Computers & Operations Research 53 (2015) 54–67

Contents lists available at ScienceDirect

Computers & Operations Research

journal homepage: www.elsevier.com/locate/caor

A decision support tool for generating shipping data for the Hunter Valley coal chain ☆

Natashia Boland, Martin Savelsbergh *, Hamish Waterer

School of Mathematical and Physical Sciences, University of Newcastle, Australia

ARTICLE INFO

Available online 8 August 2014

Keywords:
Strategic planning
Coal export supply chain
Blending
Data generation
Decision support
Hierarchical optimization
Integer programming

ABSTRACT

Strategic capacity planning is a core activity for the Hunter Valley Coal Chain Coordinator as demand for coal is expected to double in the next decade. Optimization and simulation models are used to suggest and evaluate infrastructure expansions and operating policy changes. These models require input data in the form of shipping stems, which are arrival streams of ships at the port, together with their cargo types and composition. Creating shipping stems that accurately represent future demand scenarios has been a time-consuming and daunting challenge. We describe an optimization-based decision support tool that facilitates and enhances this process, and which has become an integral part of the company's work flow. The tool embeds sampling to enable the generation of multiple shipping stems for a single demand scenario, employs targets, and desirable and permissable ranges to specify and control the characteristics of the shipping stems, and uses integer programming in a hierarchical fashion to generate shipping stems that best meet the set goals.

1. Introduction

The Australian coal export industry is the largest in the world. Australian coal exports play a significant role in supporting the Australian economy, and in meeting the world's need for coal. Australia is the largest exporter of metallurgical coal in the world, and the second-largest exporter of thermal coal, with about 60% of world market share in the former and about 20% in the latter. These exports make up almost 15% of Australia's total export earnings.

The Hunter Valley coal region, exporting coal through the Port of Newcastle, is a major contributor to this economic powerhouse. It is the primary source of thermal coal exports in Australia. In 2008 Newcastle port's throughput was around 92 million tonnes, which constituted more than 10% of the world's total trade in thermal coal.

Accommodating future growth is one of the most pressing challenges, if not the most pressing challenge, facing the Hunter Valley coal chain. Demand for coal is expected to continue to grow in the next decade due to an increasing need for coal from power generation in emerging and expanding economies. In response to this challenge, the producers and logistics providers created a collaborative logistics company, the Hunter Valley Coal Chain Coordinator (HVCCC), to centralize and coordinate analysis and planning for the system as a whole.

Increasing the coal chain's annual throughput will be accomplished in two ways: (1) by expanding the infrastructure, and (2) by improving the operational efficiency. HVCCC carries out the analysis critical to planning both infrastructure and operational change using a suite of optimization and simulations models (e.g. [11,12]). Their long-term models assess the likely impact of infrastructure capacity expansions on throughput in future years, providing the basis on which the producers commit to infrastructure projects. Such projects constitute many millions and even billions of dollars of investment, underlining the importance of the HVCCC's analysis. The HVCCC also uses shorter-term models to assess the system capacity in the coming year, providing the basis on which producers contract with their customers and with logistics providers to ensure they get their products to market in the coming year.

However, the HVCCC found it increasingly difficult to generate the shipping demand data required as input for their optimization and simulation models. Prior to having access to the tool discussed in this paper, all shipping demand data was generated by hand. As we discuss later in more detail, this has always been a highly involved task, taking substantial time to develop a single (annual) data set. However, two factors were making the task virtually impossible: (1) there was increasing pressure to run all models with multiple possible demand data sets to ensure decisions are robust to likely variations in the nature and timing of demand; and (2) with export volumes anticipated to rise by as much as 100%

☆This research was supported by the ARC Linkage Grant no. LP0990739.
* Corresponding author.

http://dx.doi.org/10.1016/j.cor.2014.07.016
0305-0548/

over the next decade, the larger scale was making it increasingly difficult to construct demand data sets that meet all requirements. (We discuss later the nature of these requirements, and why meeting them is highly complex.) The planning team was simply unable to achieve the required demand data characteristics using a manual process, and was spending untenable amounts of their time in creating the data needed to carry out their analysis mission.

The decision support tool for automatic shipping stem generation discussed in this paper requires negligible planner time and consistently generates demand data sets that meet all requirements, even for the largest export volume scenarios. The tool uses a multi-phase approach that relies on a combination of integer programming and sampling, thus allowing multiple demand data sets to be generated for the same export volume scenario. The tool is in daily use by planners at the HVCCC as an integral part of their analysis framework, providing an industry critical to both the Australian economy and the world's energy supply with more accurate and more robust quantitative assessment of capacity, and greater confidence in the quality of its advice.

1.1. The hunter valley coal chain

The Hunter Valley coal chain (HVCC) refers to the inland portion of the coal export supply chain in the Hunter Valley, New South Wales, Australia. The coal is mined and stored either at a railway siding located at the mine or at a coal loading facility (used by several mines). The coal is then transported to one of the three terminals at the Port of Newcastle, almost exclusively by rail. The coal is dumped and stacked at a terminal to form stockpiles. Once a ship arrives at berth, coal intended for that ship is reclaimed from stockpiles and transferred via conveyor belt and ship loading equipment into the ship's holds. The ship then takes the coal to its final destination.

Each coal mine is owned by a *producer*, who may own more than one. Each producer has individual contracts with *customers* to supply one or more *brands* of coal. The brand of coal dictates its characteristics, for example the range in which its calorific value, ash, moisture and/or sulphur content lies. The majority of customer contracts are long-term and on-going. In order to deliver a brand required by a customer, coal from different mines, producing coal with different characteristics, is "mixed" in a stockpile to obtain a blended product meeting the required brand characteristics. Coal is transported to the customers using *vessels*. Each vessel arriving at the port collects quantities of one or more brands of coal. A *cargo* is defined by a brand and the quantity of that brand to be carried on a single vessel. Each vessel belongs to a *vessel class*, largely characterized by its capacity. A vessel berths at a coal *terminal*. Each terminal is managed by a terminal *operator*. Each producer has individual agreements with the terminal operators to export a certain amount of coal through one or more of their terminals.

1.2. Shipping stems

Future demand scenarios for the HVCC are determined by yearly mine production forecasts (provided by the producers). As the HVCC can be viewed as operating as a pull-system, in which the nomination of a vessel, i.e., the announcement of the vessel's pending arrival at the port two to three weeks in advance, triggers the production, transportation, storage, and ultimately the loading of coal onto the vessel, the mine production forecast needs to be mapped into a stream of vessel arrivals at the port. Such a stream of vessel arrivals is referred to as a *shipping stem*. Each vessel arrival, referred to as a *trip*, is characterized by the following: an arrival time; the terminal at which the vessel is to be loaded; a *cargo-profile*, which specifies the brands of coal to be carried and the tonnage of each, constituting the vessel's cargoes; and, for each brand carried, its associated *brand-recipe*, which specifies the mines that will contribute to the brand, and the proportion from each to be mixed in the brand (blend). An example of a trip is given in Fig. 1. The coal in a cargo sourced from one mine is referred to as a *component*, so in the example, Brand 74 has two components, one from Mine 25 and the other from Mine 37. The brand-recipe for Brand 74 in this trip is 13.9% from Mine 25 and 86.1% from Mine 37.

The importance of generating shipping stems that are representative of what the future may bring cannot be overemphasized: many millions, if not billions, of dollars worth of decisions will be made based on models taking these shipping stems as input. Furthermore, generating such shipping stems is not simply a matter of scaling a historical stem. The shipping stem has to reflect that the HVCC is changing; new mines are brought on line, existing mines are (temporarily) shut down or operating at reduced capacity, new brands and new brand-recipes are introduced, new terminals start their operations, and demand management changes the arrival patterns of vessels.

1.3. Relevant literature

In general the challenges associated with generating appropriate data sets for computational studies are well-known and recognized by the optimization community, but relatively few papers focus exclusively on issues surrounding data generation. An exception is the recent paper by Hall and Posner [6] which succinctly captures the importance of data generation: "In many experiments, the methods chosen to generate synthetic data can significantly affect the results of an experiment." A different, but equally important, perspective is provided by Reilly [10] who emphasizes the importance of understanding and properly inducing correlations between characteristics of data, which is one of the major drivers and, at the same time, one of the major challenges in our stem generation research.

The impact of ship arrival patterns on the productivity of a port has been acknowledged widely and port infrastructure designs and port operating procedures should ideally be tailored to the anticipated ship arrival patterns. Van Asperen et al. [1] study the impact of three different ship arrival patterns on port efficiency. The first assumes Poisson interarrival times and represents an "uncontrolled" arrival process. The second assumes equidistant arrivals, in which ships arrive perfectly smoothly during the year, at equally spaced intervals. The third assumes scheduled arrivals, in which ships' arrival times are coordinated with activities at the port. The latter two represent "controlled" arrival processes. Small random perturbations are introduced in the controlled arrival process to add more realism. Van Asperen et al. simulate port

arrival time	11 October 2008 13:40
terminal ID	3
vessel ID	438
vessel class	Panamax

		component ID		
brand ID	brand tonnage	2	25	37
42	15,086		1.000	
62	36,433	1.000		
74	38,983		0.139	0.861
trip tonnage	90,502			

Fig. 1. An example of a trip.