20th International Congress on Modelling and Simulation, Adelaide, Australia, 1–6 December 2013
www.mssanz.org.au/modsim2013

Modelling the Capacity of the Hunter Valley Coal Chain to Support Capacity Alignment of Maintenance Activities

N. Boland [a], **B. McGowan** [b], **A. Mendes** [c] **and F. Rigterink** [a]

[a] *School of Mathematical and Physical Sciences, The University of Newcastle, New South Wales*
Email: natashia.boland@newcastle.edu.au
[b] *Hunter Valley Coal Chain Coordinator, 45 Lambton Rd, Broadmeadow, New South Wales*
[c] *School of Electrical Engineering and Computer Science, The University of Newcastle, New South Wales*

Abstract: The Hunter Valley coal supply chain (HVCC) is the system of logistics facilities – principally a network of rail track and three coal handling terminals – enabling coal mined by producers in the Hunter Valley to be transported, assembled, and loaded onto ships for export. The HVCC serves around 11 producers operating through more than 30 coal load points in the Hunter Valley, transporting coal over rail track extending around 450 km inland, managed by two track owner/operators, via rolling stock from four rail haulage providers that make around 22,000 train trips for approximately 1,400 vessels per year. The HVCC now delivers around 140 million tonnes of coal per annum (Mtpa), with the port of Newcastle exporting more coal by volume than any other facility in the world.

The Hunter Valley Coal Chain Coordinator P/L (HVCCC) is the organization at the heart of this logistics operation. In a landmark for collaborative logistics, the HVCCC was established by HVCC stakeholders to plan and manage the valuable shared infrastructure of the system. The HVCCC provides a range of services vital to the planning and delivery of coal through the logistics system, with its core task to improve the capacity of the coal chain through a centralised planning process.

One of the key ways in which this task is achieved is through the alignment of maintenance activities. All key assets in the HVCC (e.g. rail track sections, coal stacking machinery, terminal conveyor systems) undergo regular preventive maintenance, planned well in advance. While undergoing maintenance, an asset cannot function to deliver coal (or can function only with reduced capacity), thus reducing the capacity of the system. However astute scheduling of these planned maintenance activities releases latent capacity. Such astute scheduling is referred to as *capacity* or *maintenance alignment*, and is a core function of the HVCCC.

The maintenance alignment process at the HVCCC is supported by a model of the system capacity, which quantifies the impact of maintenance activities on the system. This was originally achieved with the assistance of a manual model created in Microsoft Excel, which used as input the impact of each maintenance activity on key assets in the HVCC in terms of the reduction in tonnes per hour that the asset could handle. This was further developed by HVCCC in-house to the current production application known as the Annual Capacity Model (ACM), written in Microsoft C#.net stored in a Microsoft SQL database with business rules stored in Common Knowledge. This application, whilst fit for purpose, has limited scenario testing capability and no optimisation functionality. These issues were a catalyst for collaboration of the HVCCC with the University of Newcastle, leading to the development of two separate but symbiotic prototype optimisation applications: the Capacity Evaluator and the Maintenance Optimiser. The former builds on the model concepts and logic of the current HVCCC ACM to estimate system capacity for a given maintenance schedule. The latter reschedules maintenance activities so as to maximize system capacity.

This paper focusses on the Capacity Evaluator (CE) application, and the recent enhancements to it that have facilitated its adoption. Based on a Linear Programming (LP) model, the CE constructs flows of coal over time so as to maximize total throughput, and offers new features, such as the ability to integrate in-bound and out-bound flows of coal at the terminal stockyards and account for greater complexity in the rail network. However the LP technology also presents new challenges: multiple optimal LP solutions mean that minor changes in input data can result in major variations in patterns of flow observed in the solutions. This paper reports on how this challenge was converted to an opportunity: flexibility in optimal solutions is exploited in a multiobjective approach to achieve flows consistent with contractual targets, while sacrificing little or nothing in terms of throughput. Both the mathematical modelling and decision support processes needed to achieve this, and to ensure the tool is fit for purpose, are described.

Keywords: *preventive maintenance, maintenance scheduling, multiobjective linear programming*

1. INTRODUCTION

The Hunter Valley coal supply chain (HVCC) is the system of logistics facilities – principally a network of rail track and three coal handling terminals – enabling coal mined by producers in the Hunter Valley to be transported, assembled, and loaded onto ships for export. The HVCC serves around 11 producers operating through more than 30 coal load points in the Hunter Valley, transporting coal over rail track extending around 450 km inland, managed by two track owner/operators, via rolling stock from four rail haulage providers that make around 22,000 train trips for approximately 1,400 vessels per year. The HVCC now delivers around 140 million tonnes of coal per annum (Mtpa), with the port of Newcastle exporting more coal by volume than any other facility in the world. A schematic of the HVCC is shown in Figure 1.

Figure 1: A schematic of the Hunter Valley coal chain (HVCC), showing rail track and the locations of mines and coal loading facilities

The Hunter Valley Coal Chain Coordinator P/L (HVCCC) is the organization at the heart of this logistics operation. In a landmark for collaborative logistics, the HVCCC was established by HVCC stakeholders to plan and manage the valuable shared infrastructure of the system. Established in August 2009 from its precursor, the Hunter Valley Coal Chain Logistics Team, as an independent legal entity, the HVCCC provides a range of services vital to the planning and delivery of coal through the logistics system. Its core task is to improve the capacity of the coal chain through a centralised planning process, so as to improve financial returns for HVCCC members (producers and service providers). This brief has been expanded since to include the HVCC contractual commitments of its members. An overview of the HVCC system and some of the planning challenges faced by the HVCCC is given in (Boland and Savelsbergh, 2012).

One of the key ways in which this task is achieved is through the alignment of maintenance activities. All key assets in the HVCC (e.g. rail track sections, coal stacking machinery, terminal conveyor systems) undergo regular preventive maintenance, planned well in advance. Examples of key assets for which such maintenance is planned are shown in Figure 2, with a shiploader pictured on the left and a bucket-wheel reclaimer shown on the right. While undergoing maintenance, an asset cannot function to deliver coal (or can function only with reduced capacity), thus reducing the capacity of the system.

Figure 2: Examples of terminal handling equipment that must undergo regular preventive maintenance

However astute scheduling of these planned maintenance activities releases latent capacity in the system. An illustration is given in Figure 3. Here the rail track owner, the ARTC, has three 5-hour maintenance jobs to carry out on rail track inbound to the port terminal operated by PWCS, which must carry out a 15-hour maintenance job on its rail receival equipment. By scheduling the terminal maintenance to include all 5 hours of one of the rail maintenance jobs, the impact on capacity of the rail job is absorbed by the terminal job; instead of 30 hours of lost capacity, only 25 hours is lost. (Note that even more lost capacity could be recovered if all three track jobs could be carried out in parallel with the terminal job, but this may not be possible, for example if the track work requires a specially qualified crew, who must rest between shifts.

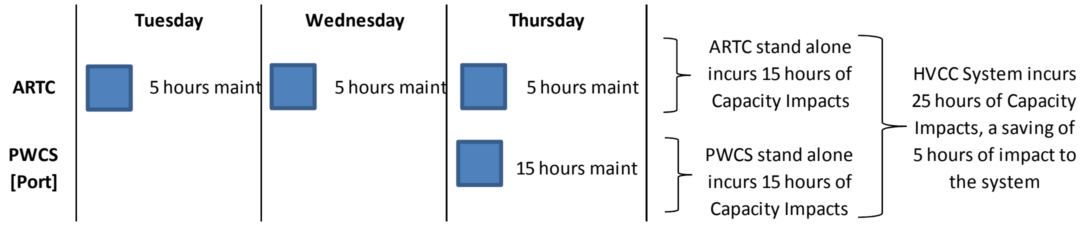

Figure 3: A simple example of how maintenance alignment releases capacity back into the system

Such astute scheduling is referred to as *capacity* or *maintenance alignment*, and is a core function of the HVCCC. In this process, the individual service providers (the terminal managers and track owner/operators) each prepare an initial maintenance schedule for the assets under its management. These are submitted to the HVCCC, which then seeks to re-time some maintenance activities so as to increase system capacity. The proposed schedule changes are negotiated with the service providers. This process is supported by a model of the system capacity that quantifies the impact of maintenance activities on the system. Originally achieved with the assistance of a manual model created in Microsoft Excel, which used as input the impact of each maintenance activity on key assets in the HVCC in terms of the reduction in tonnes per hour that the asset could handle, in 2010 the model was further developed by HVCCC in-house to the current production application known as the Annual Capacity Model (ACM), written in Microsoft C#.net stored in a Microsoft SQL database with business rules stored in Common Knowledge. The ACM, whilst fit for purpose, has limited scenario testing capability and no optimisation functionality. These issues were a catalyst for collaboration of the HVCCC with the University of Newcastle, leading to the development of two separate but symbiotic prototype optimisation applications: the Capacity Evaluator and the Maintenance Optimiser. The former builds on the concepts and logic of the ACM to estimate system capacity for a given maintenance schedule. The latter reschedules maintenance activities so as to maximize system capacity (Boland et al., 2011 & 2012). Both have undergone testing by the HVCCC over several years, with enhancements and functionality added to ensure they are fit for purpose, and in 2013 are actively in use for current planning.

This paper focusses on recent enhancements to the Capacity Evaluator (CE) application that have facilitated its adoption. Based on a Linear Programming (LP) model, the (CE) constructs flows of coal over time so as to maximize total throughput, and offers new features, such as the ability to integrate in-bound and out-bound flows of coal at the Newcastle terminal stockyards and to account for greater complexity in the rail network. However the LP technology also presents new challenges: multiple optimal LP solutions mean that minor changes in input data can result in major variations in patterns of flow observed in the solutions. This paper reports on how this challenge was converted to an opportunity: flexibility in optimal solutions is exploited in a multiobjective approach to achieve flows consistent with contractual targets, while sacrificing little or nothing in terms of throughput. Contractual targets incorporate expectations of both infrastructure usage and production forecasts for the different parts of the system, leading to the desire to balance flows through different parts of the systems in accord with prescribed ratios.

We first summarize related literature, then give a brief overview of the CE model. We proceed to explain how we model the requirement to balance flows according to prescribed ratios in the LP model, leading to *soft proportionality constraints*, and present our approach to handling the concomitant multiple objectives. The results of computational experiments showing how these objectives can be traded off are shown.

2. RELATED LITERATURE

The literature on preventive maintenance focusses primarily on maintenance policies: how often each type of maintenance should be performed, or under what conditions. For a broad perspective on this topic, we refer the reader to Sharma et al (2011) and Budai et al (2008), which discuss the advantages of considering the impact of planning maintenance on production. In Budai et al (2008), three main directions are identified: (a) costing maintenance activity, (b) carrying out maintenance at opportune moments (e.g. when a breakdown has occurred), and (c) scheduling maintenance in line with production. The HVCCC process has elements of (a) and (c): the cost of a maintenance plan is assessed in terms of its impact on the throughput and the HVCCC is seeking to schedule maintenance in line with production. On the topic of rail track maintenance schedule, Budai et al (2006) deal with multi-component system maintenance. However, they aim at minimizing the costs of maintenance activities, which are not an issue in the HVCC setting tackled here; in this setting all jobs must be done, so the cost of maintenance is a sunk cost.

On the general topic of Decision Support Systems, including those for supply chain planning, we mention the series of review papers culminating in Eom and Kim (2006), and the work of Power and Sharda (2007). Another relevant study is a recent review on supply chain performance measurement (Shepherd and Gunter, 2011), in which 89 papers are identified and analysed.

3. THE CAPACITY EVALUATOR (CE) LINEAR PROGRAMMING MODEL

3.1. Inputs

The key inputs to the Capacity Evaluator (CE) Linear Programming (LP) model for estimating the HVCC system capacity under a given preventive maintenance schedule are:

- a network representing physical and logical links along which coal can flow, together with the maximum rate of flow of coal on each arc, when it is operating normally (not undergoing maintenance);
- a planning horizon, consisting of a start date and time and end date and time (typically a calendar year);
- a set of commodities K representing immiscible coal types flowing through the network, with each having a given source node, sink node and individual flow rate limits on each arc; and
- a maintenance activity schedule, giving for each maintenance activity its start date and time, end date and time, the set of network arcs affected by the maintenance, and the fraction of its normal maximum flow rate lost as a result of the activity.

From these inputs, the planning horizon is modelled as the interval $[0,T]$ and the start and end times of each maintenance activity ordered in a sorted list $0 = t_0 < t_1 < ... < t_m = T$ so that the state of the system during any *time slice* $[t_{i-1}, t_i]$ is constant (no maintenance activity starts or ends within the interval), for all $i = 1,...,m$. A *planning network* denoted by $G = (N, A)$, with node set N and arc set $A \subseteq N \times N$, is constructed by taking a copy of the given network for each time slice, linked by arcs connecting copies of nodes modelling locations where coal can be stored from one time slice to the next. During each time slice, the capacity of each arc, i.e. the maximum total rate of flow it can sustain during the time slice, can be determined from the set of maintenance activities occurring during the time slice, the set of arcs they affect, and the impact on their normal flow rate. We use u_a to denote the maximum rate of flow possible on arc $a \in A$ modelling a given network arc, and write $i(a) \in \{1,...,m\}$ to indicate its corresponding time slice. On arcs between storage nodes in subsequent time slices, u_a is the (instantaneous) capacity of the storage location. We note that there are also other parameters, such as minimum desired amount of coal to keep in each storage location, but we omit these for simplicity of exposition.

3.2. Decision Variables

The primary decision variables are x_a^k, indicating the flow of coal type $k \in K$ along arc $a \in A$ in the planning network. There are also auxilliary variables, designed to model the flow into and out of each storage location during each day, so as to model the requirement that coal stored in the terminal is usually held for a minimum (e.g. 3 days) and no more than a maximum (e.g. 10 days) number of days, however we omit these details here for the sake of simplicity of exposition.

3.3. Constraints

There are two key classes of constraints in the model.

- Each coal type must satisfy flow balance constraints at all nodes in the planning network other than at its designated source and sink, and must satisfy its individual capacity constraints.
- The total flow on any arc in the network cannot exceed the maximum allowed by the maintenance schedule, modelled as

$$\sum_{k \in K} x_a^k \le u_a (t_{i(a)} - t_{i(a)-1}), \quad \forall a \in A \setminus S \qquad \text{and} \qquad \sum_{k \in K} x_a^k \le u_a, \quad \forall a \in S,$$

  where $S \subseteq A$ is the set of arcs in the planning network representing storage between time slices.

Other constraints include the option to set initial and final desired quantities in storage at each storage location, or to ask that final quantities stored equal initial quantities. Full details are omitted for simplicity.

3.4. The Objective Function

The goal of the maintenance alignment process is the maximize the capacity of the system to deliver coal, hence the LP model objective of the CE is simply to maximize the sum of all coal reaching its sink node:

$$MaxFlow = \max \sum_{k \in K} \sum_{a \in \delta^-(e_k)} x_a^k,$$

where e_k denotes the sink node for coal type $k \in K$ and $\delta^-(v)$ indicates the set of arcs entering node v in the planning network.

3.5. Outputs and Key Metrics

In addition to the maximum flow through the system given by the optimal LP objective value, key outputs of interest to the HVCCC are

- daily inventory levels at each storage location,
- daily flows into each terminal, and
- daily flows out of each terminal.

To facilitate capture of these, the list of time points used to define time slices always includes midnight of each day. Thus, for example, the inventory level at each storage location at the end of each day is given by the flow on the arc linking the copy of the node modelling the storage locations in the time slice immediately preceding midnight, to the copy of the node for the next time slice.

4. SOFT PROPORTIONALITY CONSTRAINTS AND MULTIPLE OBJECTIVES

The CE LP was implemented in C++, using gcc v4.6.1 (64-bit), IBM ILOG Cplex v12.4, and runs reliably in less than a minute (CPU time to read input files, build and solve the model, and save the results) on an Intel Xeon X5650 platform. It was observed in testing that multiple optimal solutions to the LP were always available. These can arise in several ways. For example, flow in a time slice of one type of coal can easily be exchanged for flow from another, provided the latter has spare capacity in its individual constraints (the constraining factor must be in the joint capacity constraints). Flow in-bound to a storage location in one time slice could also be exchanged with flow in another time slice, if a later out-bound capacity constraint is a bottleneck. Whilst multiple optimal solutions do not affect the key indicator of interest, the maximum system throughput (total flow), they can give rise to quite different outputs in terms of other indicators of interest, such as daily inventory levels, or daily flows into each terminal. Planners were thus finding that small changes in the input data could result in quite different indicators of this type. Furthermore, flows into each terminal should reflect the contractual arrangements in the system, and these had yet to be modelled. Thus the CE LP discussed in Section 3 is not fit for purpose, and additional features are needed, as follows.

4.1. Proportionality Constraints

In the HVCC, producers enter into contracts with buyers and service providers. In the former case, these contracts will, to a large extent, dictate how much coal they expect to ship through the system from each mine during the coming year. In the latter case, contracts will dictate what quantities of coal can be expected to flow through different parts of the infrastructure, for example, what can be expected to flow through each terminal. When undertaking advance planning, the HVCCC works with the expectation that, unless maintenance activities force change, these flows will be *consistent*. In other words regardless of where the bottlenecks are, if there were no maintenance, the flow in each day in each part of the system should be identical. In order for the CE model to better reflect this requirement, and the reality of contractual arrangements, the HVCCC sought the ability to balance flow as follows. Given a (small) set of n arcs (typically two or three) in the original network, which we call a *proportionality set*, and given desired proportions for the relative magnitude of flow on each of these arcs, i.e. given $p_1, p_2, ..., p_n \in [0,1]$, where $\sum_{j=1}^{n} p_j = 1$, try to ensure that

$$\frac{y_{jd}}{\sum_{h=1}^{n} y_{hd}} \approx p_j, \quad \forall j = 1,...,n, \quad \forall d \in D,$$

where y_{jd} denotes the total amount of flow on the jth arc in the set during day $d \in D$, the set of all days in the planning horizon. Note that these additional "book-keeping" variables can easily be calculated from the flow variables x in the original model, by summing over coal types and arcs in the planning network corresponding to the arc in the proportionality set during time slices within day d.

We model this requirement as a soft constraint, introducing a new objective to the problem. We introduce new variables $\mu_{jd} \geq 0$ to model the deviation of flow from the desired proportion via the constraints

$$\mu_{jd} \geq y_{jd} - p_j \sum_{h=1}^{n} y_{hd}, \quad \forall j = 1,...,n, \quad \forall d \in D,$$

with new (additional) objective

$$MinDeviation = \min \sum_{j=1}^{n} \sum_{d \in D} \mu_{jd}.$$

Note that because the proportions all sum to 1, and by non-negativity of the μ variables, this objective value is – at optimality – equivalent to

$$\frac{1}{2}\sum_{j=1}^{n}\sum_{d\in D}\left| y_{jd} - p_j \sum_{h=1}^{n} y_{hd} \right| = \frac{1}{2}\sum_{j=1}^{n}\sum_{d\in D}\left(\sum_{h=1}^{n} y_{hd}\right)\left| \frac{y_{jd}}{\sum_{h=1}^{n} y_{hd}} - p_j \right|.$$

Also note that the form of the constraints modelling deviation mean that days on which flows are higher are implicitly weighted more heavily in the objective function: as the right-hand side of the above expression shows, the relative deviation from the desired proportion is weighted by the total flow on all arcs in the set for the day. This means that there is more incentive to achieve the desired proportion on days less affected by maintenance, which is just what is required by the HVCCC.

4.2. Handling Multiple Objectives

The HVCCC would like to be able to consider one or two proportionality sets, leading to the possibility of up to three objectives: maximizing total flow, and minimizing deviation from desired proportions for each of the two proportionality sets. Currently the HVCCC is satisfied with combining the deviation for the two proportionality sets in a single objective, minimizing the sum of the two, or *total deviation*. The resulting biobjective problem is addressed with the following heuristic. First, throughput is maximized without consideration of proportionality constraints, to determine the "unconstrained" maximum throughput. The total deviation of the solution found is an upper bound on total deviation in Pareto-optimal solutions. Then throughput is maximized subject to an extra constraint forcing the total penalty to be zero. This will provide a solution with the desired proportions being perfectly met, but at a high cost in terms of lost capacity. The throughput of this solution is a lower bound on Pareto-optimal throughput values. Finally, for each throughput value equally spaced between the lower bound and the unconstrained throughput, (the number of values tested is a user-defined parameter), the total proportionality deviation is minimized subject to a constraint requiring throughput achieve the value stipulated.

5. EXPERIMENTAL RESULTS AND CONCLUSIONS

Using the HVCCC's annual capacity planning data for 2012 with 1,244 maintenance jobs and two proportionality sets, we approximated the efficient frontier for the maximum flow (throughput) versus the total deviation for the two sets by calculating Pareto-optimal solutions for 10 throughput values, equally spaced between the maximum flow, and the flow at which the deviation first has the minimal value of zero. These, together with the point generated by maximizing flow while ignoring proportionality deviation, are shown in Figure 4. The area of most interest to the HVCCC is the area to the right of the curve, where significant decreases in proportionality deviation can be achieved with little or no reduction in total flow. For the solutions corresponding to the two right-most points in the curve, the daily flow proportions on each arc in one of the proportionality sets (having three arcs) are shown in Figure 5, for

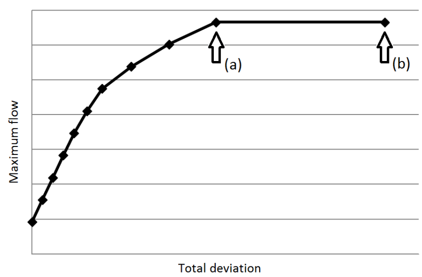

Figure 4: The efficient frontier for maximum flow (vertical axis) vs total deviation for two proportionality sets

the entire year. These show that significant reductions in the variability of the daily proportions have been achieved, without any reduction in flow. The two solutions have average relative deviation from desired proportions of 3.30% versus 1.73%, thus the relative deviation has been almost halved without any loss in flow, giving much more constant daily flows.

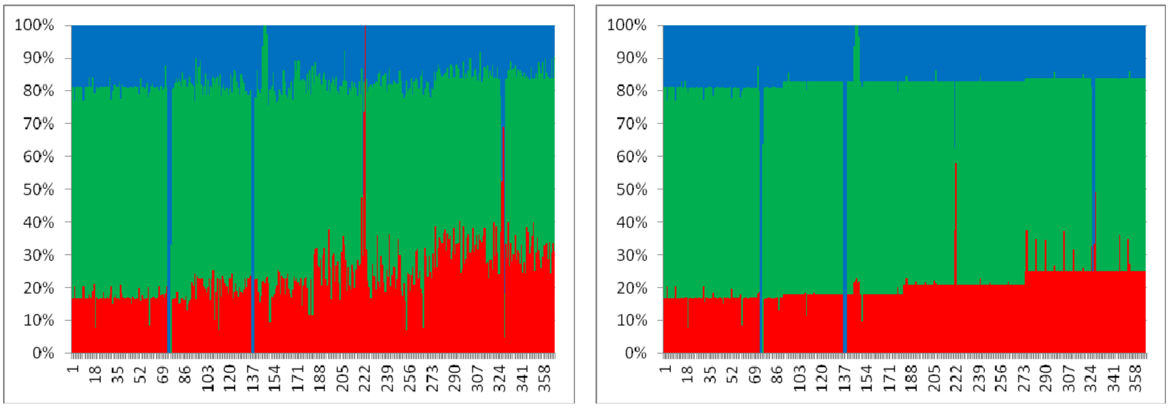

Figure 5: The effect on one proportionality set of minimizing deviation without flow loss. The graphic on the left refers to point (b) in Figure 4; and the graphic on the right to point (a). It is important to notice that both solutions have the same maximum flow, but one presents more stable daily flows than the other.

ACKNOWLEDGMENTS

This research was funded by the Australian Research Council and the Hunter Valley Coal Chain Coordinator under ARC Linkage Project grant LP0990739. We also gratefully acknowledge the generous contribution of time and expertise from many HVCCC personnel, making particular note of the inspiration and assistance of Rob Oyston, Wayne Johnson, Katherine McIlwain and Martin Jackson in the success of this research project.

REFERENCES

Boland, N., Kalinowski, T., Waterer, H., and Zheng, L. (2011). “An optimisation approach to maintenance scheduling for capacity alignment in the Hunter Valley coal chain”. In: E. Y. Baafi, R. J. Kininmonth and I. Porter (eds.), *Proceedings of the 35th APCOM Symposium: Applications of Computers and Operations Research in the Minerals Industry*, 24-30 September 2011, Wollongong, Australia, 887-897.

Boland, N., Kalinowski, T., Waterer, H., and Zheng, L. (2012). “Mixed integer programming based maintenance scheduling for the Hunter Valley coal chain”, *Journal of Scheduling*, Special Issue on Maintenance Scheduling, H. Murat Afsar, C. Artigues, E. Bourreau and E. Ozcan (eds.), available online.

Boland, N., and Savelsbergh, M. (2012). “Optimizing the Hunter Valley coal chain”. In: H. Gurnani, A. Mehrotra and S. Ray (eds.), *Supply Chain Disruptions: Theory and Practice of Managing Risk*, Springer London, 275-302.

Budai, G., Dekker, R., and Nicolai, R.P. (2008). “Maintenance and production: A review of planning models”. In Kobbacy, K.A.H., and Murthy, D.N.P. (eds.), *Complex System Maintenance Handbook, Springer Series in Reliability Engineering*, 321-344. Springer.

Budai, G., Huisman, D., and Dekker, R. (2006). “Scheduling preventive railway maintenance activities”, *Journal of the Operational Research Society*, 53:1035-1044.

Eom, S., and Kim, E. (2006). “A survey of decision support system applications (1995–2001)”, *Journal of the Operational Research Society*, 57:1264-1278.

Power, D.J., and Sharda, R. (2007). “Model-driven decision support systems: Concepts and research directions”. *Decision Support Systems*, 43:1044-1061.

Sharma, A., Yadava, G.S., and Deshmukh, S.G. (2011). “A literature review and future perspectives on maintenance optimisation”, *Journal of Quality in Maintenance Engineering*, 17:5-25.

Shepherd, C. and Gunter, H. (2011). “Measuring supply chain performance: current research and future directions”. In Fransoo et al. (eds.), *Behavioral Operations in Planning and Scheduling*, 105-121.