Submitted to *INFORMS Journal on Computing*
manuscript JOC-2016-03-OA-056

Stochastic Network Design for Planning Scheduled Transportation Services: The Value of Deterministic Solutions

Xin Wang
Department of Management Science, Lancaster University Management School, UK
Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, Trondheim, Norway, xin.wang@iot.ntnu.no

Teodor Gabriel Crainic
Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT), Québec, Canada
Département de management et technologie, École des sciences de la gestion, UQAM, Montréal, Canada, TeodorGabriel.Crainic@CIRRELT.net

Stein W. Wallace
Center for Shipping and Logistics, Department of Business and Management Science, Norwegian School of Economics, Bergen, Norway, Stein.Wallace@nhh.no

We study the value of deterministic solutions, in particular their quality and upgradability, in addressing stochastic network design problems, by analyzing their time-dependent formulations known as Scheduled Service Network Design (SSND) problems in freight transportation planning. We study several problem variants and models and investigate, for each case, the immediate quality of the deterministic solutions stemming from the 50^{th} and the 75^{th} percentile of the demand distributions. We then show that for all models, but in different ways, we are able to make effective use of parts of the deterministic solution, confirming the value of the deterministic solution in the stochastic environment, even when the deterministic solution itself performs badly.

We also investigate what makes the optimal stochastic solution better in the stochastic environment than other feasible solutions, particularly those obtained by addressing deterministic versions of the problem. We do this by quantitatively analyzing the structures of different solutions. A measurement scheme is proposed to evaluate the level of potentially beneficial structural properties (multi-path usage and path-sharing) in different solutions. We show that these structural properties are important and correlated with the performance of a solution in the stochastic environment.

Key words: Service network design; stochastic programming; value of deterministic solution

1. Introduction

Network design is the methodology of choice to assist for a large gamut of issues related to the tactical planning of operations for consolidation-based freight transportation carriers. The so-called *Scheduled Service Network Design* (*SSND*) formulations aim at deciding where, when, and how to offer transportation services to satisfy the demand for movement among the terminals operated by the carrier, generated by multiple customers with diverse products. The goal is to decide the selection, routing and scheduling of transportation services, while balancing the operating costs and service quality. Network design problems are NP-Hard for all but trivial cases and their combinatorial, mix-integer formulations are difficult to address.

This is even more so when some parameters of the model are uncertain. In freight transportation problems, the most commonly modeled uncertain phenomenon is demand. Its uncertainty can be represented by a set of *scenarios* approximating a "known" demand distribution (multi-dimensional in the case of multi-commodity or multi-source/sink problems). When the number of scenarios increases, the uncertainty of the demands is better represented (assuming the scenarios are well constructed), but the corresponding model will eventually become numerically unsolvable. One way out of this problem is to solve the stochastic program heuristically. Examples can be found in, e.g., Hoff et al. (2010) and Crainic et al. (2011).

In this paper, we do not try to develop another heuristic method for solving stochastic problems. Instead, we seek to understand the *value* that deterministic solutions (i.e., solutions from deterministic models) have in the context of stochastic service network design problems. We believe the insights derived from this study can guide in the construction of efficient heuristics for stochastic network design problems.

We investigate the quality of deterministic solutions in the stochastic environment, as well as how such deterministic solutions can be used to construct better-performing solutions to stochastic models. This is motivated by the fact that, compared to the "full" stochastic model with all scenarios considered, it is usually much easier to solve a deterministic counterpart, where all random elements take some fixed values, e.g., their means. There are, therefore, situations when an optimal (or near-optimal) deterministic solution can be found for a service network design problem, while the optimal solution to the stochastic model cannot for numerical reasons. In these situations, it is useful to know whether it is possible to extract valuable information from the deterministic solution, and how we can make use of such information in finding a good solution to the stochastic case.

We perform the study in the context of the SSND problem settings and formulations, and aim to provide a comprehensive analysis with regards to the value of deterministic solutions. Common to all SSND models are the time-dependent characteristics of services and demands represented by a time-space network. Integer and continuous decision variables stand for the selection of services at particular time instants and the routing of demand flows on the resulting network, respectively. We introduce several models with fixed and variable (integer and continuous) capacity, as well as with and without resource-management constraints. We investigate, for each model, the immediate quality of the deterministic solutions stemming from the 50^{th} and the 75^{th} percentile of the demand distributions. We then explore the possibility of *upgrading* these deterministic solutions to good solutions for the stochastic models. Finally, we show that for all models, but in different ways, we can make effective use of parts of the deterministic solutions to arrive at good solutions to the stochastic models, confirming the value of deterministic solutions in stochastic environments.

In addition, we seek to understand what structural features should a solution have in order to perform better in a stochastic environment. We do this by examining the structural differences between the optimal solutions to the stochastic models and the ones stemming from their deterministic versions, and try to find out what makes a stochastic solution behave better than its deterministic counterpart. Lium et al. (2007, 2009) have studied a version of the stochastic service network design problem with fixed capacity and resource-management constraints and indicate that certain structural features, such as multi-path usage and path-sharing, offer better solutions when there are uncertainties in demand. Inspired by these insights, we summarize and confirm these potentially beneficial features for all the models introduced in this paper, and propose a measurement scheme to *quantify* the level of such structural features for different solutions. Using the measurement scheme, we may then see how the level of the potentially important structural features of a solution is related to its performance in a stochastic environment, and thereby understand why and how a deterministic solution may be upgraded to a good solution for the stochastic case.

The contribution of this paper is to provide a first attempt on a complete and comprehensive analysis of the quality and upgradability of deterministic solutions to stochastic scheduled service network design problems. Additionally, inspired by the potentially beneficial structural features from Lium et al. (2007, 2009), we present a measurement scheme to quantitatively show how a deterministic solution may be upgraded. These analyses provide intuitive and easy to comprehend (and implement) insights that may better the understanding for stochastic models in the service network design community, e.g., how much does one lose by not using stochastic models and what kind of beneficial solution characteristics should one look for when tackling a stochastic problem. For academics, these insights may be incorporated into the development of a heuristic. For practitioners, we offer several simple alternative means that decision makers may use to produce good solutions in real applications where, most of the time, only deterministic models can be solved.

This paper is organized as follows. In Section 2, some important issues in freight transportation and service network design are reviewed. Section 3 introduces the stochastic scheduled service network design problem and presents models for several problem settings. We propose in Section 4 a set of comparison tests for the examination of quality and upgradability of deterministic solutions. Section 5 introduces the problem instances and presents the computational study. We then conclude in Section 6.

2. Literature Review

Transportation is an important domain of human activity. It supports and enables many other social and economic activities and exchanges. Freight transportation, in particular, is one of today's most important activities. Demand for freight transportation reflects the need to move goods between producers and consumers and requires a rather complex system which derives from the fact that the distances separating them are often significantly long. Crainic (2003) gives a general presentation of freight transportation players, questions, and problem classes. In an increasingly competitive environment, carriers seek to offer reliable, high quality services to their customers at a lowest possible cost, and in the mean time make a profit.

Transportation systems are often based on consolidation, where one vehicle or convoy may serve more than one customer. So, in a system where demand for transportation is represented by origin-destination (OD) pairs, freight of different OD pairs, with different origins and destinations, are combined into common vehicles. This typically happens with railways, Less-Than-Truckload (LTL) motor carriers, container shipping lines and postal services.

The underlying structure of a consolidation transportation-based system normally consists of a large network of terminals and the transportation operations are hence usually rather complex. This is in contrast to customized transportation, which provides dedicated service for each OD pair. Consolidation-based transportation carriers usually operate so-called hub-and-spoke networks to take advantage of economies of scale. In such systems, low-volume demands are first delivered to an intermediate terminal or a hub to be grouped and consolidated. High-frequency, high-capacity services are provided between the hubs, and can thus allow a much higher frequency of service between all the OD pairs. However, routing through several intermediate terminals and hubs would inevitably result in longer transport distances and more time spent at terminals and can sometimes cause serious delays. There is a great deal of literature on the subject. Surveys are presented by Christiansen et al. (2004, 2007, 2013) for maritime transportation, Cordeau et al. (1998) for rail transportation, Crainic and Laporte (1997) and Crainic (2003) for land-based long-haul transportation, Crainic and Kim (2007) for intermodal transportation, and Crainic (2000) for service network design in freight transportation.

In order to satisfy the demand of customers more timely and reliably, consolidation carriers operate a selection of *services*, each characterized by such as its route, vehicle type, frequency, and capacity. Internally, services are often collected in an operational plan (also referred to as load or transportation plan), generally accompanied by a schedule that indicates departure and arrival times at the terminals of the route (Crainic and Kim 2007). Service network design formulations are used to build such a (scheduled) transportation plan for the next operating period.

Service network design problems address a set of major issues and decisions relevant for consolidation-based carriers: the selection and scheduling of the services to operate, the routing of freight for each OD pair and the consolidation operations at terminals. The goal is to achieve profitable operations while providing timely and reliable services according to customer expectations. The corresponding models usually take the form of network design formulations. With the complicated interactions among system components and decisions, as well as the trade-offs between operating costs and service quality, service network design models are very difficult to solve, and thus heuristics are usually the solution method of choice.

Reviews on the formulation of service network design models are presented by Crainic (2000, 2003), Delorme et al. (1988) and Cordeau et al. (1998). Efforts have been made towards both static and scheduled service network design formulations. The former assume a static demand throughout the whole planning period. The time dimension of the service network design is then implicitly considered through the definition of services and inter-service operations at terminals. Such models have been proposed for multi-modal transportation (Crainic and Rousseau 1986, Crainic and Roy 1988); LTL trucking (Roy and Delorme 1989, Powell and Sheffi 1983, 1986, 1989, Powell 1986, Lamar et al. 1990), express courier services (Grünert et al. 1999, Grünert and Sebastian 2000, Büdenbender et al. 2000, Barnhart and Schneur 1996, Kim et al. 1999, Armacost et al. 2002), rail (Crainic et al. 1984, Keaton 1989, 1991, 1992, Newton et al. 1998), and shipping (Christiansen et al. 2004) etc.

Scheduled service network design formulations include an explicit representation of movements of freight in time and usually target the planning of schedules to support decisions related to when services depart from origins and intermediate terminals. A *space-time network* with a scheduling time line is usually used to represent the operations of such scheduled service network systems. The representation of the physical network is replicated at each time point. Temporal arcs then connect the same or different terminals within two time-point representations to represent, respectively, holding activities at the same terminal or actual movements of freight between terminals. The resulting models are similar to those of the static versions but on significantly larger networks due to the time dimension. The additional constraints related to scheduling also contribute to making this class of problems more difficult to solve than static versions. Such formulations have been proposed for, e.g., LTL trucking (Farvolden and Powell 1991, 1994, Farvolden et al. 1993), express courier services (Smilowitz et al. 2003), rail (Haghani 1989, Gorman 1998a,b, Andersen et al. 2009a,b, Pedersen et al. 2009, Zhu et al. 2014) and navigation (Sharypova et al. 2012). Meta-heuristics were proposed in most cases.

Another noteworthy issue is the consideration of resource management. In some applications in the freight transportation industry, the decision maker needs to take into account resource management at the same time as designing the service network, especially when the management, distribution and maintenance of the resources represent a significant part of the total cost. This might include moving empty vehicles, which follows from the imbalances between the freight supply and demand in different regions and points of the systems, resulting in imbalances between vehicle supplies and demands at the terminals. To address these imbalances, empty vehicles must be delivered to terminals where they will be needed to satisfy known or forecasted demand in the following time periods. These repositioning operations usually carry major costs, but are normally dealt with at the operational level of planning, after the network is decided (e.g., Dejax and Crainic 1987, Cordeau et al. 1998, Crainic et al. 1989). Efforts have lately been dedicated to considering asset management requirements, including vehicle repositioning, at the tactical design stage, in order to improve the overall performance of the system (e.g., Pedersen et al. 2009, Andersen et al. 2009a,b, 2011, Lium et al. 2007, 2009, Bai et al. 2014).

Service network design problems have mainly been studied under the assumption that all necessary information, particularly the demand as well as the cost and profit structure, is available before the design decisions are made. It is a general understanding, though, that in most cases, at the time when the transportation plan is made, the demand it will later face is actually uncertain. This is traditionally not explicitly taken into account during the design phase but postponed to be dealt with at the operational phase. Hence, most papers use deterministic models. Demand is usually set to some point forecast of future demand, computed through various forecasting methods or based on historical data (e.g., the "regular" demand of a "normal" week obtained by adjusting last-year's demand with this year's input from the sales department).

Under normal circumstances, the expected quality of a solution derived from a stochastic model is better than its deterministic counterpart when evaluated in the stochastic environment. The reason is that, while it is optimal for one specific scenario, the deterministic solution might be very bad in those scenarios where it is not optimal. See, for example, Wallace (2000) and Higle and Wallace (2003) for discussions. And in most cases, the deterministic design is feasible, but not necessarily optimal in the stochastic model. This badness can be measured by "the Value of the Stochastic Solution" (Birge 1982), or VSS, representing the expected gains obtained from using the stochastic rather than the deterministic solution in the stochastic environment. Previous studies have also shown that by explicitly introducing stochastic demand, the solutions produced can be qualitatively different from those stemming from deterministic models, see for example Wallace (2000). However, there are situations where the VSS is high, meaning that the deterministic solution behaves badly in the stochastic setting, yet the deterministic solution shares some properties with the corresponding stochastic solution. For example, Thapalia et al. (2011, 2012a,b) show that for the single-commodity network design problem, certain structural patterns from the deterministic solutions re-emerge in the stochastic solutions. Similar observations are made in Maggioni and Wallace (2012) for a series of other problems. Lium et al. (2007, 2009) also qualitatively study the structural changes in solutions after introducing uncertainty to a version of the service network design problem, and show that more consolidation is induced by the need to hedge against demand uncertainty. Traditionally, in consolidation-based freight transportation, consolidation is seen as a way to accommodate the fact that most vehicles would not be full with direct deliveries. Lium et al. (2009) show that consolidation, in addition, can achieve higher operational flexibility in a dynamic environment where future demands are unknown, without requiring too much extra services.

3. The Models

In this paper, we consider a stochastic, multi-commodity SSND problem, in which a periodic, cyclic schedule is built for a number of commodities (OD pairs). We present, in this section, four models with fixed and variable capacity, and with or without resource-management constraints. Section 3.1 describes the problem and the notation, while Sections 3.2 and 3.3 present mathematical formulations with fixed and variable capacity, respectively. Section 3.4 discusses the resource-management considerations.

3.1. Problem Setting and Notation

The stochastic SSND problem is set up on a space-time network consisting of nodes and arcs over a given schedule length (e.g., a week), divided into T equal-length *time periods* (e.g., a day) starting at *time points* $t = 0, ..., T - 1$. We denote by $\mathcal{T}$ the set of time points. The schedules are assumed to be cyclic (that is, time period 1 follows time period T) and repetitive for a given planning horizon (e.g., a month or a season) for which the current resource and demand conditions of the system do not change. Such schedule repetition is quite commonly seen in many types of transportation service networks, such as public transit, intercontinental liner shipping and inter-modal rail. See Andersen et al. (2009a,b), Pedersen et al. (2009) and Zhu et al. (2014) for examples of various service network design problems with repetitive schedules.

The nodes in the space-time network stand for terminals at different time points and the arcs represent services for moving commodities between these terminals across time, as well as activities of holding vehicles and freight at a terminal (the holding arcs). Let $\mathcal{N}$ represent the set of terminals. In a space-time network, these terminals are replicated at each time point. We denote by $\mathcal{A}$ the set of arcs between the nodes among which the selection for the final network is to be made. An arc $a = (i, j; t)$ represents the service departing from terminal i at time point t and arriving at terminal j, $\forall i, j \in \mathcal{N}, t = 0, ..., T-1, i \neq j$; or the holding activity at terminal i from time point t to the next if $i = j$. We use l_{ij} to represent the arc length between terminals i and j which can take one or more time periods depending on the physical distance between the two terminals if $i \neq j$, and equals one if $i = j$ for all $i \in \mathcal{N}$. Furthermore, it is assumed that the handling of freight at terminals happens within the time periods, which implies no time delay caused by terminal operations such as unloading, sorting, consolidation and loading activities.

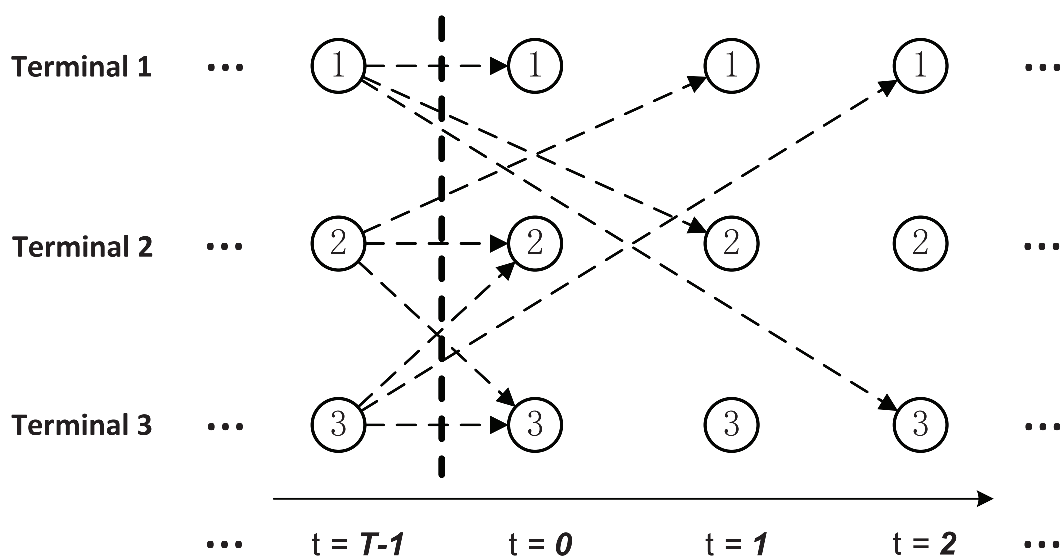

Figure 1 **An example illustrating possible arcs (represented by dashed arrows) that can be set up at $t = T-1$, i.e., the last time point of the repetitive and cyclic schedule.**

Figure 1 shows an example with three terminals and a repetitive and cyclic schedule over T time periods, and illustrates the possible arcs that can be set up at $t = T-1$, i.e., the last time point of the cyclic schedule. For example, a service departing from Terminal-1 at $t = T-1$ takes three time periods and arrives at Terminal-3 at $t = 2$ of the subsequent scheduling cycle. Three holding arcs, joining two representations of the same terminals at two consecutive time points, are also displayed in the example. Note that the cyclic feature of the space-time network is illustrated by letting the services leap over the bold division line in Figure 1.

Let $\mathcal{K}$ be the set of commodities (OD pairs) representing the origin-to-destination demands for transporting a certain quantity of freight between the respective origin and destination terminals within a certain number of time periods. For each $k \in \mathcal{K}$, the transport requirements of commodity k are defined by: $o_k, d_k \in \mathcal{N}$, its origin and destination terminals; $\sigma_k, \tau_k \in \mathcal{T}$, the time point it becomes available at its origin terminal and the time point it must be present at its destination terminal; and its demand, which is described by a continuous distribution. To be able to solve exactly the stochastic problems, the multi-dimensional demand distribution is approximated by a finite set of scenarios $\mathcal{S}$. A probability p^s is assigned to scenario $s \in \mathcal{S}$, with $\sum p^s = 1$. We use δ_k^s to denote the demand for commodity k in scenario s; thus a scenario is $|\mathcal{K}|$-dimensional and contains one demand realization for each commodity.

There is a set up cost $f_{ij;t}$ associated with opening an arc $(i,j;t)$. Also, we need to pay for commodity flows, that is, the transportation and storage of the commodities. Thus cost $e_{ij;t;k}$ associated with arc $(i,j;t)$ represents the unit flow cost incurred to move commodity k or have it wait at the terminal. Additionally, to account for demand not satisfied by the services, we denote by b_k the (usually much more expensive) unit ad hoc handling cost of commodity k whenever part (or all) of its demand cannot be satisfied by regular services. Note that such ad hoc handling can represent, depending on the application, outsourcing the unmet demand to some third-party carrier, delivering it by a different mode (outside the model), or simply represent a penalty for delaying or rejecting the demand. The additional capacity provided by ad hoc handling can be represented by ad hoc arcs that do not carry fixed costs.

The goal is to solve the stochastic optimization problem in order to find a good, if not optimal, solution that represents a periodic schedule which minimizes the expected total system cost. This corresponds to a two-stage structure in the decision process. For a detailed discussion on two-stage settings in modeling with stochastic programming, see Kall and Wallace (1994) and King and Wallace (2012). The first stage decisions, i.e., the selection of services or “the design”, are made before the realization of the random demands, where a fixed cost must be paid whenever a service is selected (set up), representing its make up or maintenance costs. Once these decisions are made, the design is used repeatedly to satisfy the observed realization of random demands. So the second stage is characterized by distributing commodity flows using the selected services with additional capacity described by the ad hoc arcs. The overall objective is thus to minimize the cost of the first stage design plus the *expected* operational and ad hoc handling costs when applying such a design to the demand realizations.

3.2. The Fixed Capacity Model

We first present the formulation with fixed (but not necessarily identical) capacity for every arc. A fixed capacity $h_{ij;t}$ is therefore associated with arc $(i,j;t)$. Let $V_{ij;t} \in \mathcal{A}$ represent the {0,1} arc selection decision variables, and $Y^s_{ij;t;k}$ the flow variables representing the continuous flow of commodity k on arc $(i,j;t)$ in scenario s. Furthermore, let Z^s_k represent the continuous volume of commodity k that uses ad hoc handling in scenario s.

Due to the cyclic nature of the network, the m^{th} time point prior to time t can be denoted as:

$$t \ominus m = (t - m + T) \bmod T \tag{1}$$

The two-stage fixed-capacity formulation of the stochastic SSND problem can then be written as:

$$\min \sum_{i\in\mathcal{N}}\sum_{j\in\mathcal{N}}\sum_{t\in\mathcal{T}} f_{ij;t}V_{ij;t} + \sum_{s\in\mathcal{S}} p^s \left(\sum_{i\in\mathcal{N}}\sum_{j\in\mathcal{N}}\sum_{t\in\mathcal{T}}\sum_{k\in\mathcal{K}} e_{ij;t;k}Y^s_{ij;t;k} + \sum_{k\in\mathcal{K}} b_k Z^s_k \right) \tag{2}$$

$$\sum_{i\in\mathcal{N}} Y^s_{ij;t\ominus l_{ij};k} - \sum_{i\in\mathcal{N}} Y^s_{ji;t;k} = \begin{cases} \delta^s_k - Z^s_k, \text{ if } j = d_k \text{ and } t = \tau_k \\ -\delta^s_k + Z^s_k, \text{ if } j = o_k \text{ and } t = \sigma_k \\ 0, \text{ other} \end{cases} \tag{3}$$

$$\forall j \in \mathcal{N}, \forall t \in \mathcal{T}, \forall k \in \mathcal{K}, \forall s \in \mathcal{S}$$

$$\sum_{k\in\mathcal{K}} Y^s_{ij;t;k} \le h_{ij;t}V_{ij;t} \qquad \forall i,j \in \mathcal{N}, \forall t \in \mathcal{T}, \forall s \in \mathcal{S} \tag{4}$$

$$V_{ij;t} \in \{0,1\} \qquad \forall i,j \in \mathcal{N}, \forall t \in \mathcal{T} \tag{5}$$

$$0 \le Z^s_k \le \delta^s_k \qquad \forall k \in \mathcal{K}, \forall s \in \mathcal{S} \tag{6}$$

$$0 \le Y^s_{ij;t;k} \le \delta^s_k \qquad \forall i,j \in \mathcal{N}, \forall t \in \mathcal{T}, \forall k \in \mathcal{K}, \forall s \in \mathcal{S} \tag{7}$$

The objective function (2) minimizes the costs for opening services plus the expected costs for moving and holding commodities, as well as using ad hoc capacity. Constraints (3) represent the conservation of flow for commodities. Constraints (4) make sure the total flow on an arc does not exceed its capacity. Constraints (5) impose the integrality requirements on the design variables. Constraints (6) limit the use of ad hoc capacity to the observed actual scenario demand, while constraints (7) limit the flow of a commodity to its corresponding demand on all arcs.

3.3. The Variable Capacity Model

We also introduce a slightly different model with variable capacity, where every service, when selected (opened), has a *maximum* capacity, $h_{ij;t}, (i,j;t) \in \mathcal{A}$, limiting the number of flow-carrying units it may haul. This concerns, e.g., rail cars making up a block or train, trailers in a multi-trailer trucking service, or barges in a barge-train. For simplicity's sake, we assume all units making up a service have equal capacity, $u_{ij;t}$. The cost of adding one unit of service capacity is represented by $c_{ij;t}, (i,j;t) \in \mathcal{A}$. To our best knowledge, this problem setting has not been studied before.

We define the integer decision variables $X_{ij;t}$ to represent the number of units of capacity provided on arc $(i,j;t) \in \mathcal{A}$. The other decision variables are the same as in the fixed-capacity model capturing the service selection choices, indicating whether the service is selected and leaves at the specified time point ($V_{ij;t}$), the continuous flow of commodity k on arc $(i,j;t)$ in scenario s ($Y^s_{ij;t;k}$), and the continuous volume of commodity k that uses ad hoc handling in scenario s (Z^s_k). The formulation then becomes:

$$\begin{aligned} \min \sum_{i\in\mathcal{N}}\sum_{j\in\mathcal{N}}\sum_{t\in\mathcal{T}} f_{ij;t}V_{ij;t} + \sum_{i\in\mathcal{N}}\sum_{j\in\mathcal{N}}\sum_{t\in\mathcal{T}} c_{ij;t}X_{ij;t} \\ + \sum_{s\in\mathcal{S}} p^s\Big(\sum_{i\in\mathcal{N}}\sum_{j\in\mathcal{N}}\sum_{t\in\mathcal{T}}\sum_{k\in\mathcal{K}} e_{ij;t;k}Y^s_{ij;t;k} + \sum_{k\in\mathcal{K}} b_k Z^s_k\Big) \end{aligned} \tag{8}$$

$$\sum_{i\in\mathcal{N}} Y^s_{ij;t\ominus l_{ij};k} - \sum_{i\in\mathcal{N}} Y^s_{ji;t;k} = \begin{cases} \delta^s_k - Z^s_k, \text{ if } j = d_k \text{ and } t = \tau_k \\ -\delta^s_k + Z^s_k, \text{ if } j = o_k \text{ and } t = \sigma_k \\ 0, \text{ other} \end{cases} \tag{9}$$

$$\forall j \in \mathcal{N}, \forall t \in \mathcal{T}, \forall k \in \mathcal{K}, \forall s \in \mathcal{S}$$

$$\sum_{k\in\mathcal{K}} Y^s_{ij;t;k} \le u_{ij;t} X_{ij;t} \qquad \forall i,j\in\mathcal{N}, \forall t\in\mathcal{T}, \forall s\in\mathcal{S} \tag{10}$$

$$0 \le X_{ij;t} \le h_{ij;t} V_{ij;t} \qquad \forall i,j\in\mathcal{N}, \forall t\in\mathcal{T} \tag{11}$$

$$V_{ij;t} \in \{0,1\} \qquad \forall i,j\in\mathcal{N}, \forall t\in\mathcal{T} \tag{12}$$

$$X_{ij;t} \in \mathbb{Z}_+ \qquad \forall i,j\in\mathcal{N}, \forall t\in\mathcal{T} \tag{13}$$

$$0 \le Z^s_k \le \delta^s_k \qquad \forall k\in\mathcal{K}, \forall s\in\mathcal{S} \tag{14}$$

$$0 \le Y^s_{ij;t;k} \le \delta^s_k \qquad \forall i,j\in\mathcal{N}, \forall t\in\mathcal{T}, \forall k\in\mathcal{K}, \forall s\in\mathcal{S} \tag{15}$$

The objective function (8) minimizes the total cost for offering services and providing service capacities, plus the expected cost for moving or holding commodities and using ad hoc handling. Constraints (4) in the fixed capacity model were replaced by constraints (10) and (11). Constraints (10) make sure the total flow on each arc does not exceed the provided capacity, which is now also a decision variable. Constraints (11) ensure the maximum number of units of capacity on each service is respected.

3.4. Resource Management Considerations

The management of resources used in the transport operations, such as the movement of vehicles, power units and crews, is traditionally not explicitly included in service network design models. Although the management, distribution and maintenance of the relevant assets are in their own right important, in the literature, they are usually dealt with in a separate problem. Furthermore, in some applications, such resources are primarily acquired from external providers, and their reallocations are therefore irrelevant, e.g., oil/gas companies that plan regular services to ship the products to their customers using external ships on time or voyage charters.

The simultaneous determination of service networks and resource movements is, however, receiving more attention in the literature as it can lead to more efficient utilization of the resources. We therefore introduce asset-balance constraints for the fixed and variable-capacity models.

The asset-balance requirements in the fixed-capacity case take the form

$$\sum_{i\in\mathcal{N}} V_{ij;t\ominus l_{ij}} = \sum_{i\in\mathcal{N}} V_{ji;t} \qquad \forall j\in\mathcal{N}, \forall t\in\mathcal{T}, \tag{16}$$

where one assumes a single unit of resource required to operate each service (as used, e.g., in Lium et al. 2007, 2009).

When the assets controlled correspond to the number of services (e.g., power units, ships, etc.), equation (16) may be also used within the variable-capacity formulation. When, on the other hand, the controlled assets are the units of capacity, the constraints have to be written in the appropriate units as in equation (17) where, for simplicity of presentation we assume all units are the same for all services.

$$\sum_{i\in\mathcal{N}} X_{ij;t\ominus l_{ij}} = \sum_{i\in\mathcal{N}} X_{ji;t} \qquad \forall j\in\mathcal{N}, \forall t\in\mathcal{T} \tag{17}$$

4. Comparison Tests

In this section we propose a set of comparison tests, based on solving the deterministic versions of the presented stochastic models, with the intention of investigating the value of the deterministic solution in a stochastic setting. Note that the stochastic formulations may easily be transformed into their respective deterministic versions, by inputting only one scenario. In particular, we start by evaluating the quality of the deterministic solution in the stochastic environment. We then seek to construct other solutions, using parts of the deterministic solution, to see if the performance can be improved.

We first propose two comparison tests inspired by Thapalia et al. (2012b) that apply to both fixed and variable capacity models. Each comparison test can be seen as a particular means to use the information obtained from solving the deterministic version, and to arrive at a solution to the stochastic problem. The results of these comparison tests, i.e., the performances of the corresponding solutions, can then be benchmarked against the result of solving exactly the stochastic problem (i.e., the performance of the optimal solution). These two comparison tests are as follows:

1. Deterministic design used in the stochastic model: (**Determ**)

This is the standard VSS evaluation (Birge 1982), where the deterministic solution is directly carried over to the stochastic model and the quality of the deterministic solution is evaluated. When applied to fixed and variable capacity models, this comparison test is similar but with some differences:

(a) For the fixed capacity models, we first solve the deterministic version of the problem and observe which arcs are open. We keep these arcs open and close all other arcs in the network, i.e., fix the first stage decision variables $V_{ij;t}$. We then run the remaining LP (linear program) to obtain the flow variables of the stochastic model.

(b) For the variable capacity models, the process is similar except that the first stage design consists of $V_{ij;t}$ as well as the capacities offered on these opened services $X_{ij;t}$. We therefore fix both $V_{ij;t}$ and $X_{ij;t}$, and follow up by solving an LP to set the flow variables of the stochastic model.

2. Deterministic design with extra services and capacities: (**Upgrade**)

When applying this comparison test to fixed and variable capacity models, the differences are:

(a) For the fixed capacity models, we keep those arcs obtained from the deterministic solution open, but do not close other arcs in the network. We then run the stochastic problem again to allow new arcs to be set up in addition to those opened in the deterministic solution.

(b) For the variable capacity models, we view both arcs set up and capacities provided on these arcs as "invested" and they cannot be undone. However, we still allow more capacity to be offered on these selected arcs, as long as their corresponding maximum capacities are respected. Additionally, we also allow extra arcs to be set up apart from the selected arcs, as in the fixed-capacity case.

The tests of **Determ** and **Upgrade** are performed to check the immediate performance and upgradability, respectively, of the deterministic solution in the stochastic setting. Using the **Upgrade** test we try to gain insights with regards to the problem we are investigating: can the deterministic solution be upgraded (by extra investments) to a reasonably good solution in the stochastic environment, or are we already lost after implementing the deterministic solution? Both conclusions are possible as demonstrated by Maggioni and Wallace (2012) for some other types of stochastic problems.

For the variable capacity model, we propose an extra test to try to upgrade the deterministic solution in a different way. Different from the fixed capacity model, the information of a particular deterministic solution to the variable capacity model contains two components: the service selection information; and the capacities provided on these selected services. We use the term *skeleton* to represent the former component, and therefore the deterministic skeleton and the corresponding capacities installed on such a skeleton make up the full deterministic design decisions. We thus propose the following test, extracting only the skeleton of a deterministic design:

3. Deterministic skeleton with capacities set by the variable-capacity stochastic model: (**Skeleton**)

We start by solving the deterministic problem and observe which arcs are open. We then only fix the service selection variables $V_{ij;t}$ by keeping these arcs open, and run a MIP (mixed integer program) to set the capacities $X_{ij;t}$ and the flows.

It is important to notice that the **Skeleton** tests do not allow extra arcs to be opened (as tests **Upgrade** do), but only allow the capacities on the deterministic skeleton to be set using the stochastic model. The **Upgrade** is therefore not an efficient numerical procedure as the second step is to solve a stochastic program of the same complexity as the original stochastic SSND problem. However, the **Skeleton** is operated on a reduced network with all other arcs (apart from the ones on the deterministic skeleton) closed. It is particularly interesting when capacities are continuous (discussed in Section 5.5), as the stochastic program used to set capacities will then be a stochastic LP instead of a stochastic MIP.

5. Computational Study

To perform the computational study, we use two sets of problem instances constructed following a random-based procedure: a *Small* set of relatively small-size instances and a *Larger* set in which problems are larger. We introduce these instances in Section 5.1.

The *Small* set of instances is such that most stochastic problems, especially in the variable capacity cases, can be solved within 2 hours using CPLEX 12.6.1 without terminating with large optimality gaps (the gap is at most 5.66%, which occurs in solving one stochastic problem with variable integer capacity). The performances of other approaches, such as **Determ**, **Skeleton** and **Upgrade**, are therefore well benchmarked for these *Small* instances. Sections 5.2 and 5.3 report these results and discuss the value of deterministic solutions for the fixed- and variable-capacity models, respectively.

Section 5.4 uses a quantitative approach to examine the structural improvements that take place when upgrading the deterministic solutions. Section 5.5 summarizes the results for the *Small* set, including the cases of variable continuous capacities and models with asset-balance constraints.

In Section 5.6, similar comparison tests are performed on the *Larger* set of instances, while increasing the maximal run time to 5 hours. Using instances of larger and various sizes, we seek to further check the consistency of the insights drawn from the previous *Small* set of instances, and to investigate the potential usage of deterministic solution based methods on larger problems where stochastic programs are not numerically solvable.

5.1. Test Instances

Table 1 Summary of all problem instances

Set	#	$\lvert\mathcal{N}\rvert$	$\lvert\mathcal{K}\rvert$	$\lvert\mathcal{T}\rvert$	$\lvert\mathcal{S}\rvert$	Demand Distribution
Small	10	6	16	7	30	symmetrical/randomly skewed
Larger	24	6/10	16/20	7/9	40/60	symmetrical/randomly skewed

We first summarize the parameter settings for the two instance sets in Table 1: the *Small* set consists of 10 instances with 6 terminals, 16 commodities, 7 time points and 30 scenarios; and the *Larger* set consists of 24 instances with 6 or 10 terminals, 16 or 20 commodities, 7 or 9 time periods and 40 or 60 scenarios (a complete list of the *Larger* instances can be found later in Table 4). In the following, we start with the random-based procedure that is used to construct the instances. We then introduce and discuss our scenario generation process: the distributions of the stochastic demands (symmetrical or skewed), their correlations and the appropriateness of the number of scenarios chosen for our problem instances.

Given the numbers of terminals, commodities and time points, we start with randomly generating values for the coordinates of all the terminals, evenly spread inside a square-shaped area. Direct services are allowed between any two terminals, which indicates a potentially complete service network. The service length between two different terminals is decided according to their physical distance, such that for any $i, j \in \mathcal{N}$ and $i \neq j$, service length l_{ij} is assigned an integer value: 1, 2 or 3 in our experiments. We assume the storage of a commodity is always possible at no cost in any terminal, i.e., all holding arcs in the space-time network already exist ($V_{ij;t} = 1$ for all $i, j \in \mathcal{N}$ and $i = j$) and are uncapacitated. For service (non-holding) arcs in a fixed capacity model, their (fixed) capacities are all set to 6. In the variable capacity models, the maximal capacity $h_{ij;t}$ of every service arc is also set to 6, while one unit capacity $u_{ij;t}$ is equal to 1. This setting also means that we assume there are six levels to which the capacity of a service may be set, where each level represents around 16% of the total capacity.

The values of the unit flow costs $e_{ij;t;k}$ and of the unit ad hoc handling costs b_k, associated with commodity k, are set proportional to the distance between the terminals, with the latter being eight times as high as the former. This is firstly because ad hoc capacities do not incur set-up costs. Secondly, the ad hoc handling costs are carefully set such that the following two undesirable extreme situations are avoided: if the ad hoc handling is too expensive, it is as if we do not have ad hoc handling in the first place, and most likely too much capacity is installed; on the other hand, if the ad hoc handling is too cheap, we may start to replace regular services with ad hoc capacity, which contradicts the whole idea of the model. This is discussed further in Section 5.2.

For every commodity, its origin and destination terminals are both selected randomly. A commodity's time span (from the time point it becomes available to the time point it has to be delivered) ranges from 2 to 5, and is not shorter than the service length between the associated terminals. The stochastic demand of a commodity is either subject to a *symmetrical* triangular distribution or a skewed triangular distribution with *random skewness*, i.e., it is randomly selected to be left- or right-skewed. In the symmetrical case, the distribution has a lower limit 0, mode (peak value) $b/2$ and upper limit b (the resulting coefficient of variation is 0.4), where b is set to 4 in our instances. Recall that the maximal capacity of a service is set to 6, hence allowing three commodities with average demand to pass at the same time. Also, the demand of one commodity will never fill a service. In the skewed case, the lower and upper limits are always 0 and b, respectively, but the mode is set at random to $b/4$ or $3b/4$. When constructing the instances, we let half of all instances have demands that are all symmetrically distributed (common in many applications, such as less-than-truckload trucking), and the other half have demands with random skewness. More specifically, for the *Small* set, in five of the instances all commodities have symmetrical distributions, whereas in the other five instances each commodity has an either left- or right-skewed demand distribution. For the *Larger* set, this is indicated with s or r when referring to a specific instance later in this paper (see Table 4 for examples), which stand for *symmetrical* and *random skewness*, respectively. Correlation matrices are generated randomly, of course making sure that they are indeed positive semi-definite, a necessary condition for a correlation matrix.

We discretize the demand distributions by generating scenarios to represent the stochasticity. The scenario generation process is performed using the moment-matching method introduced by Høyland et al. (2003). This method takes as input the first four marginal moments and correlations of the random variables, and generates (if there are enough scenarios – and in our case there are) scenarios that replicate all these properties. The scenarios have equal probabilities, i.e., $p^s = 1/\mid \mathcal{S} \mid$ for all $s \in \mathcal{S}$.

The more scenarios, the better the representation of the demand distribution. But as we increase the number of scenarios, the difficulty of obtaining an optimal solution increases as well. Our scenario generation method is stochastic (the way sampling is stochastic, though we are not sampling). Hence, if the procedure is rerun with the same input of marginal distributions and correlations, the resulting scenario tree will be different. To ensure that the results are not much affected by the particular scenario trees used, we check the *in-sample* and *out-of-sample stability* (Kaut and Wallace 2007). We generate 20 scenario trees of the same size and solve the stochastic problem for each of them. In the in-sample stability test, for 30 scenarios, we observe a difference between the highest and lowest objective function values of 3.58%. To test for out-of-sample stability, we sample a much larger scenario tree with 5000 scenarios to represent the uncertainty of the "true" problem. We calculate the "true" objective function values corresponding to the 20 solutions (design decisions) coming from the different (smaller) scenario trees, and observe a difference of 3.26% between the highest and lowest. In addition, the average in- and out-of-sample values are only 0.37% apart. These stability calculations yield acceptable difference values and confirm that our model using 30 scenarios is in- and out-of-sample stable (at about a 4% level).

5.2. Value of the Deterministic Solution - Fixed Capacity Model

Based on the fixed capacity model presented in Section 3.2, the first two comparison tests (**Determ** and **Upgrade**) are performed for all *Small* instances and benchmarked against solving the stochastic problems exactly. The aggregated results are shown in Figure 2. The bars show the losses produced by the **Determ** and **Upgrade** tests, relative to the optimal solutions of the stochastic program. The ***Min.Loss*** and ***Max.Loss*** indicate the best and worst cases for the two tests across the instances. The ***Avg.Loss*** columns show the mean losses.

From Figure 2 we see that although losses can go as high as 60%, the **Determ** approach produces an average loss of around 17%, which is rather small compared with some other stochastic network design problems reported in the literature (see Thapalia et al. 2012a, Maggioni and Wallace 2012). But more importantly, with extra services, the deterministic solutions can be greatly improved as suggested by the dramatic decrease in all three losses

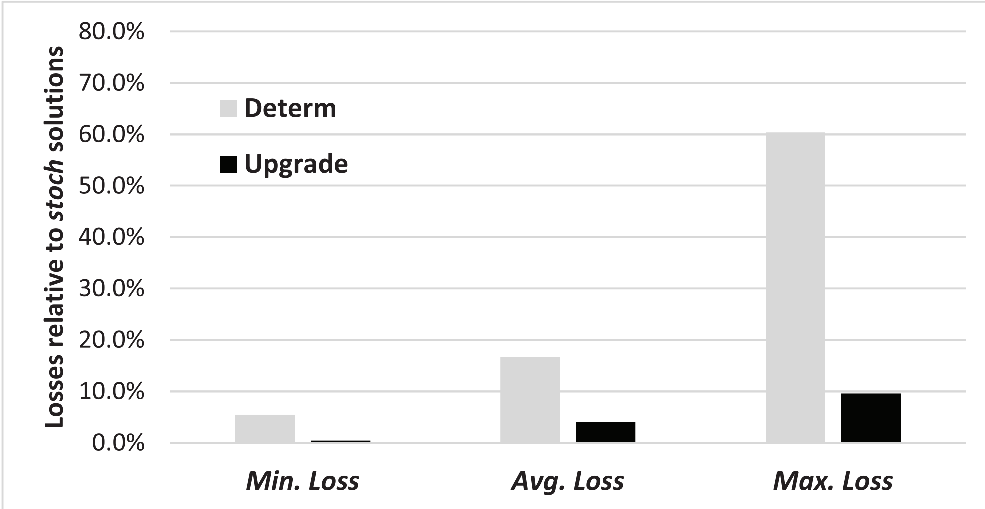

Figure 2 **Comparison of the Determ and Upgrade tests in the fixed capacity model. Results are measured by minimum, average and maximum losses relative to the stochastic (optimal) solution.**

from **Determ** to **Upgrade**. In other words, it shows that adding extra services to the deterministic design is beneficial and effective in most circumstances (loss is under 10% even for the worst case).

In the tests we characterize demand stochasticity for each commodity using a (symmetrical or skewed) triangular distribution, which is replaced by its 50^{th} percentile in the deterministic case. In those scenarios where some demands cannot be satisfied with the deterministic design, one must use the expensive ad hoc capacity which translates into the losses reflected in Figure 2 for **Determ**: about 17% on average and 60% at the highest. However, if the deterministic design is allowed to be expanded with extra services, these unmet demands may use the relatively cheaper extra services instead of ad hoc capacities, hence the lower losses for **Upgrade**.

The losses from using the deterministic designs in the stochastic environment are primarily caused by insufficient capacities. We therefore test deterministic designs produced using the 75^{th} percentile of the demand distributions. This is common practice in many industries. In this case, the average loss from using the deterministic solution in the stochastic environment (test **Determ**) drops from 16.63% to 8.86%. For the **Upgrade** tests, an average loss of 2.16% is observed, which is also an improvement compared to the 50^{th} percentile counterpart, which had an average loss of 4.03%. The detailed results are summarized and reported in Table 2 in Section 5.5.

Notice that the absolute values of these losses are affected by the cost structure of the problem, especially the setting of ad hoc handling costs relative to transportation costs. This is a problem shared with all stochastic programs using soft constraints, and is unavoidable unless the penalties are dictated by the problem at hand – something they rarely are. In this paper, therefore, we do not focus on the absolute values of the losses, and no general statements should be drawn from these values with regards to the absolute "badness" of certain designs in a stochastic environment. Instead, we focus more on the comparison of relative performances, for example when comparing forecasts using the 50^{th} or 75^{th} percentile of the distribution in the deterministic model.

Mathematically speaking, the difficulty of performing the **Upgrade** test is on par with solving the original stochastic problem to optimality. The actual difficulty of course depends on the specific instance. But on a complete service network, its complexity is not reduced much by fixing a relatively small number of $\{0,1\}$ decision variables, as it is still a large MIP when all the other $\{0,1\}$ decisions are still to be determined. However, this test is primarily included for model understanding, not numerical efficiency. The fact that the deterministic design can be upgraded into a very good solution shows that the investments in the deterministic design are not wasted. In a highly dynamic transportation industry, it means that decision makers can sometimes safely invest into some services well ahead of time, especially if a discount is applicable by doing so. This is also a good way to reduce risks when the costs of setting up services are highly uncertain in the future, and especially when such costs are likely to go up closer to the time when one has to make the final plan. If the investment period is long, it is safe to start by setting up the services from the deterministic design as they can be expanded later. Also note that if the **Upgrade** test shows good results, one may develop constructive heuristcs based on the deterministic solution rather than starting from scratch.

Similar observations are made with the expected value approach for some other types of problems; we refer the interested readers to Maggioni and Wallace (2012) for more details. Note that it is not at all obvious that deterministic solutions are upgradable, which is also illustrated by Maggioni and Wallace and which underlines the interest of our results.

5.3. Value of Deterministic Solution - Variable Capacity Model

For the model with variable capacity, it is also possible to perform the **Skeleton** test in addition to the ones performed in the fixed capacity case. Remember that for such a test, we start by solving the deterministic problem; fix the "skeleton" variables $V_{ij;t}$ only; then solve the remaining MIP to set the capacities $X_{ij;t}$ and distribute the flows $Y^s_{ij;t;k}$.

The same ten *Small* instances are used to perform the tests here. We compare the deterministic solution (**Determ**), the deterministic skeleton with updated capacities (**Skeleton**) and the deterministic design with extra services (**Upgrade**). Note again that in the **Upgrade** test for the variable capacity model, we see both services set up and capacities provided on these services as "invested". We still allow more capacity to be offered on these selected services, however, as long as their corresponding capacity limits are respected. We also allow extra services to be set up apart from the selected services.

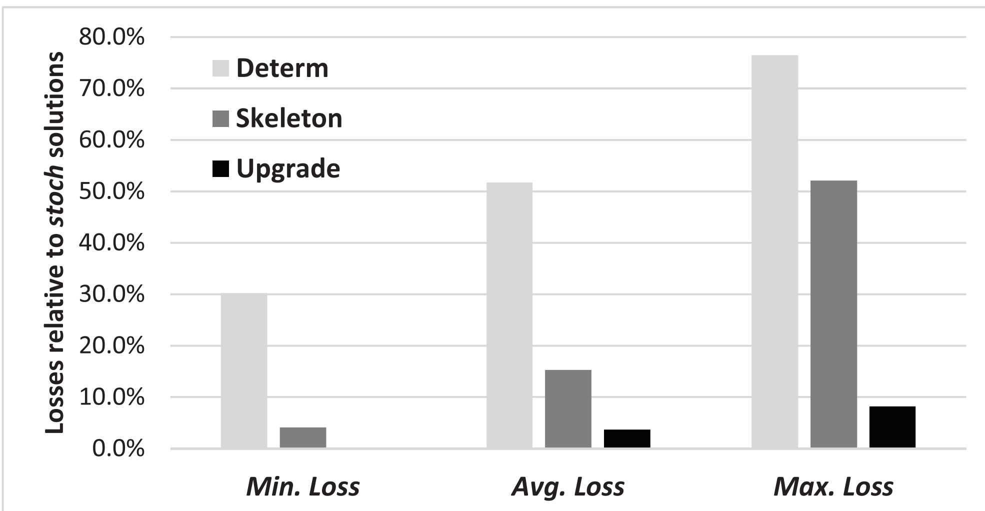

Figure 3 **Comparison of Determ, Skeleton and Upgrade designs for the variable capacity model. Results are measured by minimum, average and maximum losses relative to the stochastic (optimal) solution.**

Figure 3 shows that the deterministic solution (**Determ**) is quite bad in the stochastic setting, while **Skeleton** behaves much better. Although the maximum loss of **Skeleton** is still high (around 52%), its average loss (around 15%) is quite acceptable. On the other hand, in the **Determ** tests, the average loss goes over 50%, and even the minimum loss is around 30%. Again, the **Upgrade** results show that deterministic solutions can be upgraded to very good solutions for the stochastic models. Again, this is useful information about the models, even though, numerically, this is not an efficient procedure.

These results show that, in general, the deterministic solution does not handle demand uncertainty well when capacities are not fixed a priori. Using the skeleton, however, is beneficial in most circumstances. This is very well illustrated by comparing the **Determ** performances in Figures 2 and 3. It may be explained by the possibility in the variable-capacity model to closely adjust the supplied capacity to demand. This capability is very useful for a deterministic setting but not when evaluating the deterministic solution in a stochastic setting. Indeed, adjusting the capacity to the estimated demand results in little extra capacity available when the observed demand is higher than the prediction, which comes at the price of much ad hoc capacity. The results reported later in Table 2 in Section 5.5 are extremely telling in this context, the performance of the **Determ** approach improving dramatically (more than fourfold) when the 75^{th} percentile of the demand distribution is used as forecast. The performance of the skeleton-based solution is still better, but the two approaches are more at par in the 75^{th} situation, as the improvement of **Skeleton** is smaller. Notice that the last observation points to the fact that the **Determ** 75^{th} approach could be more "forgiving" of a bad demand estimation. On the other hand, the performance of **Upgrade** is fundamentally constant, with small differences between 50^{th} and 75^{th} percentile cases. Also notice that the relative values of fixed set up costs $f_{ij;t}$ and capacity costs $c_{ij;t}$ are important drivers in the differences between **Determ** and **Skeleton** in Figure 3. Consider an extreme situation where all $c_{ij;t}$ are zero, implying that capacities on the opened services are provided at no charge. In that case, the variable-capacity model can, for all practical purposes, be considered a fixed-capacity model, and the **Skeleton** approach is identical to **Determ** in the fixed-capacity case.

5.4. Structural Differences

We now investigate the structural differences between solutions obtained in the different tests and try to find out if any systematic "structural improvements" can be observed as the solutions get better. This is in order to gain insights into why and how the deterministic solutions may be upgraded and therefore of value to the stochastic problems.

Inspired by the beneficial features discovered by Lium et al. (2009) where, in particular, more hub-and-spoke structures were observed after demand stochasticity was explicitly considered, we have also observed in our experiments more consolidation activities in stochastic solutions than in their deterministic counterparts. Therefore, if the level of consolidation in any particular solution can be measured quantitatively, we may find a correlation between the solution's consolidation level and its performance (in terms of total costs) in the stochastic setting. For example, we may compare these two measures (consolidation level and performance) of the **Determ**, **Skeleton**, **Upgrade** and optimal solutions, and see if some trend shows up.

However, to precisely define the “the level of consolidation” allowed by a design is difficult. We therefore propose a scheme to measure two substitute phenomena: the levels of *multi-path usage* and *path-sharing* when the design is used in a stochastic environment. If more commodities are using multiple paths to reach their respective destinations, and more services in the network are shared by several commodities, then, most likely, more consolidation activities have taken place.

We start by measuring the levels of multi-path usage. For a given solution, we record the highest number of paths each commodity uses *across all scenarios* and then produce a histogram to display the frequencies (in terms of number of commodities) *with all instances added up*. For example, we say that a commodity uses two paths if this commodity travels on, at the most, two paths in at least one of the scenarios; and if we have ten instances, each with 16 commodities, we count this as 160 commodities in the statistics. We then count how many of these commodities travel on one, two, three, and so on, paths.

Figure 4 presents the level of multi-path usage measured by commodity counts, in the **Determ**, **Upgrade** and **Stoch** (optimal stochastic solution) cases, for the fixed capacity model. In the **Determ** case, i.e., Figure 4a, there are 152 commodities using only one path, seven commodities using two paths, and one commodity using three paths to reach their respective destinations. In the **Stoch** case, i.e., Figure 4c, the number of commodities using two paths rises to 39, and there are 14 commodities using three paths and even two commodities using four paths while the number for a single path has dropped from 152 to 105. For the **Upgrade** case, as shown in Figure 4b, we can also see a significant increase in the number of commodities using multiple paths compared to the **Determ** case, yet lower compared to the **Stoch** case.

If we study the levels of multiple-path usage for **Determ**, **Upgrade** and **Stoch** and compare those with their respective performances we see a trend (see Figure 2 for the performances of **Determ** and **Upgrade**; **Stoch**, as the optimal solution, will of course produce 0% losses). That is, the better the solution performs the higher the level of multiple-path usage it has. Considering the great improvement in performance from **Determ** to **Upgrade**, this also indicates that with some new arcs opened, the deterministic design

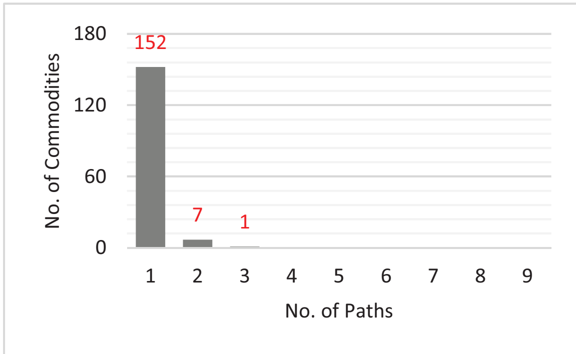

(a) Determ

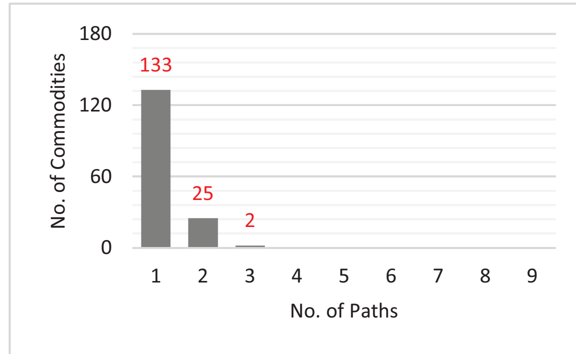

(b) Upgrade

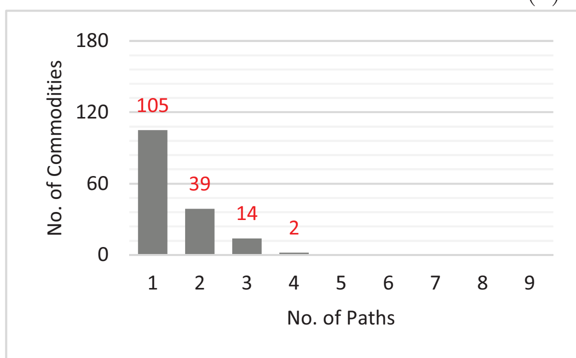

(c) Stoch

Figure 4 The level of multi-path usage for the fixed capacity model.

is able to evolve to a structurally different design that allows a higher level of multi-path usage and becomes very competitive for the stochastic problem.

Similar insights can be drawn when measuring the levels of path-sharing. We do this by counting the number of commodities routed through each opened service (not including the holding arcs). Note that a commodity may be routed though a number of services to reach its destination. We thus say that if two commodities have at least one service in common, they are sharing paths.

Path-sharing measurements for the fixed capacity model are displayed in Figure 5. In the **Determ** case, 141 arcs are shared by two commodities and 46 arcs are shared by three commodities. These two counts increase to 148 and 78 in the **Upgrade** case. In the **Stoch** case, the number of arcs shared by two commodities stays at a similar level (136) while the number of arcs shared by three commodities increases further to 87, and the number of arcs shared by four commodities reaches 60. In general, we can see a right shift of the frequency curve, from **Determ** to **Upgrade** and then to **Stoch**, while the performance

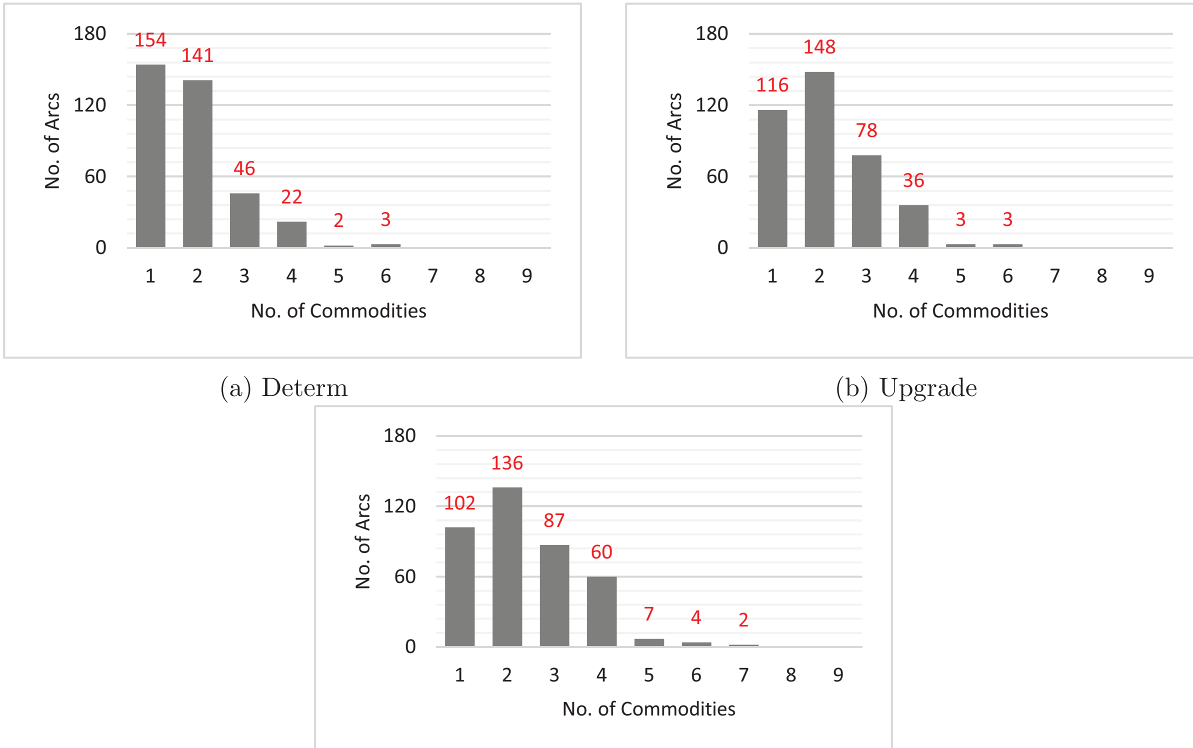

(a) Determ (b) Upgrade

(c) Stoch

Figure 5 The level of path-sharing for the fixed capacity model.

of the corresponding solutions improve in the stochastic environment, indicating that the better the solution performs the higher level of path-sharing it has.

The above results confirm two structural features for the fixed capacity model: it is potentially beneficial to have a design structure that allows high levels of multi-path usage and path-sharing. Furthermore, with some extra services, the deterministic solution can be structurally changed in terms of its potential to allow higher levels of these two phenomena, and become much better suited to handle the stochastic demands. So how many extra services are required to make the change?

First of all, our results show that 70%-90% of the arcs selected by the deterministic solutions reappear in the corresponding stochastic solutions. It means that the stochastic solution shares with its deterministic counterpart most of the service selection decisions, but the stochastic solution also includes some unique arcs to obtain a structure with much higher flexibility to handle demand variations through higher levels of multi-path usage and path-sharing.

Our numbers also show that, on average, around 15% extra arcs are added to the deterministic design in the **Upgrade** test. Therefore, by adding a limited number of extra arcs, the deterministic design can become structurally different, and much better suited for the stochastic environment. So what can we do to find the right extra arcs? As mentioned earlier, on a complete network, the difficulty of finding these extra arcs numerically are on par with solving the original stochastic program. However, this idea of trying to find the "correct" extra arcs based on the deterministic solution can be used in the development of a heuristic approach, for example by targeting those arcs that are likely to increase the levels of multi-path usage and path-sharing in the network.

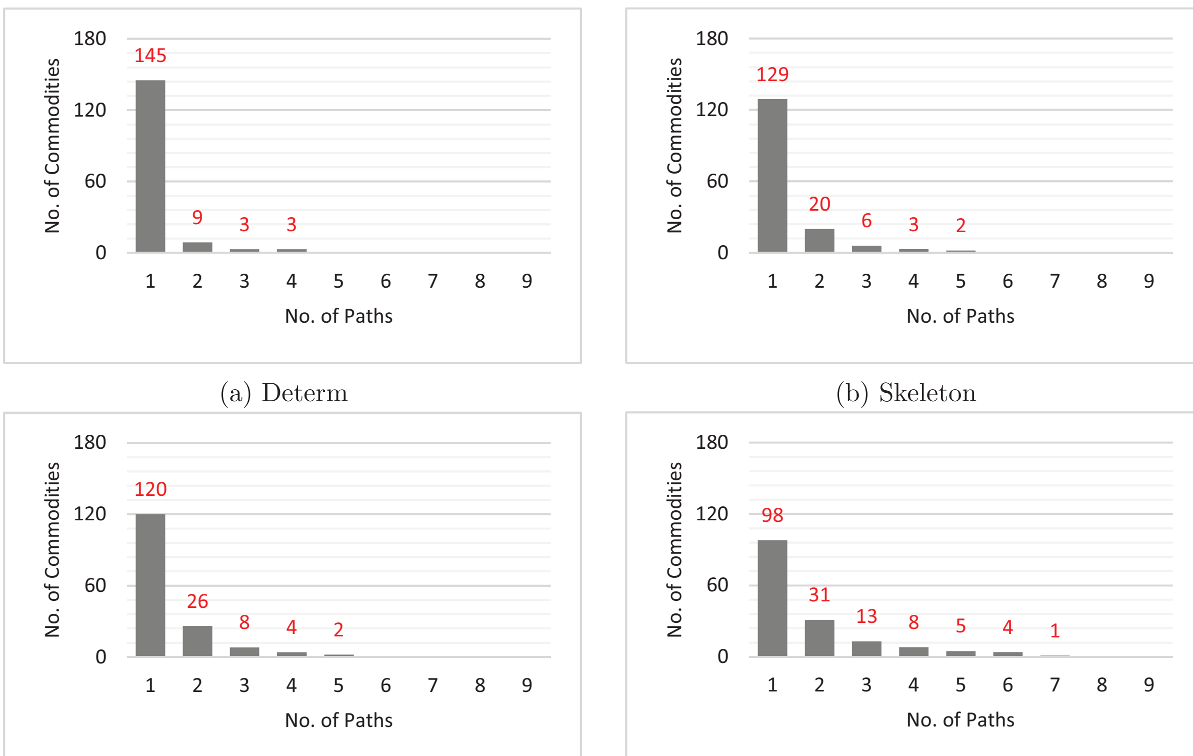

(a) Determ (b) Skeleton

(c) Upgrade (d) Stoch

Figure 6 The level of multi-path usage for the variable capacity model.

Following a similar thinking, we also apply the measurement scheme to the variable capacity model, with the extra test **Skeleton**. Similar conclusions may be drawn from the results displayed in Figures 6 and 7: the better the solution performs, the higher levels of multiple-path usage and path-sharing it has. This is clearly visible from the charts and the numbers.

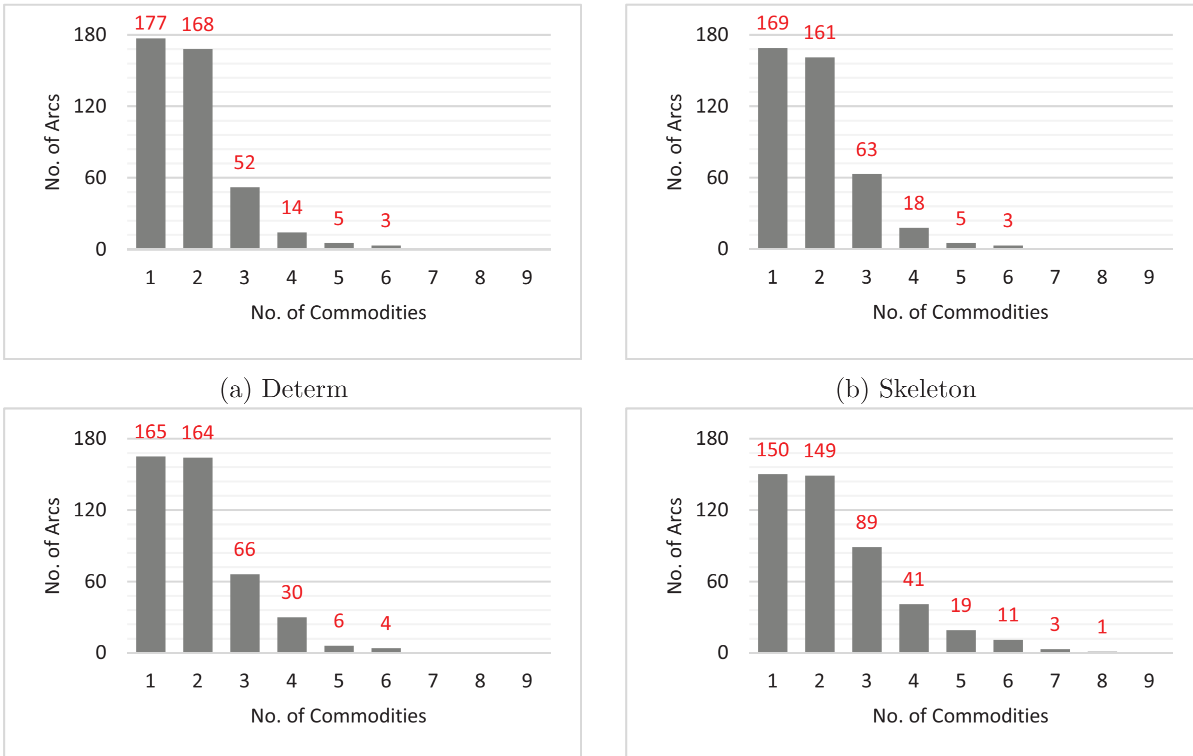

(a) Determ (b) Skeleton

(c) Upgrade (d) Stoch

Figure 7 The level of path-sharing for the variable capacity model.

But if we consider the changes from **Determ** to **Skeleton** (they have the same service selection decisions, but provide different capacities), we can see some interesting similarities, in contrast to the updates from **Determ** to **Upgrade** in the fixed capacity model. Rather than allowing other services to be opened, **Skeleton** merely changes the capacities provided on the already selected services. It still brings out similar *structural improvements*, allowing higher levels of multi-path usage and path-sharing. We conclude that a design based on the deterministic skeleton is able to adapt itself structurally to uncertainty even when its options are highly limited.

In **Skeleton**, capacities are only allowed on the deterministic skeleton. So, essentially, the original complete network is “shrunk” to a smaller skeleton network. Then, the possibility of finding a new path for a given commodity depends on whether there is another combination of services (apart from the deterministic one) on the reduced network to take it from its origin to its destination. If yes, then in those scenarios where the commodity’s demand is very high, it might use the new path as long as there is free capacity. In Figure 6, we see the number of commodities using two paths more than doubled, from nine to 20.

A noteworthy observation is that while the better the solution performs, the higher the levels of multiple-path usage and path-sharing are; the reverse is not always true. There is an obvious counter-example. If we enforce very tight capacity limits on all possible services, we can obtain a solution with an extremely high level of multiple-path usage and path-sharing, as all the commodities would have to find many paths trying to avoid expensive ad hoc handling. This might result in opening a large number of services, and very poor performance.

5.5. Summarizing the Comparison Tests on the *Small* set

The **Determ**, **Skeleton** and **Upgrade** approaches represent different ways of using the information obtained from the deterministic solution in the stochastic environment. In this section, based on the ten instances in the *Small* set, we compare the average loss of these three approaches applied on all models introduced in this paper, including those with and without asset balance considerations.

In addition, we include results for the variable capacity models, where we relax the integrality constraints on the capacities, i.e., constraints (13) in Section 3.3. This may correspond to an approximation of actual integer capacities (could be appropriate when capacities are large and their integrality is less important) or to applications in some fields where capacities are actually continuous, e.g., in bulk shipping and railways where the capacity is often in meters/feet or tons. The computational effort to perform the **Skeleton** test is much lower if the capacity variables $X_{ij;t}$ are continuous. Given the deterministic solution, the **Skeleton** method fixes the service selection variables $V_{ij;t}$ and determines the capacities by solving a stochastic LP. Therefore this approach can be seen as a viable heuristic if the results are strong.

The performances of the **Determ**, **Skeleton** and **Upgrade** approaches for the variable capacity model with integer and continuous capacities, together with the results from the fixed capacity model, are displayed in Table 2. Test results with deterministic demands set at the 50^{th} and 75^{th} percentiles of their corresponding distributions are shown. Asset-balance constraints are not included in the models tested here. The same instances are used for every row in the table.

When comparing the average losses for the designs corresponding to the 50^{th} and 75^{th} percentile deterministic demands, we see that the 75^{th} percentile designs always perform better. This of course depends on problem parameters, in particular, how much more

Table 2 Comparing different ways of using deterministic information in the stochastic environment (without asset balance) for the Small cases

Model and Parameter Setting	Average Loss		
	Determ	**Skeleton**	**Upgrade**
Fixed Capacity, 50^{th}	16.63%		4.03%
Fixed Capacity, 75^{th}	8.86%		2.16%
Variable Integer Capacity, 50^{th}	51.71%	15.28%	3.73%
Variable Integer Capacity, 75^{th}	11.46%	6.41%	2.34%
Variable Continuous Capacity, 50^{th}	49.76%	15.82%	3.90%
Variable Continuous Capacity, 75^{th}	11.22%	5.67%	2.48%

expensive the ad hoc capacity is. This approach is in line with what is done in many industries, where demands well above the mean are used as forecasts.

Table 3 Comparing different ways of using deterministic information in the stochastic environment (with asset balance) for the Small cases

Model and Parameter Setting	Average Loss		
(with asset balance)	**Determ**	**Skeleton**	**Upgrade**
Fixed Capacity, 50^{th}	15.21%		3.83%
Fixed Capacity, 75^{th}	9.28%		2.39%
Variable Integer Capacity, 50^{th}	30.81%	13.85%	3.32%
Variable Integer Capacity, 75^{th}	9.92%	5.20%	2.88%
Variable Continuous Capacity, 50^{th}	27.08%	11.51%	3.24%
Variable Continuous Capacity, 75^{th}	9.35 %	6.09%	3.12%

Table 3 displays the corresponding results for the cases with asset-balance constraints, as introduced in Section 3.4. Again, using the 75^{th} percentile of the demand distribution is a better choice when obtaining the deterministic solution. For the fixed and the variable capacity models, with both integer and continuous capacity settings, **Determ** (75^{th}) produces average losses that are all less than 10%. The **Skeleton** method can further improve the performance of the deterministic solution with not much computational effort: a much smaller MIP in the integer capacity case and an LP in the continuous capacity case, both on a reduced skeleton network.

By comparing the numbers in Tables 2 and 3, we observe that, in general, the average losses for models with asset balance (Table 3) are lower. Such a significant drop in losses (especially for the cases “Variable Integer Capacity, 50^{th}” and “Variable Continuous Capacity, 50^{th}”) actually represents the difference between two situations in real applications: the situation where the resources are booked from elsewhere, and the situation where the decision maker owns the resources and thus needs to manage also the empty moves. In the latter situation, the movements of empty vehicles typically leads to an increase in the expected transport capability on the space-time network (considered by the decision maker) in a stochastic setting. For example, in many applications within maritime transportation the ships are owned by the decision maker and therefore ballast sailings need to be performed in order to reposition the empty ships. Then there is a chance that the ballast leg of one ship may be able to pick up the unmet demand of some commodity serviced by another ship, in some scenarios where there is a surge in demand of that commodity.

5.6. Performance of Deterministic Approaches on Larger Instances

We now perform similar comparison tests on the *Larger* set of instances while extending the maximal run time to 5 hours. Using instances of larger and various sizes, we seek to further check the consistency of the insights drawn from the previous *Small* set of instances, and to investigate the potential usage of deterministic solution based methods on larger problems where stochastic programs are not numerically solvable.

Table 4 displays the comparison, for all instances in the *Larger* set, between CPLEX and the best-performing deterministic based approach (according to Table 2) for each model: the **Determ** 75^{th} approach for the fixed capacity model, and the **Skeleton** 75^{th} approach for both variable integer- and continuous capacity models. The **Upgrade** approaches are excluded here since they are not numerically efficient, as discussed in previous sections. The size of an instance is indicated by the number of terminals (N), commodities (K), time periods (T) and scenarios (S), and the “r/s” column indicates whether the demands are with random skewness (r) or all symmetrical (s). We record the optimality gaps between the best found solution and best bound reported by CPLEX under the “Gap” columns, and, compared with these best found solutions, the relative loss of deterministic based approaches under the “Loss” columns. Also note that the times recorded for **Determ** 75^{th} and **Skeleton** 75^{th} approaches include the time for solving the deterministic problem (also with 5 hours time limit) as well as the time taken for the improvement afterwards.

Table 4 Comparing selected approaches and CPLEX on the Larger instance set

Instance					Fixed Capacity				Variable Integer Cap.				Variable Continuous Cap.			
					CPLEX		Determ 75^{th}		CPLEX		Skel. 75^{th}		CPLEX		Skel. 75^{th}	
N	K	T	S	r/s	Gap (%)	Time (min)	Loss (%)	Time (min)	Gap (%)	Time (min)	Loss (%)	Time (min)	Gap (%)	Time (min)	Loss (%)	Time (min)
6	16	7	40	r	0.01	35.4	+6.0	0.01	2.14	300.0	+3.7	11.7	1.02	300.0	+1.8	0.02
6	16	7	40	s	0.00	22.7	+8.8	0.01	6.87	300.0	+1.0	6.6	4.08	300.0	+1.7	0.01
6	16	7	60	r	0.00	5.9	+9.3	0.02	6.70	300.0	+1.9	2.0	2.46	300.0	+4.0	0.02
6	16	7	60	s	0.01	2.4	+17.2	0.01	2.01	300.0	+3.9	0.6	0.00	232.5	+4.2	0.01
6	16	9	40	r	1.62	300.0	+8.5	0.02	6.22	300.0	+2.6	32.7	4.91	300.0	+2.6	0.03
6	16	9	40	s	0.01	60.6	+7.2	0.02	1.52	300.0	+4.2	47.1	1.55	300.0	+2.4	0.02
6	16	9	60	r	0.01	98.5	+10.3	0.02	4.60	300.0	+4.3	34.6	1.51	300.0	+4.6	0.02
6	16	9	60	s	0.99	300.0	+7.3	0.02	4.12	300.0	+2.3	115.4	7.52	300.0	-2.6	0.02
6	20	7	40	r	0.01	4.6	+5.0	0.02	2.56	300.0	+2.6	64.0	0.82	300.0	+4.2	0.02
6	20	7	40	s	0.01	52.9	+6.2	0.05	4.95	300.0	+4.5	81.1	7.64	300.0	+0.0	0.06
6	20	7	60	r	1.06	300.0	+8.6	0.03	9.62	300.0	+1.7	62.0	12.90	300.0	-3.2	0.04
6	20	7	60	s	0.01	51.8	+5.1	0.05	4.57	300.0	+4.5	101.7	3.17	300.0	+6.2	0.04
6	20	9	40	r	0.25	300.0	+1.9	0.02	9.09	300.0	+1.6	301.9	17.36	300.0	-5.9	0.11
6	20	9	40	s	0.01	63.7	+6.8	0.04	4.39	300.0	+5.6	92.9	2.49	300.0	+2.9	0.05
6	20	9	60	r	0.47	300.0	+6.7	0.03	37.99	300.0	-31.2	300.1	30.75	300.0	-9.6	0.03
6	20	9	60	s	0.01	287.4	+7.7	0.11	21.25	300.0	-16.2	300.2	14.17	300.0	-4.1	0.12
10	16	7	40	r	0.01	44.0	+11.8	0.02	6.42	300.0	+4.8	188.1	6.01	300.0	-1.3	0.03
10	16	7	40	s	0.01	92.9	+5.4	0.02	1.40	300.0	+3.7	231.6	0.30	300.0	+2.9	0.03
10	16	7	60	r	0.41	300.0	+12.3	0.02	48.52	300.0	-41.6	300.1	76.95	300.0	-73.4	0.03
10	16	7	60	s	0.87	300.0	+3.8	0.04	38.11	300.0	-32.3	300.8	76.97	300.0	-75.0	0.05
10	16	9	40	r	0.00	5.0	+4.8	0.02	4.67	300.0	+2.3	69.3	2.02	300.0	+4.2	0.03
10	16	9	40	s	0.00	8.5	+7.0	0.03	34.11	300.0	-25.2	21.8	74.74	300.0	-66.7	0.03
10	16	9	60	r	0.02	300.0	+5.5	0.05	53.73	300.0	-48.7	300.0	79.20	300.0	-72.4	0.05
10	16	9	60	s	0.02	300.0	+2.9	0.04	27.37	300.0	-17.1	300.1	27.13	300.0	-12.9	0.04
Average					0.24	147.3	+7.3	0.03	14.29	300.0	-6.5	136.2	18.99	297.2	-11.9	0.04

For the fixed capacity model, Table 4 shows that CPLEX manages to solve or find near-optimal solutions to all stochastic problems within 5 hours. In contrast to CPLEX, the **Determ** 75^{th} approach produces on average a loss of 7.3%. However, it takes no more than 10 seconds for each instance, which is a fraction of the time needed by CPLEX.

For the variable integer capacity model, the **Skeleton** 75^{th} approach provides solutions with good quality, some greatly outperform the best solutions found by CPLEX (hence with significantly negative losses), especially on the instances with larger size (such as instance "10 16 9 60 r"). Additionally, this is achieved within much shorter time (136.2 minutes on average), which is in fact primarily spent on solving the deterministic versions of the problems; by contrast the resolving part of **Skeleton** 75^{th} (based on fixed service selections and adjusting the capacities) only takes a few minutes. When the size of the instance increases, CPLEX starts to struggle with solving the deterministic problems to optimality (with gaps up to 4.89%), and hence uses up the 5 hours limit for some instances, e.g., instance "10 16 7 60 s" where the total run time is 300.8 but the resolving part only takes 0.8 minutes. This is also an indication that we can speed up the **Skeleton** 75^{th} approach by reducing the time limit for solving the problem deterministically.

For the variable continuous capacity model, Table 4 shows that the **Skeleton** 75^{th} approach is far superior. It manages to obtain solutions with 11.9% *improvement* on average relative to the best solutions found by CPLEX after 5 hours, all finished within one minute. Also note that in this case where capacity is continuous (compared to the integer capacity case), solving the deterministic versions of the stochastic problems become very fast; the stochastic problems themselves, on the other hand, are equally hard to solve numerically.

In Table 5, we further show the relative (average) performance of different deterministic solution based approaches, i.e., **Determ** 75^{th}, **Skeleton** 50^{th} and **Skeleton** 75^{th}, benchmarked against **Determ** 50^{th}. The results for both the *Small* and *Larger* sets are shown. Consistent with the observations made on the *Small* set of instances, we see significant improvements from using the 75^{th} percentile rather than the 50^{th} percentile of the demand distributions, and that the **Skeleton** approaches in general perform better than the **Determ** ones.

Table 5 Comparing different approaches on both sets of instances

	Fixed Capacity		Variable Integer Cap.				Variable Continuous Cap.			
Sets	Det.50^{th}	Det.75^{th}	Det.50^{th}	Det.75^{th}	Skel.50^{th}	Skel.75^{th}	Det.50^{th}	Det.75^{th}	Skel.50^{th}	Skel.75^{th}
Small	0.0%	-5.8%	0.0%	-24.0%	-26.2%	-29.4%	0.0%	-22.8%	-23.3%	-27.1%
Larger	0.0%	-6.7%	0.0%	-26.4%	-29.8%	-33.0%	0.0%	-23.7%	-26.1%	-30.2%
Average	0.0%	-6.2%	0.0%	-25.2%	-28.0%	-31.2%	0.0%	-23.2%	-24.7%	-28.6%

6. Conclusion

In this paper, we have discussed the value and the upgradability of deterministic solutions in scheduled stochastic service network design problems, for fixed and variable capacity models with both integer and continuous capacity settings. In those situations where deterministic solutions can be found, optimally or heuristically, we may upgrade these solutions into well performing solutions to the corresponding stochastic problems.

For the fixed capacity model, by adding a limited number of extra services, the deterministic design can become structurally different, and much better suited for the stochastic environment. For the variable capacity model, this can also be achieved by using part of the deterministic design information (the skeleton) and also with not much computational effort. In particular, when the capacities are continuous or when the integrality of the capacities can, for all practical purposes, be relaxed, the **Skeleton** method becomes an LP on a reduced skeleton network. We also show that it is a better practice for our problems to use the 75^{th} percentile of the random demands when obtaining the deterministic solutions. Note that this also depends on problem parameters, in particular, how much more expensive the ad hoc capacity is.

To quantitatively investigate the structural improvements from the deterministic design to better performing solutions in the stochastic environment, a measurement scheme has been used to evaluate the level of the potentially beneficial structural features: multi-path usage and path-sharing. It was concluded that, in general, the better the solution performs in the stochastic environment, the higher the levels of multiple-path usage and path-sharing it displays. The reverse is not true, but still, this might lead to possible ways to develop heuristic approaches for the stochastic problem.

Therefore, an interesting direction of future research may be to explore whether we find the “correct” extra services based on the deterministic solution (or even a feasible solution), using the beneficial structural features confirmed in this paper. For example, if certain services increase the level of multi-path usage and path-sharing in the network, then these might be the potentially “correct” extra services for the stochastic problem.

Another research avenue is to investigate the existence of similar upgradability of deterministic solutions in other network design problems. As mentioned earlier, such upgradability is not obvious at all in some other stochastic problems (Maggioni and Wallace 2012). We may be able to determine under what circumstances the deterministic solution is useful in the stochastic environment, and to see if a certain modeling factor is found to have great impact on the upgradability of the deterministic solution.

Acknowledgments

Partial funding for this project has been provided by the Natural Sciences and Engineering Council of Canada (NSERC), through its Industrial Research Chair and Discovery Grant programs. We also gratefully acknowledge the support of Fonds de recherche du Québec through their infrastructure grants. Partial support was also received from SNF – The Centre for Applied Research at NHH – via its support for the Center for Shipping and Logistics at NHH. Finally, we thank one referee for a set of valuable comments that improved the paper.

References

Andersen J, Christiansen M, Crainic TG, Grønhaug R (2011) Branch and Price for Service Network Design with Asset Management Constraints. *Transportation Science* 45(1):33–49, ISSN 0041-1655, URL `http://dx.doi.org/10.1287/trsc.1100.0333`.

Andersen J, Crainic TG, Christiansen M (2009a) Service network design with asset management: Formulations and comparative analyses. *Transportation Research Part C: Emerging Technologies* 17(2):197–207, ISSN 0968090X, URL `http://dx.doi.org/10.1016/j.trc.2008.10.005`.

Andersen J, Crainic TG, Christiansen M (2009b) Service network design with management and coordination of multiple fleets. *European Journal of Operational Research* 193(2):377–389, ISSN 03772217, URL `http://dx.doi.org/10.1016/j.ejor.2007.10.057`.

Armacost AP, Barnhart C, Ware KA (2002) Composite Variable Formulations for Express Shipment Service Network Design. *Transportation Science* 36(1):1–20, ISSN 0041-1655, URL `http://dx.doi.org/10.1287/trsc.36.1.1.571`.

Bai R, Wallace SW, Li J, Chong AYL (2014) Stochastic service network design with rerouting. *Transportation Research Part B: Methodological* 60:50–65, ISSN 01912615, URL `http://dx.doi.org/10.1016/j.trb.2013.11.001`.

Barnhart C, Schneur RR (1996) Air Network Design for Express Shipment Service. *Operations Research* 44(6):852–863, ISSN 0030-364X, URL `http://dx.doi.org/10.1287/opre.44.6.852`.

Birge JR (1982) The value of the stochastic solution in stochastic linear programs with fixed recourse. *Mathematical Programming* 24(1):314–325, ISSN 0025-5610, URL `http://dx.doi.org/10.1007/BF01585113`.

Büdenbender K, Grünert T, Sebastian HJ (2000) A Hybrid Tabu Search/Branch-and-Bound Algorithm for the Direct Flight Network Design Problem. *Transportation Science* 34(4):364–380, ISSN 0041-1655, URL `http://dx.doi.org/10.1287/trsc.34.4.364.12319`.

Chopra VK, Ziemba WT (2011) The effect of errors in means, variances, and covariances on optimal portfolio choice. MacLean LC, Thorp EO, Ziemba WT, eds., *The Kelly Capital Growth Investment Criterion: Theory and Practice*, volume 3 of *World Scientific Handbook in Financial Economics Series*, 249–257 (World Scientific Publishing).

Christiansen M, Fagerholt K, Nygreen B, Ronen D (2007) Chapter 4 Maritime Transportation. Barnhart C, Laporte G, eds., *Handbooks in operations research and management science*, volume 14, 189–284 (Elsevier), URL `http://dx.doi.org/10.1016/S0927-0507(06)14004-9`.

Christiansen M, Fagerholt K, Nygreen B, Ronen D (2013) Ship routing and scheduling in the new millennium. *European Journal of Operational Research* 228(3):467–483, ISSN 0377-2217, URL `http://dx.doi.org/10.1016/j.ejor.2012.12.002`.

Christiansen M, Fagerholt K, Ronen D (2004) Ship routing and scheduling: Status and perspectives. *Transportation science* 38(1):1–18, ISSN 0041-1655, URL `http://dx.doi.org/10.1287/trsc.1030.0036`.

Cordeau JF, Toth P, Vigo D (1998) A Survey of Optimization Models for Train Routing and Scheduling. *Transportation Science* 32(4):380–404, ISSN 0041-1655, URL `http://dx.doi.org/10.1287/trsc.32.4.380`.

Crainic TG (2000) Service network design in freight transportation. *European Journal of Operational Research* 122(2):272–288, ISSN 03772217, URL `http://dx.doi.org/10.1016/S0377-2217(99)00233-7`.

Crainic TG (2003) Long-Haul Freight Transportation. Hall RW, ed., *Handbook of Transportation Science*, 451–516 (Boston: Kluwer Academic Publishers), second edition, URL `http://dx.doi.org/10.1007/0-306-48058-1_13`.

Crainic TG, Dejax P, Delorme L (1989) Models for multimode multicommodity location problems with interdepot balancing requirements. *Annals of Operations Research* 18(1):277–302, ISSN 0254-5330, URL `http://dx.doi.org/10.1007/BF02097809`.

Crainic TG, Ferland JA, Rousseau JM (1984) A Tactical Planning Model for Rail Freight Transportation. *Transportation Science* 18(2):165–184, ISSN 0041-1655, URL `http://dx.doi.org/10.1287/trsc.18.2.165`.

Crainic TG, Fu X, Gendreau M, Rei W, Wallace SW (2011) Progressive hedging-based metaheuristics for stochastic network design. *Networks* 58(2):114–124, ISSN 00283045, URL `http://dx.doi.org/10.1002/net.20456`.

Crainic TG, Kim KH (2007) Chapter 8 Intermodal Transportation. Barnhart C, Laporte G, eds., *Transportation*, volume 14 of *Handbooks in Operations Research and Management Science*, 467–537 (North-Holland, Amsterdam), URL `http://dx.doi.org/10.1016/S0927-0507(06)14008-6`.

Crainic TG, Laporte G (1997) Planning models for freight transportation. *European Journal of Operational Research* 97(3):409–438, ISSN 03772217, URL `http://dx.doi.org/10.1016/S0377-2217(96)00298-6`.

Crainic TG, Rousseau JM (1986) Multicommodity, multimode freight transportation: A general modeling and algorithmic framework for the service network design problem. *Transportation Research Part B: Methodological* 20(3):225–242, ISSN 01912615, URL http://dx.doi.org/10.1016/0191-2615(86)90019-6.

Crainic TG, Roy J (1988) OR tools for tactical freight transportation planning. *European Journal of Operational Research* 33(3):290–297, ISSN 03772217, URL http://dx.doi.org/10.1016/0377-2217(88)90172-5.

Dejax P, Crainic TG (1987) A Review of Empty Flows and Fleet Management Models in Freight Transportation. *Transportation Science* 21(4):227–248, ISSN 0041-1655, URL http://dx.doi.org/10.1287/trsc.21.4.227.

Delorme L, Roy J, Rousseau JM (1988) Motor-Carriers Operations Planning Models: A State of the Art. Bianco L, Bella AL, eds., *Freight Transport Planning and Logistics*, 510–545 (Springer-Verlag, Berlin), URL http://dx.doi.org/10.1007/978-3-662-02551-2_17.

Farvolden JM, Powell WB (1991) A Dynamic Network Model for Less-Than-Truckload Motor Carrier Operations. Working Paper 90-05, Department of Industrial Engineering, University of Toronto, Toronto, ON, Canada.

Farvolden JM, Powell WB (1994) Subgradient Methods for the Service Network Design Problem. *Transportation Science* 28(3):256–272, ISSN 0041-1655, URL http://dx.doi.org/10.1287/trsc.28.3.256.

Farvolden JM, Powell WB, Lustig IJ (1993) A Primal Partitioning Solution for the Arc-Chain Formulation of a Multicommodity Network Flow Problem. *Operations Research* 41(4):669–693, ISSN 0030-364X, URL http://dx.doi.org/10.1287/opre.41.4.669.

Gorman MF (1998a) An application of genetic and tabu searches to the freight railroad operating plan problem. *Annals of Operations Research* 78:51–69, ISSN 02545330, URL http://dx.doi.org/10.1023/A:1018906301828.

Gorman MF (1998b) Santa Fe Railway uses an operating-plan model to improve its service design. *Interfaces* 28(4):1–12.

Grünert T, Sebastian HJ (2000) Planning models for long-haul operations of postal and express shipment companies. *European Journal of Operational Research* 122(2):289–309, ISSN 03772217, URL http://dx.doi.org/10.1016/S0377-2217(99)00234-9.

Grünert T, Sebastian HJ, Tharigen M (1999) The design of a letter-mail transportation network by intelligent techniques. Sprague R, ed., *Proceedings of the 32nd Annual Hawaii International Conference on Systems Sciences. 1999. HICSS-32.*, 16 (IEEE Comput. Soc), ISBN 0-7695-0001-3, URL http://dx.doi.org/10.1109/HICSS.1999.772628.

Haghani AE (1989) Formulation and solution of a combined train routing and makeup, and empty car distribution model. *Transportation Research Part B: Methodological* 23(6):433–452, ISSN 01912615, URL http://dx.doi.org/10.1016/0191-2615(89)90043-X.

Higle JL, Wallace SW (2003) Sensitivity Analysis and Uncertainty in Linear Programming. *Interfaces* 33(4):53–60, ISSN 0092-2102, URL http://dx.doi.org/10.1287/inte.33.4.53.16370.

Hoff A, Lium AG, Løkketangen A, Crainic TG (2010) A metaheuristic for stochastic service network design. *Journal of Heuristics* 16(5):653–679, ISSN 1381-1231, URL http://dx.doi.org/10.1007/s10732-009-9112-8.

Høyland K, Kaut M, Wallace SW (2003) A heuristic for moment-matching scenario generation. *Computational Optimization and Applications* 24(2/3):169–185, ISSN 09266003, URL http://dx.doi.org/10.1023/A:1021853807313.

Kall P, Wallace SW (1994) *Stochastic Programming* (Chichester: John Wiley & Sons), 2 edition.

Kaut M, Wallace SW (2007) Evaluation of scenario-generation methods for stochastic programming. *Pacific Journal of Optimization* 3(2):257–271.

Keaton MH (1989) Designing optimal railroad operating plans: Lagrangian relaxation and heuristic approaches. *Transportation Research Part B: Methodological* 23(6):415–431, ISSN 01912615, URL http://dx.doi.org/10.1016/0191-2615(89)90042-8.

Keaton MH (1991) Service-cost tradeoffs for carload freight traffic in the U.S. rail industry. *Transportation Research Part A: General* 25(6):363–374, ISSN 01912607, URL http://dx.doi.org/10.1016/0191-2607(91)90014-H.

Keaton MH (1992) Designing Railroad Operating Plans: A Dual Adjustment Method for Implementing Lagrangian Relaxation. *Transportation Science* 26(4):263–279, ISSN 0041-1655, URL http://dx.doi.org/10.1287/trsc.26.4.263.

Kim D, Barnhart C, Ware KA, Reinhardt G (1999) Multimodal Express Package Delivery: A Service Network Design Application. *Transportation Science* 33(4):391–407, ISSN 0041-1655, URL http://dx.doi.org/10.1287/trsc.33.4.391.

King AJ, Wallace SW (2012) *Modeling with Stochastic Programming*. Springer Series in Operations Research and Financial Engineering (Springer New York), ISBN 978-0-387-87816-4, URL http://dx.doi.org/10.1007/978-0-387-87817-1.

Lamar BW, Sheffi Y, Powell WB (1990) A Capacity Improvement Lower Bound for Fixed Charge Network Design Problems. *Operations Research* 38(4):704–710, ISSN 0030-364X, URL http://dx.doi.org/10.1287/opre.38.4.704.

Lium AG, Crainic TG, Wallace SW (2007) Correlations in stochastic programming: A case from stochastic service network design. *Asia-Pacific Journal of Operational Research* 24(02):161–179, ISSN 0217-5959, URL http://dx.doi.org/10.1142/S0217595907001206.

Lium AG, Crainic TG, Wallace SW (2009) A Study of Demand Stochasticity in Service Network Design. *Transportation Science* 43(2):144–157, ISSN 0041-1655, URL http://dx.doi.org/10.1287/trsc.1090.0265.

Maggioni F, Wallace SW (2012) Analyzing the quality of the expected value solution in stochastic programming. *Annals of Operations Research* 200(1):37–54, ISSN 0254-5330, URL http://dx.doi.org/10.1007/s10479-010-0807-x.

Newton HN, Barnhart C, Vance PH (1998) Constructing Railroad Blocking Plans to Minimize Handling Costs. *Transportation Science* 32(4):330–345, ISSN 0041-1655, URL http://dx.doi.org/10.1287/trsc.32.4.330.

Pantuso G, Fagerholt K, Wallace SW (2017) Which uncertainty is important in multistage stochastic programmes? A case from maritime transportation. *IMA Journal of Management Mathematics* 28(1):5–17, ISSN 1471-678X, URL http://dx.doi.org/10.1093/imaman/dpu026.

Pedersen MB, Crainic TG, Madsen OBG (2009) Models and Tabu Search Metaheuristics for Service Network Design with Asset-Balance Requirements. *Transportation Science* 43(2):158–177, ISSN 0041-1655, URL http://dx.doi.org/10.1287/trsc.1080.0234.

Powell WB (1986) A Local Improvement Heuristic for the Design of Less-than-Truckload Motor Carrier Networks. *Transportation Science* 20(4):246–257, ISSN 0041-1655, URL http://dx.doi.org/10.1287/trsc.20.4.246.

Powell WB, Sheffi Y (1983) The load planning problem of motor carriers: Problem description and a proposed solution approach. *Transportation Research Part A: General* 17(6):471–480, ISSN 01912607, URL http://dx.doi.org/10.1016/0191-2607(83)90167-X.

Powell WB, Sheffi Y (1986) Interactive Optimization for Motor Carrier Load Planning. *Journal of Business Logistics* 7(2):64–90.

Powell WB, Sheffi Y (1989) Design and Implementation of an Interactive Optimization System for Network Design in the Motor Carrier Industry. *Operations Research* 37(1):12–29, ISSN 0030-364X, URL http://dx.doi.org/10.1287/opre.37.1.12.

Roy J, Delorme L (1989) Netplan: A Network Optimization Model For Tactical Planning In The Less-Than-Truckload Motor-Carrier Industry. *INFOR: Information Systems and Operational Research* 27(1):22–35, ISSN 0315-5986, URL http://dx.doi.org/10.1080/03155986.1989.11732079.

Sharypova K, Crainic TG, van Woensel T, Fransoo JC (2012) Scheduled service network design with synchronization and transshipment constraints for intermodal container transportation networks. Technical Report CIRRELT-2012-77, Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport, Montréal, QC, Canada.

Smilowitz KR, Atamtürk A, Daganzo CF (2003) Deferred item and vehicle routing within integrated networks. *Transportation Research Part E: Logistics and Transportation Review* 39(4):305–323, ISSN 13665545, URL http://dx.doi.org/10.1016/S1366-5545(02)00048-0.

Thapalia BK, Crainic TG, Kaut M, Wallace SW (2011) Single-Commodity Stochastic Network Design with Multiple Sources and Sinks. *INFOR: Information Systems and Operational Research* 49(3):193–211, URL http://dx.doi.org/10.3138/infor.49.3.003.

Thapalia BK, Crainic TG, Kaut M, Wallace SW (2012a) Single-commodity network design with random edge capacities. *European Journal of Operational Research* 220(2):394–403, ISSN 03772217, URL `http://dx.doi.org/10.1016/j.ejor.2012.01.026`.

Thapalia BK, Wallace SW, Kaut M, Crainic TG (2012b) Single source single-commodity stochastic network design. *Computational Management Science* 9(1):139–160, ISSN 1619-697X, URL `http://dx.doi.org/10.1007/s10287-010-0129-0`.

Wallace SW (2000) Decision Making Under Uncertainty: Is Sensitivity Analysis of Any Use? *Operations Research* 48(1):20–25, ISSN 0030-364X, URL `http://dx.doi.org/10.1287/opre.48.1.20.12441`.

Zhu E, Crainic TG, Gendreau M (2014) Scheduled Service Network Design for Freight Rail Transportation. *Operations Research* 62(2):383–400, ISSN 0030-364X, URL `http://dx.doi.org/10.1287/opre.2013.1254`.