Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport

Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation

Stochastic Scheduled Service Network Design: The Value of Deterministic Solutions

Xin Wang
Teodor Gabriel Crainic
Stein W. Wallace

March 2016

CIRRELT-2016-14

Bureaux de Montréal :
Université de Montréal
Pavillon André-Aisenstadt
C.P. 6128, succursale Centre-ville
Montréal (Québec)
Canada H3C 3J7
Téléphone : 514 343-7575
Télécopie : 514 343-7121

Bureaux de Québec :
Université Laval
Pavillon Palasis-Prince
2325, de la Terrasse, bureau 2642
Québec (Québec)
Canada G1V 0A6
Téléphone : 418 656-2073
Télécopie : 418 656-2624

www.cirrelt.ca

Stochastic Scheduled Service Network Design: The Value of Deterministic Solutions†

Xin Wang[1], Teodor Gabriel Crainic[2,*], Stein W. Wallace[3]

[1] Department of Management Science,Lancaster University Management School, Bailrigg, Lancaster, LA1 4YX, United Kingdom

[2] Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT) and Department of Management and Technology, Université du Québec à Montréal, P.O. Box 8888, Station Centre-Ville, Montréal, Canada H3C 3P8

[3] Department of Business and Management Science, Norwegian School of Economics, NHH, Hellevein 30, NO-5045, Bergen, Norway

Abstract. We study the value of deterministic solutions, in particular their quality and upgradeability, to stochastic scheduled service network design problems. We present several service network design models and investigate, for each case, the immediate quality of the deterministic solutions stemming from the 50th and the 75th 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.

Keywords: Scheduled service network design, stochastic programming, value of deterministic solution.

Acknowledgements. 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 – Nature et technologie (FRQNT) 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.

†Revised version of the CIRRELT-2014-46.

---

* Corresponding author: teodorgabriel.crainic@cirrelt.net

Dépôt légal – Bibliothèque et Archives nationales du Québec
Bibliothèque et Archives Canada, 2016

1 Introduction

In the freight transportation industry, service network design (Crainic, 2000) helps provide solutions to a set of tactical problems: where, when and how to offer services for the delivery of the demands. The goal is to decide the selection, routing and scheduling of services, while balancing the operating costs and service quality. Discrete decision variables are normally involved in such processes and the resulting programming models are thus usually very complex.

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, however, 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, by investigating their quality 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 the deterministic counterparts, 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 or not 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 therefore present models for several service network design problem settings, each representing an important set of applications in the freight transportation industry, and aim to provide a comprehensive analysis with regards to the value of deterministic solutions in each case. More specifically, we introduce, in the context of time-dependent scheduled service network design, 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 are 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 upgradeability of deterministic solutions to stochastic scheduled service network design problems. It adds to the discussion of whether deterministic solutions are of any value to stochastic models, and provides insights into 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. Additionally, inspired by the potentially beneficial structural features from the literature, we present a measurement scheme to quantitatively show how a deterministic solution may be upgraded.

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. The proposed comparison tests and the instance generation procedure are discussed in Section 4. Section 5 presents and analyzes the computational results. 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) 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 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 tradeoffs 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 interservice operations at terminals. Such models have been proposed for multimodal 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; Buedenbender 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, 1996; 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., 1992), 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., 2013) 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 estimations 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, it 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 (2010). 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, scheduled service network design (SSND) problem, in which a periodic, cyclic schedule is built for a number of commodities (OD pairs). We present, in this section, variant models with fixed and variable capacity, and the additional 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. (2013) 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. An arc $a = (i, j; t)$ represents the service departing from terminal i at time point t and arriving at terminal j. A service can be set up at any time point between any pair of terminals $(i, j), \forall i, j \in \mathcal{N}, i \neq j$, in either direction. It is assumed that a service can take one or more time periods, depending on the physical distance between the two terminals. We use l_{ij} to represent the service length between terminals i and j, i.e., the number of time periods required for transporting goods between the two terminals. 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.

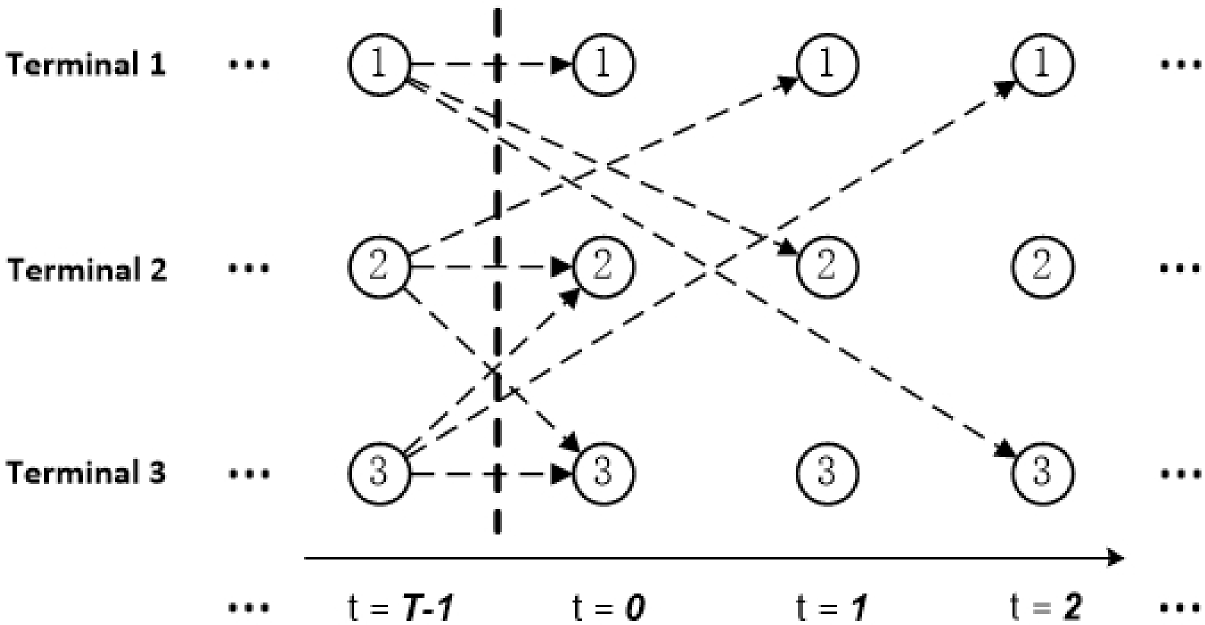

Figure 1: An example illustrating possible services (represented by dashed arrows) that can be set up at $t = T - 1$, i.e. the last time point of the repetitive and cyclic schedule.

Figure 1 shows an example with three terminals and a repetitive and cyclic schedule over T time periods, and illustrates the possible services 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 and the time point by which it must be delivered; 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 represented by a 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 $(i, j; t) \in \mathcal{A}$. Also, we need to pay for commodity flows, that is, the transportation and storage of the commodities. Thus costs $e_{ij;t}$ associated to an arc $(i, j; t)$ represent the unit flow costs incurred to move commodities on the opened services or have them wait at terminals. 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, in the following, the formulation with fixed (but not identical) capacity for every opened service. A fixed capacity $h_{ij;t}$ is therefore associated with arc $(i, j; t)$. Let $V_{ij;t}$ represent the {0,1} service 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.

Note that we use {0,1} decision variables to capture the service selection choices, indicating whether or not the service leaves at the specified time point. Therefore, only one service, with specified characteristics, is allowed at a given time point from one terminal to another. When several departures are possible in the same time period, general (non-negative) integer variables must be used. However, by making the time periods appropriately small, one can always use {0,1} variables to address multiple departures within a certain period of time. Therefore, without loss of generality, we use only {0,1} variables to represent service selection decisions for all models presented in this paper.

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:

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

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

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

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

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

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

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 or not 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:

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

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

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

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

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

$$0 \le Z^s_k \le \delta^s_k \qquad \forall k \in \mathcal{K}, \forall s \in \mathcal{S} \tag{13}$$

$$0 \le Y^s_{ij;t;k} \le \delta^s_k \qquad \forall i, j \in \mathcal{N}, \forall t \in \mathcal{T}, \forall k \in \mathcal{K}, \forall s \in \mathcal{S} \tag{14}$$

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 was 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

$$\sum_{i \in \mathcal{N}} V_{ij;t \ominus l_{ij}} = \sum_{i \in \mathcal{N}} V_{ji;t} \qquad \forall j \in \mathcal{N}, \forall t \in \mathcal{T}, \qquad (15)$$

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 (15) 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 (16) where, for simplicity of presentation we assume all units are the same for all services.

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

4 Comparison tests and instance generation

In this section, we propose a set of comparison tests with the intention of investigating the value of the deterministic solution in a stochastic setting. 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.

The comparison tests are presented in Section 4.1, and are based on solving the deterministic versions of the presented stochastic models. The stochastic formulations may easily be transformed into their respective deterministic versions, by inputting one single scenario only. Furthermore, we present in Section 4.2 how we generated the test instances used in the computational study.

4.1 Performance Comparison

For a given scenario tree, we first use three comparison tests inspired by Thapalia et al. (2012b) to compare the performances of different designs in the stochastic environment:

1. Stochastic solution: (optimal solution)

   We solve the stochastic problem. This is used as a benchmark for other solutions.

2. Deterministic design used in the stochastic model: (**VSS**)

   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. This applies to both fixed and variable capacity models:

   (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.

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

   This also applies to both fixed and variable capacity models:

   (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 **VSS** and **Upgrade** are performed to check the immediate performance and upgradeability, 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:

4. 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 5.4), as the stochastic program used to set capacities will then be a stochastic LP instead of a stochastic MIP.

4.2 Parameter setting and instance generation

To perform the computational study with the proposed comparison tests, we use a set of problem instances generated following a random-based procedure. As it is preferable to use the "true" stochastic solution as benchmark, we only study cases where optimality can be found numerically (using CPLEX) for all comparison tests.

We consider the stochastic SSND problem on a space-time network with 7 time points, 6 terminals and 16 commodities. To start with, we randomly generate 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 lengths are decided according to the physical distances between the two terminals of the service, such that for any $i, j \in \mathcal{N}$ and $i \neq j$, service length l_{ij} is assigned with an integer value: 1, 2 or 3 in our experiments. The values of the unit flow costs $e_{ij;t}$ 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 much higher than 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 5.1.

For every commodity, its origin and destination terminals are both selected randomly. The time span (from the time point it becomes available to the time point it has to be delivered) ranges from 2 to 5. The stochastic demands of all commodities are subject to symmetric triangular distributions with a standard deviation equal to 40% of the mean.

We discretize the demand distributions by generating scenarios with equal probabilities to represent the stochasticity. The scenario generation process is performed using the moment-matching method introduced by Høyland et al. (2003). The demand correlation matrix needed to generate the scenarios for every instance is created as follows: (a) the commodities are equally (or almost equally, if there is an odd number of commodities) divided into two groups; (b) if two commodities are in the same group, their demands are assumed to be positively correlated with a correlation coefficient randomly chosen from [0.00,0.50); otherwise, if the two commodities are in different groups, their demands are negatively correlated with a correlation coefficient randomly chosen from (-0.50,0.00]; (c) the resulting matrix has to be positive semi-definite to ensure its validity as a correlation matrix; if not, step (b) is repeated. This way of constructing correlation matrices normally leads directly to positive semi-definiteness.

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. Thus there is a trade-off between the stability of the stochastic solution and the problem growing too large. In our experiments, we use 30 scenarios to represent the stochasticity. The in-sample stability tests (Kaut and Wallace, 2007) give a difference of less than 5%, which is acceptable, between the highest and lowest optimal objective function values on a large number of different scenario trees.

5 Computational results

For the computational study presented in this section, we use ten problem instances randomly generated following the procedure described in Section 4.2. We then generate 30 scenarios for each instance to represent the demand stochasticity. All programs are solved to optimality using CPLEX 12.5.

Sections 5.1 and 5.2 report the results of performing the comparison tests and discuss the value of deterministic solutions for the fixed and variable models, respectively. Section 5.3 examines the potential structural improvements that take place when upgrading the deterministic solutions using a quantitative approach. We also investigate the value of deterministic solutions in the cases of using continuous capacities and including asset-balance constraints in Sections 5.4 and 5.5, respectively.

5.1 Value of Deterministic Solution - fixed capacity model

Based on the fixed capacity model presented in Section 3.2, the first three comparison tests (the complete stochastic program, **VSS** and **Upgrade**) are performed for all problem instances and the aggregated results are shown in Figure 2. The bars show the losses produced by the **VSS** 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 respectively, out of the ten instances. The ***Avg.Loss*** are the mean losses the two tests produce across all the instances.

From the comparisons in Figure 2 we can see that although losses can go as high as nearly 55%, the deterministic solutions produce an average loss of around 20%, which is rather small compared with some other stochastic network design problems (see Thapalia et al. 2012a; Maggioni and Wallace 2012). Moreover, with extra services, the deterministic solutions can be greatly improved. The **Upgrade** test 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 above, we characterize demand stochasticity for each commodity using a symmetric distribution, which is replaced by its mean (i.e.,the 50^{th} percentile of the distribution) in the deterministic case. In those scenarios where some demands cannot be satisfied with the deterministic design, the more expensive ad hoc capacity must be used, which translates into the losses reflected in Figure 2 for **VSS**: about 20% on average and 55% 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 fact that all losses for **Upgrade** are extremely small therefore shows the upgradeability of deterministic

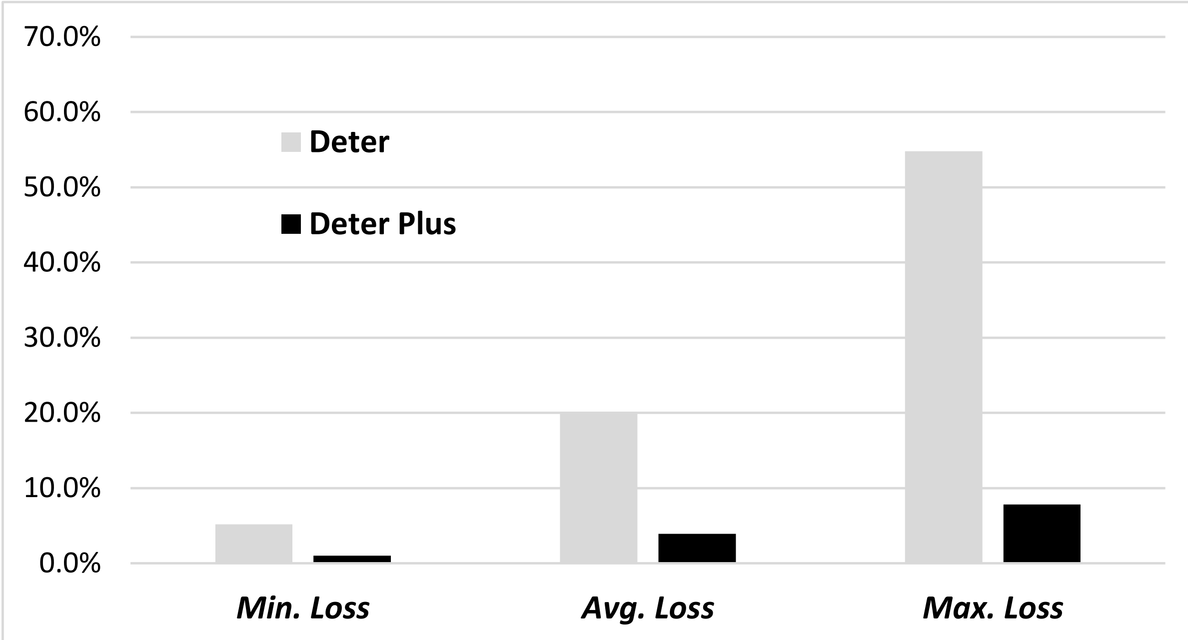

Figure 2: Comparison of performances of the **VSS** and **Upgrade** tests in the fixed capacity model. Results are measured by minimum, average and maximum losses relative to the stochastic (optimal) solution.

solutions.

In our experiments, losses of the deterministic designs in the stochastic environment are primarily caused by insufficient capacities. We therefore test for deterministic designs produced with the demand of each commodity taking the value of the 75^{th} percentile of its corresponding distribution. This is common practice in many industries. Our results show that when using the 75^{th} percentile, the average loss from using the deterministic solution in the stochastic environment (test **VSS**) drops from 19.96% to 11.23%. The **Upgrade** tests are also performed for the 75^{th} percentile cases, and an average loss of 2.55% is observed, which also indicates upgradeability of the deterministic solutions. The detailed results are reported in Table 1 in Section 5.4.

It is important to notice that the absolute values of these losses are affected by the cost structure in the network, 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. 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 setting. Instead, we focus more on the comparison of relative performances in the different tests.

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 big MIP when we allow other services to be opened. However, 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, or even 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.

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 upgradeable, which is also illustrated by Maggioni and Wallace and which underlines the interest of our results.

5.2 Value of Deterministic Solution - variable capacity model

For the model with variable capacity, it is also possible to test for **Skeleton** in addition to the tests 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 run a MIP to set the capacities $X_{ij;t}$ and distribute the flows $Y^s_{ij;t;k}$.

The same ten instances are used to perform the tests here. We compare the performances of the following three types of solutions in the stochastic environment: the deterministic solution (**VSS**), the deterministic skeleton with updated capacities (**Skeleton**) and the deterministic design with extra services (**Upgrade**). The alternative tests with the deterministic demand taking the value of the 75^{th} percentile of the associated distribution are also performed, whose results are reported in Table 1 in Section 5.4. 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 shows that the deterministic solution (**VSS**) is quite bad in the stochastic setting, while **Skeleton** behaves much better. Although the maximum loss of **Skeleton** is still high (over 40%), its average loss (around 15%) is quite acceptable. On the other hand, in the **VSS** tests, the average loss goes over 55%, and even the minimum loss

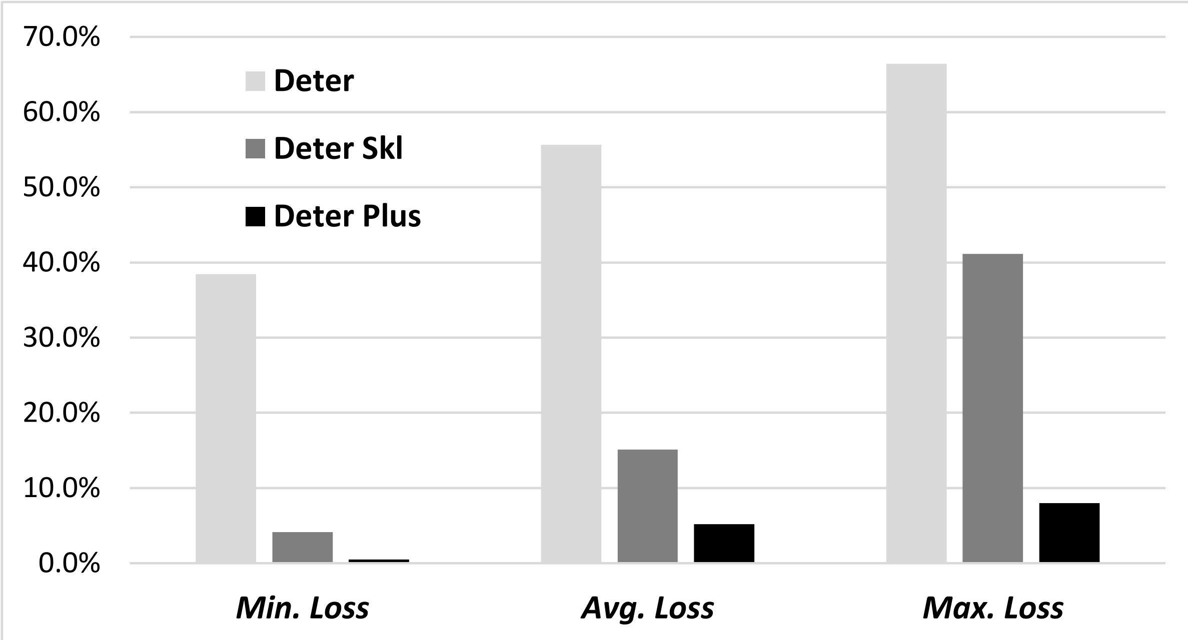

Figure 3: Comparison of **VSS**, **Skeleton** and **Upgrade** designs for the variable capacity model. Results are measured by minimum, average and maximum losses relative to the stochastic (optimal) solution.

is nearly 40%, when the deterministic solution is used directly. Again, the **Upgrade** approach offers the best performance but its computational effort could be high.

The results show that, in general, the deterministic solution does not handle well demand uncertainty (compared to the fixed-capacity case) when the capacity is not fixed a priori, but inheriting the skeleton of the deterministic solution is beneficial in most circumstances. This is very well illustrated by comparing the **VSS** performances in Figures 2 and 3. This may be explained by the possibility in the variable-capacity model to closely adjust the supplied capacity to the 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 come at the price of much ad-hoc capacity used. The results reported in Table 1 are extremely telling in this context, the performance of the **VSS** approach improving dramatically (threefold) 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 that situation, as the improvement of **Skeleton** is less important. Notice that the last observation points to the fact that this approach could be more "forgiving" of a bad demand estimation. On the other hand, the performance of **Upgrade** is fundamentally constant.

5.3 Structural Differences

We now investigate the structural differences between the solutions obtained using the proposed tests and try to find out if any "structural improvements" are taking place when the performances of the solutions are getting better. This is in order to gain insights as to why and how the deterministic solution may be upgraded and therefore of value to the stochastic problem.

Inspired by the beneficial features discovered by Lium et al. (2009) where, in particular, more hub-and-spoke structures are observed after demand stochasticity is explicitly considered, we have also observed in our experiments more consolidation activities in stochastic solutions compared to their deterministic counterparts. Therefore, if the consolidation levels of such as **VSS**, **Skeleton**, **Upgrade** and the optimal solution can be quantitatively measured, we may find a correlation between the level of consolidation and the performance of the associated solution.

However, to precisely define the "the level of consolidation" allowed by a design is difficult. It is thus hard to find a straightforward way to quantitatively determine the potential of a design to allow a higher level of consolidation. Therefore, we propose a scheme to measure two substitute phenomena: the levels of *multi-path usage* and *path-sharing* when the design is used in the 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 potentially more consolidation activities should take place.

We start by measuring the levels of multi-path usage. For a given solution, we count the number of paths each commodity is using and then produce a histogram to display the frequencies (in terms of number of commodities) *with all instances added up.* For example, if we have 10 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. The counting results are presented in Figures 4 and 5 for the fixed capacity model, and Figures 6 and 7 for the variable capacity model.

Figure 4 presents the level of multi-path usage measured by commodity counts, in **VSS**, **Upgrade** and the stochastic solution cases, for the fixed capacity model. In the first case, there are 153 commodities using only one path and 7 commodities which use 2 paths to reach their destinations. The number of commodities using 2 paths rises to 54 in the stochastic case, and there are even 12 commodities using 3 paths and 1 commodity using 4 paths while the number for a single path has dropped from 153 to 93. From Figure 4 (b) we can also see a significant increase in the number of commodities using multiple paths compared to the **VSS** case, yet lower compared to the stochastic case.

If we define the levels of multiple-paths usage measured for **VSS**, **Upgrade** and

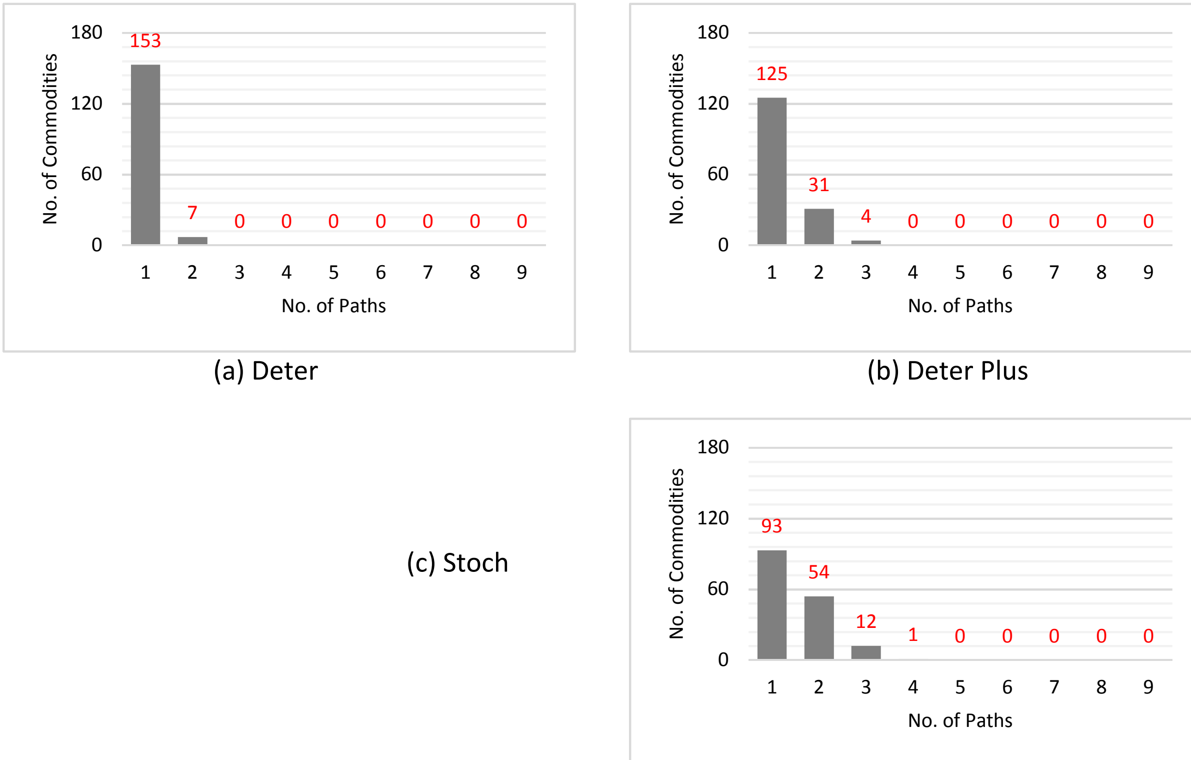

Figure 4: The level of multi-path usage for the fixed capacity model.

Stoch as low, medium and high, then when compared against their performances in the stochastic environment we see a trend (see Figure 2 for the performances of **VSS** and **Upgrade**; **Stoch**, as the optimal solution, will of course produce 0% losses). The better the solution performance, the higher the level of multiple-path usage. Considering the great improvement in performance from **VSS** to **Upgrade**, this also indicates that with some new arcs opened, the deterministic design 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 level of path-sharing. We do this by counting the number of commodities routed through each opened service. 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.

The results of path-sharing measurements for the fixed capacity model are displayed in Figure 5. In the **VSS** case, 129 arcs are shared by 2 commodities and 48 arcs are shared by 3 commodities. These two counts increase to 143 and 72 in the **Upgrade** case. In the **Stoch** case, the number of arcs shared by 2 commodities stays at a similar level (134) while the number of arcs shared by 3 commodities increases further to 93, and the

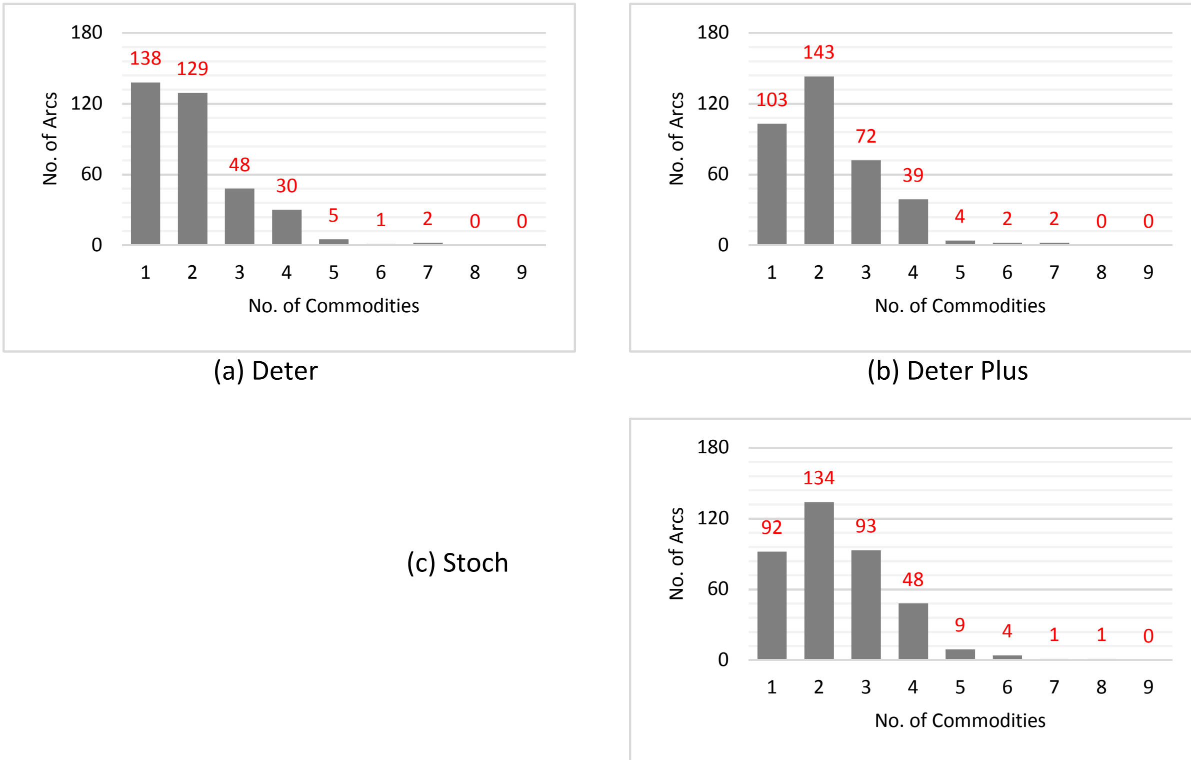

Figure 5: The level of path-sharing for the fixed capacity model.

number of arcs shared by 4 commodities reaches 48. In general, we can see a right shift of the frequency curve, from **VSS** to **Upgrade** and then to the stochastic case, while the performance 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 most of the service selection decisions of the stochastic solution are shared with the deterministic counterpart, but it includes additional arcs (services) 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 can be on par with solving the original stochastic program. However, if a heuristic approach is used to obtain the solution to the deterministic version of the problem, we may already have some potentially useful information to start with. For example, one can target those arcs which are not part of the final solution but had the longest stay inside the incumbent solutions during the process of the heuristic, or, have entered the candidate list with the highest frequencies.

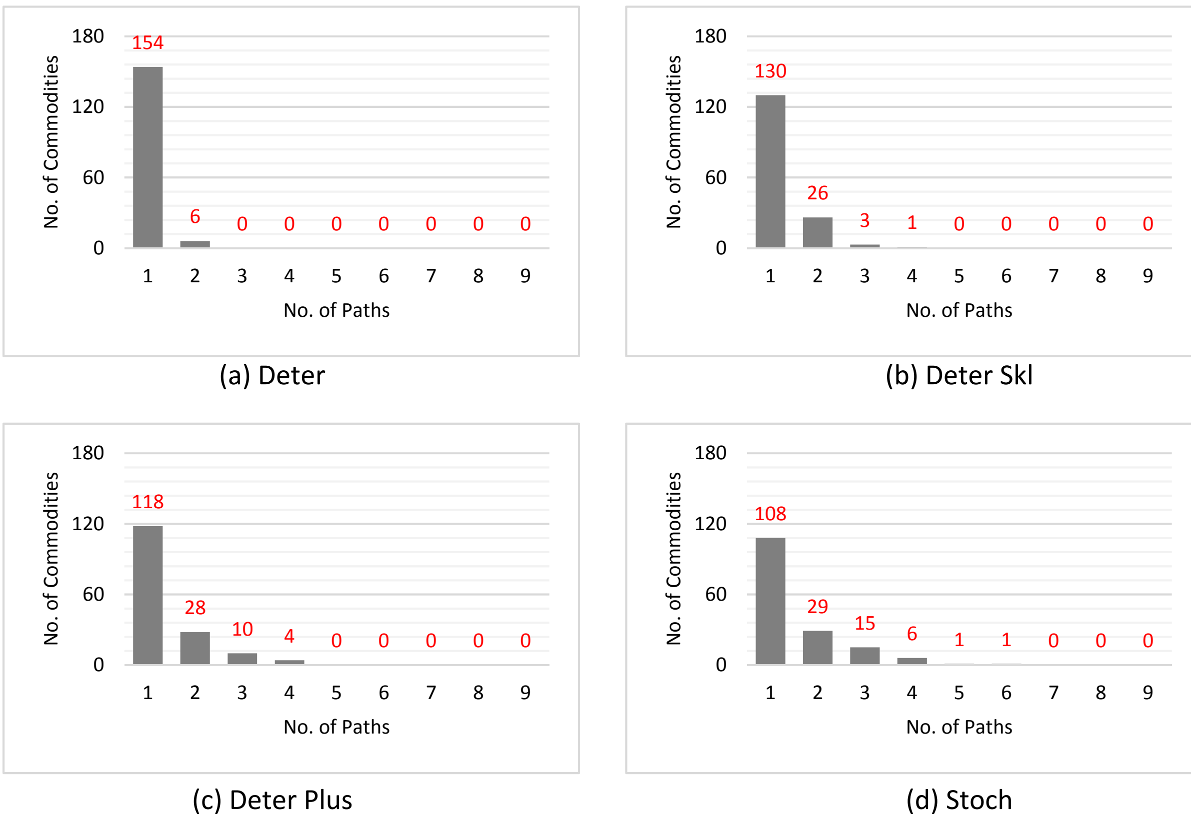

Figure 6: The level of multi-path usage for the variable capacity model.

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 level of multiple-paths usage and path-sharing it has. This is clearly visible from the numbers and the charts.

But if we consider the changes from **VSS** 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 **VSS** to **Upgrade** in the fixed capacity model.

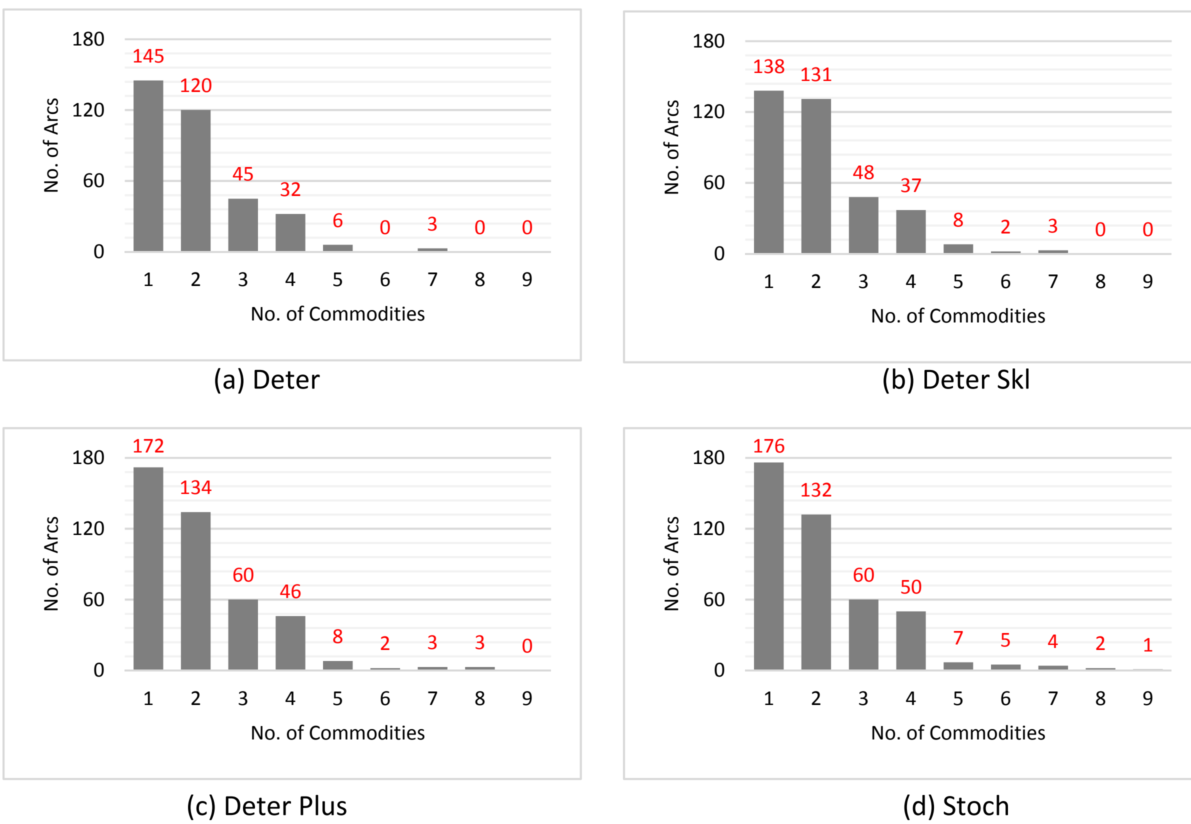

Figure 7: The level of path-sharing for the variable capacity model.

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 (how much smaller depends on the problem instance and its ratio of total demand to the total capacity that may be offered on the service network). Then, the possibility of finding a new path for a given commodity depends on whether or not there happens to be another combination of services (apart from the deterministic ones) 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 on this path. In Figure 6, we see the number of commodities using two paths quadrupled from 6 to 26.

A noteworthy observation is that the better the solution performs, the higher the levels of multiple-path usage and path-sharing are, but *not* vice versa. 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.

5.4 Effect of Continuous Capacities

We now turn to the continuous-capacity case for the variable capacity model, relaxing the integrality constraints on the capacities. 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 (which yields well-performing designs based on deterministic solutions) is much lower if 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.

The performances of the **VSS**, **Skeleton** and **Upgrade** approaches for the variable capacity model with integer and continuous capacities, together with the results from the

Table 1: Average loss in the stochastic environment

Model and Parameter Setting	Average Loss VSS	Skeleton	Upgrade
Fixed Capacity, 50^{th}	19.96%		3.91%
Fixed Capacity, 75^{th}	11.23%		2.55%
Variable Integer Capacity, 50^{th}	55.67%	15.14%	5.18%
Variable Integer Capacity, 75^{th}	12.63%	8.22%	4.37%
Variable Continuous Capacity, 50^{th}	57.77%	15.90%	4.96%
Variable Continuous Capacity, 75^{th}	13.48%	9.09%	4.72%

fixed capacity model, are displayed in Table 1. Test results with deterministic demands set at the 50^{th} and 75^{th} percentiles of their corresponding distributions are shown. The same instances are used for every row in the table.

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

5.5 Asset Balance Considerations

Table 2: Average loss in the stochastic environment with asset balance

Model and Parameter Setting **(with asset balance)**	Average Loss **VSS**	**Skeleton**	**Upgrade**
Fixed Capacity, 50^{th}	18.61%		4.62%
Fixed Capacity, 75^{th}	9.01%		3.77%
Variable Integer Capacity, 50^{th}	27.19%	11.91%	5.03%
Variable Integer Capacity, 75^{th}	8.47%	6.68%	3.85%
Variable Continuous Capacity, 50^{th}	26.36%	12.80%	4.75%
Variable Continuous Capacity, 75^{th}	9.55 %	6.20%	4.13%

Table 2 displays the results of the experimentation performed with the modified formulations with the asset-balance constraints 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, **VSS** (75^{th}) produces average losses that are all less than 10%. The **Skeleton** method can further improve the performance of the solution with not much computational efforts: a much smaller MIP for the integer capacity case and an LP for the continuous capacity case, both on a reduced skeleton network.

By comparing the numbers in Tables 1 and 2, we observe that, in general, the average losses for models with asset balance (Table 2) 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 happen to be able to pick up the unmet demand of some commodity serviced by another ship, in some scenarios in which there is a surge in demand of that commodity.

6 Conclusion

In this paper, we have discussed the value and the upgradeability 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 much better performing ones to the stochastic problem.

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 efforts. In particular, when the capacities are continuous, the **Skeleton** method becomes an LP on a reduced skeleton network. We also show that it is a better practice to use the 75^{th} percentile of the random demands when obtaining the deterministic solutions.

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 upgradeability of deterministic solutions in other network design problems. As mentioned earlier, such upgradeability 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 upgradeability 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.

References

Andersen, J., Christiansen, M., Crainic, T. G., and Gronhaug, R. (2011). Branch and price for service network design with asset management constraints. *Transportation Science*, 45(1):33–49.

Andersen, J., Crainic, T. G., and Christiansen, M. (2009a). Service network design with asset management: Formulations and comparative analyses. *Transportation Research Part C: Emerging Technologies*, 17(2):197–207.

Andersen, J., Crainic, T. G., and Christiansen, M. (2009b). Service network design with management and coordination of multiple fleets. *European Journal of Operational Research*, 193(2):377–389.

Armacost, A., Barnhart, C., and Ware, K. (2002). Composite variable formulations for express shipment service network design. *Transportation Science*, 36(1):1–20.

Bai, R., Wallace, S. W., Li, J., and Chong, A. Y. (2014). Stochastic service network design with rerouting. *Transportation Research B: Methodological*, 60:50–65.

Barnhart, C. and Schneur, R. (1996). Network design for express freight service. *Operations Research*, 44(6):852–863.

Birge, J. R. (1982). The value of the stochastic solution in stochastic linear programming with fixed recourse. *Mathematical Programming*, 24:314–325.

Buedenbender, K., Grünert, T., and Sebastian, H.-J. (2000). A hybrid tabu search/branch and bound algorithm for the direct flight network design problem. *Transportation Science*, 34(4):364–380.

Christiansen, M., Fagerholt, K., Nygreen, B., and Ronen, D. (2007). Maritime transportation. In Barnhart, C. and Laporte, G., editors, *Transportation*, volume 14 of *Handbooks in Operations Research and Management Science*, chapter 4, pages 189–284. North-Holland, Amsterdam.

Christiansen, M., Fagerholt, K., and Ronen, D. (2004). Ship routing and scheduling: Status and perspectives. *Transportation Science*, 38(1):1–18.

Cordeau, J.-F., Toth, P., and Vigo, D. (1998). A survey of optimization models for train routing and scheduling. *Transportation Science*, 32(4):380–404.

Crainic, T. G. (2000). Service network design in freight transportation. *European Journal of Operational Research*, 122(2):272–288.

Crainic, T. G. (2003). Long haul freight transportation. In Hall, R., editor, *Handbook of Transportation Science*, pages 451–516. Kluwer Academic Publishers, Norwell, MA, second edition.

Crainic, T. G., Dejax, P., and Delorme, L. (1989). Models for multimode multicommodity location problems with interdepot balancing requirements. *Annals of Operations Research*, 18:279–302.

Crainic, T. G., Ferland, J.-A., and Rousseau, J.-M. (1984). A tactical planning model for rail freight transportation. *Transportation Science*, 18(2):165–184.

Crainic, T. G., Fu, X., Gendreau, M., Rei, W., and Wallace, S. W. (2011). Progressive hedging-based metaheuristics for stochastic network design. *Networks*, 58(2):114–124.

Crainic, T. G. and Kim, K. H. (2007). Intermodal transportation. In Barnhart, C. and Laporte, G., editors, *Transportation*, volume 14 of *Handbooks in Operations Research and Management Science*, chapter 8, pages 467–537. North-Holland, Amsterdam.

Crainic, T. G. and Laporte, G. (1997). *Design and Operation of Civil and Environmental Engineering S ystems*, chapter Planning Models for Freight Transportation, pages 343–394. Wiley-Interscience, New York.

Crainic, T. G. and Rousseau, J.-M. (1986). Multicommodity, multimode freight transportation: A general modeling and algorithmic framework for the service network design problem. *Transportation Research B: Methodology*, 20B:225–242.

Crainic, T. G. and Roy, J. (1988). O.R. Tools for tactical freight transportation planning. *European Journal of Operational Research*, 33(3):290–297.

Dejax, P. and Crainic, T. G. (1987). A review of empty flows and fleet management models in freight transportation. *Transportation Science*, 21(4):227–247.

Delorme, L., Roy, J., and Rousseau, J.-M. (1988). Motor-carrier operation planning models: A state of the art. In Bianco, L. and Bella, A. L., editors, *Freight Transport Planning and Logistics*, pages 510–545. Springer-Verlag, Berlin.

Farvolden, J. and Powell, W. (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, J. and Powell, W. (1994). Subgradient methods for the service network design problem. *Transportation Science*, 28(3):256–272.

Farvolden, J., Powell, W., and Lustig, I. (1992). A Primal Partitioning Solution for the Arc–Chain Formulation of a Multicommodity Network Flow Problem. *Operations Research*, 41(4):669–694.

Gorman, M. (1998a). An application of genetic and tabu searches to the freight railroad operating plan problem. *Annals of Operations Research*, 78:51–69.

Gorman, M. (1998b). Santa fe railway uses an operating-plan model to improve its service design. *Interfaces*, 28(4):1–12.

Grünert, T. and Sebastian, H.-J. (2000). Planning models for long-haul operations of postal and express shipment companies. *European Journal of Operational Research*, 122:289–309.

Grünert, T., Sebastian, H.-J., and Thäerigen, M. (1999). The design of a letter-mail transportation network by intelligent techniques. In Sprague, R., editor, *Proceedings Hawaii International Conference on System Sciences 32*.

Haghani, A. (1989). Formulation and solution of combined train routing and makeup, and empty car distribution model. *Transportation Research B: Methodological*, 23B(6):433–452.

Higle, J. L. and Wallace, S. W. (2003). Sensitivity analysis and uncertainty in linear programming. *Interfaces*, 33:53–60.

Hoff, A., Lium, A.-G., Løkketangen, A., and Crainic, T. G. (2010). A metaheuristic for stochastic service network design. *Journal of Heuristics*, 16(5):653–679.

Høyland, K., Kaut, M., and Wallace, S. W. (2003). A heuristic for moment-matching scenario generation. *Computational Optimization and Applications*, 24(2–3):169–185.

Kall, P. and Wallace, S. W. (1994). *Stochastic Programming.* John Wiley & Sons, Chichester.

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

Keaton, M. (1989). Designing optimal railroad operating plans: Lagrangian relaxation and heuristic approaches. *Transportation Research B: Methodological*, 23B(6):415–431.

Keaton, M. (1991). Service–cost tradeoffs for carload freight traffic in the u.s. rail industry. *Transportation Research A: Policy and Practice*, 25A(6):363–374.

Keaton, M. (1992). Designing optimal railroad operating plans: A dual adjustment method for implementing lagrangian relaxation. *Transportation Science*, 26:263–279.

Kim, D., Barnhart, C., Ware, K., and Reinhardt, G. (1999). Multimodal express package delivery: A service network design application. *Transportation Science*, 33(4):391–407.

King, A. J. and Wallace, S. W. (2012). *Modeling with Stochastic Programming.* Springer Series in Operations Research and Financial Engineering. Springer.

Lamar, B., Sheffi, Y., and Powell, W. (1990). A capacity improvement lower bound for fixed charge network design problems. *Operations Research*, 38(4):704–710.

Lium, A.-G., Crainic, T. G., and Wallace, S. W. (2007). Correlations in stochastic programming: A case from stochastic service network design. *Asia-Pacific Journal of Operational Research*, 24(2):161–179.

Lium, A.-G., Crainic, T. G., and Wallace, S. W. (2009). A study of demand stochasticity in stochastic network design. *Transportation Science*, 43(2):144–157.

Maggioni, F. and Wallace, S. W. (2012). Analyzing the quality of the expected value solution in stochastic programming. *Annals of Operations Research*, 200(1).

Newton, H. (1996). *Network Design Under Budget Constraints with Application to the Railroad Blocking Problem*. PhD thesis, Industrial and Systems Engineering, Auburn University, Auburn, Alabama, U.S.A.

Newton, H., Barnhart, C., and Vance, P. (1998). Constructing railroad blocking plans to minimize handling costs. *Transportation Science*, 32(4):330–345.

Pedersen, M., Crainic, T. G., and Madsen, O. (2009). Models and tabu search metaheuristics for service network design with asset-balance requirements. *Transportation Science*, 43(2):158–177.

Powell, W. (1986). A local improvement heuristic for the design of less-than-truckload motor carrier networks. *Transportation Science*, 20(4):246–357.

Powell, W. and Sheffi, Y. (1983). The load-planning problem of motor carriers: Problem description and a proposed solution approach. *Transportation Research A: Policy and Practice*, 17(6):471–480.

Powell, W. and Sheffi, Y. (1986). Interactive optimization for motor carrier load planning. *Journal of Business Logistics*, 7(2):64–90.

Powell, W. and Sheffi, Y. (1989). Design and implementation of an interactive optimization system for the network design in the motor carrier industry. *Operations Research*, 37(1):12–29.

Roy, J. and Delorme, L. (1989). Netplan: A network optimization model for tactical planning in the less-than-truckload motor-carrier industry. *INFOR*, 27(1):22–35.

Sharypova, M., Crainic, T. G., van Woensel, T., and Fransoo, J. (2012). Scheduled Service Network Design with Synchronization and Transshipment Constraints for Intermodal Container Transportation Networks. Publication CIRRELT-2012-77, Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport, Montréal, QC, Canada.

Smilowitz, K., Atamtürk, A., and Daganzo, C. (2003). Deferred item and vehicle routing within integrated networks. *Transportation Research Part E: Logistics and Transportation*, 39:305–323.

Thapalia, B. K., Crainic, T. G., Kaut, M., and Wallace, S. W. (2011). Single-commodity stochastic network design with multiple sources and sinks. *INFOR: Information Systems and Operational Research*, 49(3):193–211.

Thapalia, B. K., Crainic, T. G., Kaut, M., and Wallace, S. W. (2012a). Single-commodity network design with random edge capacities. *European Journal of Operational Research*, 220(2):394–403.

Thapalia, B. K., Crainic, T. G., Kaut, M., and Wallace, S. W. (2012b). Single source single-commodity stochastic network design. *Computational Management Science*, 9(1):139–160. Special issue on 'Optimal decision making under uncertainty'.

Wallace, S. (2000). Decision making under uncertainty: Is sensitivity analysis of any use? *Operations Research*, 48:20–25.

Wallace, S. W. (2010). Stochastic programming and the option to do it differently. *Annals of Operations Research*, 177(1):3–8.

Zhu, E., Crainic, T. G., and Gendreau, M. (2013). Scheduled service network design for freight rail transportation. *Operations Research*, 62(2):383–400.