, so that given its location on the pad, the reclaimer can reach it from its previous task in time, without getting too close (or passing!) its neighbouring reclaimer. Furthermore, the number of reclaimers permitted to operate at any one time cannot exceed the number of ship-loaders. Finally, a vessel must be at berth some buffer time prior to the start reclaiming time of its first stockpile, but if it is a large ship, its departure time from the berth may be later than the end time of reclaiming its last stockpile, as it may have to wait for the next high tide to leave. There is a channel limit on the number of large vessels that can depart on one high tide. Thus vessels may be held at one of a limited number of berths. So vessel reclaiming time assignment must also take this into account.

## 4. SOLUTION ALGORITHMS

### 4.1. A Greedy Approach

Vessels are considered in order of ETA. For each vessel v in order, and for each of its stockpiles $s \in S(v)$, in order, a stockyard location (pad and position) and day is sought so that (i) on that day, there is some spare capacity to bring a trainload of coal for s to the pad, and (ii) the pad length and currently planned stockpiles can accommodate the length of the stockpile in that location on that day. The method does not attempt to "fill holes", and only considers locations that are unrestricted by currently planned stockpiles for all future times. Of all such locations the one with the earliest day is selected. Ties are broken by imposing a priority ordering on stockyard locations, e.g. pads are ordered from bottom to top and pad positions from right to left, and selecting the first location in the order. Once the location and start date of the stockpile is fixed, trainloads are brought for it as quickly as is possible given the residual load point and stacker capacity. Once all stockpiles have been assigned a pad location and train schedule, the earliest time at which the vessel can berth and reclaiming can start is known. Then each stockpile on the vessel is, in order, assigned a reclaimer, at the earliest time at which both a reclaimer that can access the stockpile's location, and a ship-loader, are available. Once the reclaimer schedule for a vessel is decided, its departure time can be calculated, as well as the delay incurred. Delay of a vessel is taken to be the difference between its scheduled departure time, and the time it would have departed if it had been served with no other vessels competing for resources.

### 4.2. A Randomized Greedy Algorithm

It is not difficult to construct examples in which a purely greedy method is short-sighted, and causes unnecessary delays. However the complexity of the planning process and interaction of the constraints makes anything other than a sequential approach to allocating resources to a vessel challenging. Thus changing the vessel sequence is considered. At any point in the process, the next k vessels in ETA order are considered, and one of them chosen at random to schedule next. The method can be run multiple times, and the best result selected. Experiments showed that k=3 and 50 runs provided a good balance between run time and solution quality.

### 4.3. A Fast-Slow Cargo Build Time Partition Method

Greedy methods are by their nature myopic, and can't foresee when or where bottlenecks in capacity might occur. However bottlenecks in stacker and reclaimer capacity may be mitigated if the load on them could be more evenly spread. Since their load depends on the location of the stockpiles in the yard, one approach is to partition the yard space, reserving some areas for cargos that impose a high load, and some a low load, on this equipment. In particular, some load points are more constrained than others, meaning that coal from these load points arrives more slowly, leading to a longer build time for the cargo which needs it. Cargo can be classified accordingly as slow, medium and fast build. The latter will impose greater load on the stackers than the former. Thus one proposal is to partition the yard in a "chequer-board" fashion, with each pad divided in two halves, one designated for fast and one for slow cargo, in an alternating pattern. (Medium cargo is still free to be assigned in any location.) Such a strategy can easily be implemented within the greedy approach, simply by altering the priority ordering on stockyard locations, and including cargo-dependence.

## 5. COMPUTATIONAL STUDY

### 5.1. Data sets

The basis of our computational study is infrastructure and shipping demand data provided by the HVCCC, reflecting the system settings and activity in 2006, for a single cargo assembly terminal. The historical vessel data had about 6 months of vessel arrivals (339 vessels). About 59% of cargos were classified as medium build time, 16% fast and 25% slow. The stockyard was similar to that shown in Figure 1.

In order to understand the behavior of alternative algorithms and strategies, the vessel data was manipulated to create additional data sets. The vessel arrival stream is permuted in a "local" way. For each pair of arriving vessels in turn, their ETA is swapped with probability 0.5. The choice of pairs moves one vessel along the stream in each step, so with some probability a vessel could be swapped multiple times. This could move a vessel far from its original ETA, but with exponentially decreasing probability. Thus only local changes are favoured. This approach to generating multiple data sets ensures that the planned system load, and any regularity or patterns in it are largely preserved, and hence realistic. Importantly, it ensures conclusions drawn about algorithm performance are robust to the natural variability of vessel ETAs, which can vary by a few days depending on weather and other local conditions.

### 5.2. Computational Results

Table **1**. Results of numerical experiments, showing the average, maximum and standard deviation of delay time to the nearest hour over all vessels, for each algorithm. In the case of Randomized Greedy, these statistics are averaged over 50 runs, and are reported for the best schedule produced in terms of mean delay.

| Data set | Greedy | | | Fast-Slow Partition | | | Randomized Greedy | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | | Ave over all runs | | | Best over all runs | | |
| | Ave | Max | SD | Ave | Max | SD | Ave | Max | SD | Ave | Max | SD |
| History | 220 | 398 | 98 | 230 | 510 | 108 | 289 | 577 | 91 | 91 | 287 | 63 |
| 1 | 259 | 426 | 115 | 71 | 284 | 40 | 249 | 520 | 114 | 121 | 297 | 60 |
| 2 | 286 | 522 | 131 | 275 | 482 | 140 | 280 | 554 | 125 | 122 | 328 | 68 |
| 3 | 336 | 597 | 164 | 277 | 474 | 141 | 267 | 563 | 133 | 147 | 435 | 74 |
| 4 | 222 | 398 | 100 | 194 | 422 | 98 | 281 | 563 | 128 | 157 | 402 | 80 |
| 5 | 209 | 373 | 86 | 286 | 483 | 125 | 253 | 536 | 125 | 94 | 284 | 52 |
| 6 | 358 | 721 | 195 | 292 | 596 | 155 | 264 | 550 | 122 | 150 | 409 | 80 |
| 7 | 152 | 322 | 67 | 336 | 593 | 167 | 291 | 590 | 135 | 122 | 297 | 63 |
| 8 | 260 | 485 | 118 | 376 | 720 | 195 | 242 | 521 | 116 | 128 | 373 | 54 |
| 9 | 281 | 497 | 149 | 108 | 284 | 48 | 254 | 524 | 115 | 138 | 374 | 64 |
| Ave | 258 | 474 | 122 | 245 | 485 | 122 | 267 | 550 | 124 | 127 | 348 | 66 |

Table 1 shows the results of running the Greedy, Fast-Slow Partition, and Randomized Greedy algorithms on the historical data, plus ten alternative data sets, generated by the permutation method described above.

All results were run on an Intel i7-Q740 1.73GHz Quadcore processor with 6GB RAM, running 64-bit Ubuntu 10.10, using only a single thread. Typical CPU times for the Greedy and Fast-Slow Partition were around 35-40 seconds per run, and showed little variation from one data set to another. Typical CPU times for each run of Randomized Greedy were about the same: for a set of 50 runs, the total time was a little over half an hour, again with little variation in run time from one data set to another.

A striking feature of the results is the high degree of variability from one data set to the next for the pure Greedy and Fast/Slow methods, showing the importance of vessel arrival order. The latter method performed slightly better overall in terms of average vessel delays, but at the expense of maximum delay. Randomization reduces variability, and taking the best of a number of random runs clearly gives the best result in terms of all vessel delay statistics: average delay and standard deviation can be halved, and the maximum delay reduced by about a quarter.

## 6. DISCUSSION AND CONCLUSIONS

The model has demonstrated its utility for carrying out what-if analysis of alternative stockyard management strategies. Implementing alternative strategies is simply a matter of changing the priority order of location-stockpile pairs, and any strategy can be tested with very little investment of computation time.

There is potential for improving the algorithms, in a number of respects. Whilst ETA vessel order has the benefit of fairness, avoiding bringing forward some vessels at the expense of others, it may not be the most efficient for the system. Thus alternative approaches should be considered. Furthermore, more intelligence in stockpile position, rail/load point capacity assignment, and reclaimer scheduling could be introduced. Further investigation is essential before a decision support tool could be introduced into the planning process.

## ACKNOWLEDGMENTS

We would like to acknowledge the valuable contributions of Jonathon Vandervoort, Rob Oyston, Janelle Endacott, Tracey Giles, Craig Long, Jeff Mullard and the Capacity Planning Team from the Hunter Valley Coal Chain Coordinator (HVCCC) P/L. Without their patience, support, and feedback, this research could not have occurred. We also thank the HVCCC and the Australian Research Council for their joint funding under the ARC Linkage Grant no. LP0990739.

## REFERENCES

Abdekhodaee, A., Dunstall, S., Ernst, A.T., and Lam, L. (2004). Integration of Stockyard and Rail Network: A Scheduling Case Study. *Proceedings 5th Asia-Pacific Industrial Engineering and Management Systems Conference*, E. Kozan (Ed.) , Gold Coast, Australia, December 12-15, 2004, 1-16.

Ayu and Cardew-Hall, (2002). An application of infinitesimal perturbation analysis (IPA) approach to the optimization of a mining port stockyard system. *Proceedings of the Sixth International Workshop on Discrete Event Systems* (WODES'02), 2 - 4 Oct, 2002, Zaragoza, Spain 264–269.

Bay, M., Crama, Y., Langer, Y., and Rigo, P. (2010). Space and time allocation in a shipyard assembly hall. *Annals of Operations Research*, 179, 57–76.

Binkowski, M., and McCarragher, B.J. (1999). A queueing model for the design and analysis of a mining stockyard. *Discrete Event Dynamic Systems: Theory and Applications*, 9, 75–98.

Conradie, D.G., Morison, L.E., and Joubert, J.W. (2006). Scheduling at coal handling facilities using simulated annealing. *Mathematical Methods of Operations Research*, 68, 277-293.

Glassrock, C.D. and Hoare, R.T. (1995). An application of dynamic simulation to a bulk solids handling facility. *Proceedings of the 5th International Conference on Bulk Materials Storage, Handling and Transportation*, Newcastle, July 1995, Barton, ACT: Institution of Engineers, Australia, 291-297.

Hopper, E. and Turton, B. C. H. (2001). An empirical investigation of meta-heuristic and heuristic algorithms for a 2D packing problem. *European Journal of Operational Research*, 128, 34-57.

Lodi, A., Martello, S. and Monaci, M. (2002). Two-dimensional packing problems: A survey. *European Journal of Operational Research*, 141(2), 241-252.

Marasani, R., Dawood, N. N. and Hobbs, B. (2001). Stockyard layout planning in precast concrete products industry: a case study and proposed framework. *Construction Management and Economics*, 19, 365–377.

Steenken, D., Voß, S., and Stahlbock, R. (2004). Container terminal operation and operations research - a classification and literature review. *OR Spectrum*, 26(1), 3-49.

Newman, A. M., Rubio, E., Caro, R., Weintraub, A., and Eurek, K. (2010). A review of operations research in mine planning. *Interfaces*, 40(3), 222-245.

Welgama, P.S. and Oyston, R. (2003). Study of options to increase the throughput of the Hunter Valley coal chain. *Proceedings of MODSIM 2003*, Townsville, QLD, 14-17 July 2003, 1841-1846.

Wong, A., and Kozan, E. (2010). Optimization of container process at seaport terminals. *Journal of the Operational Research Society*, 61(4), 658-665.