CIRRELT

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

Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation

A Study of Demand Stochasticity in Service Network Design

Arnt-Gunnar Lium
Teodor Gabriel Crainic
Stein W. Wallace

July 2007

CIRRELT-2007-22

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

Bureaux de Québec :
Université Laval
Pavillon Palasis-Prince, local 2642
Québec (Québec)
Canada G1K 7P4
Téléphone : 418 656-2073
Télécopie : 418 656-2624

www.cirrelt.ca

A Study of Demand Stochasticity in Service Network Design

Arnt-Gunnar Lium[1], Teodor Gabriel Crainic[2,*], Stein W. Wallace[1]

[1] Molde University College, P.O. Box 2110, NO-6402 Molde, Norway

[2] Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT), Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada H3C 3J7 and Chaire de recherche industrielle du CRSNG en management logistique, ESG, Université du Québec à Montréal, C.P. 8888, succursale Centre-ville, Montréal, Canada H3C 3P8

Abstract. The objective of this paper is to investigate the importance of taking uncertainty explicitly into account for service network design. We study how solutions based on uncertain demand differ from solutions based on deterministic demand and provide qualitative descriptions of the structural differences. Some of these structural differences provide a hedge against uncertainty by using consolidation. This way we get consolidation as output from the model rather than as an a priori required property. Service networks with such properties are robust, as seen by the customers, by providing operational flexibility.

Keywords. Stochastic programming, scheduled service network design, flexibility, robustness.

Acknowledgements. We would like to thank Michal Kaut for modifying his scenario generation procedure to suite our needs. Stein W. Wallace was supported by grant 156315/530 from the Research Council of Norway. Partial funding for this research has been provided by Natural Sciences and Engineering Research Council of Canada (NSERC) through its Discovery Grants and Chairs and Faculty Support programs. While working on this project, Dr. Crainic was Adjunct Professor at the Molde University College, Molde, Norway, and member of CIRRELT.

---

* Corresponding author: crainic.teodor@uqam.ca

Dépôt légal – Bibliothèque nationale du Québec,
Bibliothèque nationale du Canada, 2007

1 Introduction

Service network design formulations are generally associated with tactical planning of operations for consolidation carriers, that is, carriers letting more than one commodity (or passenger) share the capacity of their vehicles. The goal of the planning process is to determine the routes on which services will be offered, the type of services that will be offered, as well as the frequency and schedules of services. The selected services and the schedule constitute a transportation or load plan. The schedule is also published for the benefit of the potential users of the system. In building the plan, one aims for an efficient operation in terms of total system cost, given the available resources, the known demands, and the level of service quality that the carrier desires to achieve.

There is quite a significant body of literature on the subject. Service network design formulations have been proposed for many types of multimodal (e.g., Crainic and Rousseau 1986; Crainic and Roy 1988) or single mode transportation: rail (e.g., Crainic, Ferland, and Rousseau 1984; Haghani 1989; Gorman 1998a, 1998b; Keaton 1989, 1991, and 1992; Newton 1996; Newton, Barnhart, and Vance 1998); less-than-truckload (LTL; e.g., Delorme and Roy 1989; Powell and Sheffi 1983, 1986, 1989; Powell, 1986a; Lamar, Sheffi, and Powell 1990; Farvolden and Powell 1991, 1994), maritime navigation (Christiansen, Fagerholt, and Ronen 2004), express courier services (e.g., Grünert and Sebastian 2000; Grünert, Sebastian, and Thärigen 1999; Buedenbender, Grünert, and Sebastian 2000; Armacost, Barnhart, and Ware 2002; Barnhart and Schneur 1996; Smilowitz, Atamtürk, and Daganzo 2003); and so on. See Christiansen *et al.* (2007), Cordeau, Toth, and Vigo (1998), Crainic (2000; 2003), and Crainic and Kim (2007) for recent reviews.

The service network literature and, to the best of our knowledge, the systems implemented at various carriers assume complete knowledge. That is, all these formulations are deterministic. This is not to say that the research community and the transportation professionals ignore the uncertainties that accompany actual operations. On the contrary, several papers clearly state that one works with forecasted demand, that the transportation plan is built for a “regular” operation period, usually a week, and that the plan is to be adjusted during actual operations to account for too high or too low demand.

This approach may be cast as a two-phase procedure. The first phase solves the service network design using point forecasts of the demand. In the second phase the uncertainty is resolved, the actual values of demand are observed, and the plan is modified accordingly. The first phase is performed once for the contemplated time horizon, while the second one is repeated every period. (The same treatment is equally applied to other critical factors, such as transportation times, but for simplicity, we will restrict the scope of this paper to demand only.)

One may then ask: What is lost by not integrating information about the stochastic nature of demand directly into the tactical planning methodology? Would the integration of such information lead to different service design patterns, either in the services selected or the consolidation strategies used? Would the resulting transportation plan and schedule be more robust than what we presently obtain by providing more flexibility in the day-to-day operations? And if the answer is yes, can we characterize what it is that creates robustness? Why and how are solutions to stochastic models better than those from deterministic ones? Answers to these two questions are the main contributions of this paper.

The objective of this paper is to contribute to answering such questions. These questions are rather general in nature. They are relevant to all situations where we use deterministic optimization models in contexts where stochastics is obviously present. We need to ask: Will solving the deterministic model, rather than the stochastic one, result in significant errors, or are we just talking about minor adjustments? The computational as well as modeling efforts needed to address the stochastic case are substantial, so we need to be sure that it is worth while. Or at least – and that is the setting of this paper – that the potential errors must be such that checking the actual effects must be worth the cost.

Keywords here are flexibility and robustness. By robustness we understand the ability to withstand random effects, while flexibility implies the ability to adjust to them. A tree is robust while a straw is flexible when facing high winds. In our context we wish robust schedules as seen by the customers. Specifically, that means the ability to operate the specified origin-destination pairs regardless of changes in demand. This will be achieved by flexibility in how commodities are sent in day-to-day operations. We say that a schedule is more robust, the more cost effectively it deals with these varying demands, hence the lower expected costs it leads to. The choice that schedules should be robust and the flows of commodities from origin to destination are to be flexible is a conscious choice the modeler has to make. Schedules do not simply happen to be robust, the requirement must be put into the model. For the model presented in this paper, the design coming from the deterministic model is feasible, but not necessarily optimal, in the stochastic model. Hence, by construction, the stochastic model is best in the stochastic framework.

There are many objections to this simple argument. The two most important are

- What happens if our estimation of (stochastic) future demand is off? Is it not safer to use expected demand?

- Even if the stochastic solution obviously is better in the described setting, are we sure the difference is of importance?

In the end, we cannot prove that a given stochastic model is better than its deterministic counterpart. Obviously, if the data used in the model, or the model itself, is totally off the mark, a deterministic model can yield results which, in the real world, are better than those of a stochastic model.

What we can do, though, is to assume that our description of the problem is correct, and then investigate how the two models behave relative to one another. That is the purpose of this paper. The first important observation is that the expected behavior of the solution coming from a deterministic model can be arbitrarily worse than the one coming from the stochastic model. This is discussed in Wallace (2000) and Higle and Wallace (2003). It is there shown that the solutions of stochastic models are structurally different from those of deterministic models. Simply stated, flexibility costs money, and in a deterministic world, you will therefore never buy flexibility. In the language of this paper: In a deterministic world, you will fit your network design perfectly to the future demand you are shown, and not at all (of course) take into account that there might be other possible futures where it may fit very badly.

In most service network design models, decisions and structures are enforced upon the solution before any optimization takes place at all. It is common to decide a priori whether to use consolidation or not, and if this consolidation should be performed in some kind of hub-and-spoke system or in other ways. There is no one-to-one relationship between consolidation and hub-and-spoke. While a hub-and-spoke system implies consolidation, the reverse is not true. Typically, these strategic decisions are based on the idea that operating a hub-and-spoke network will bring consolidation, making operations more efficient due to, for example, economies of scale. One might then ask, will this type of network that uses consolidation provide us with some sort of robustness when dealing with uncertain demand? To test this, we propose a model that does not have any a priori bindings with respect to the structure of the solution, but rather lets the structure be a result of the optimization process with its stochastic demand. Whatever structure is obtained, we know that it will be based on quantitative aspects only, and not be biased by "rules-of-thumb" or politics, which – unfortunately – are more common than not in the transport industry.

Most textbooks in operations research – and obviously many vendors of software as well – tell us to perform "what-if-sessions" to address the uncertainty about the future. The idea is that if we study these solutions – one solution for each possible future – we have in front of us a collection of (in our case) designs that gives us the total picture of what we might do. We might for example calculate the expected behavior for each potential design, and keep the best one. Or we may try to combine them into something even better – perform a kind of "convex" combination to arrive at an overall very good (if not optimal) solution. To realize that this is a false assurance is possibly the most critical point in understanding what stochastics does to an optimization problem. This approach does not take into account that what we need is robust solutions, and what we have is a collection of non-robust ones. Even though we have *many* solutions in front of us, none of them is robust, in the sense that they are all made under assumptions of a known future. More details, with a worked out example, about this subtle but crucial point can be found in Higle and Wallace (2003).

We may also set this argument in the framework of real options. A robust design will typically contain options in the sense of investments that make the future more easy to handle. But options cost money, and we never buy an option if we know we will not need it. And if we know we shall need what the option offers, we do not buy the option, but the real object immediately. This is true unless the option comes for free. If you have three decisions in a simple investment problem: Do nothing, build a factory or buy the right to build a factory, a deterministic model will always pick one of the first two as long as there is a cost associated with the option (and in real life there is). So, the observation is: The robust solutions contain options, the deterministic ones do not, and solving many deterministic problems will not overcome that problem. We must tell the model explicitly that the future is stochastic. It will then pick up good options – if they exist – and hence provide flexibility in the operations. So, only in problems without implicit options (if they exist), will deterministic models work well. And this is why we turn to stochastic programs, despite the heavy burden that places on us with respect to both data collection and solution abilities (in addition to the more complicated modeling itself).

For a more general look at stochastic programming, see, for example, Kall and Wallace (1994) and Volume 10 of *Handbooks in Operations Research and Management Science* by Ruszczyński and Sharpiro (2003).

To initiate the study of the impact of introducing stochastic elements into service network design formulations, we take a simplified version of the problem faced by LTL carriers in which periodic schedules are built for a number of vehicles and where only the demand may vary stochastically. We chose a problem size allowing the use of standard mixed-integer software. The results clearly indicate that 1) yes, integrating stochastic elements in the service network design model is beneficial, 2) the resulting transportation plan is more robust by introducing operational flexibility, and 3) we are able to describe qualitatively what are the structural differences, that is, what are the options.

The plan of the paper is as follows. Section 2 recalls the main components of a service network design model and introduces the simplified problem we use and the associated deterministic formulation. Section 3 displays the stochastic formulation and introduces the solution strategy employed. The experimental setting, including a detailed discussion of how we generate scenarios, is presented in Section 4, while the computational results are presented in Section 5. Analyses of some of the characteristics that can be found in stochastic schedules are shown in Section 6. We conclude in Section 7.

2 A Service Network Design Case

Consolidation transportation carriers are usually organized as so-called hub-and-spoke networks, where service is offered between a much larger number of origin-destination pairs than that of direct, origin-to-destination services operated by the carrier. Low-volume demands are then moved first to an intermediate point, a hub, there to be consolidated with loads from other customers into vehicles and convoys and moved to other hubs by high frequency and capacity services. More than one consolidation-transfer operation may occur during a trip. Such an organization allows a higher quality service for all origin-destination pairs, in terms of frequency of service, and a more efficient utilization of resources (hence, lower costs). The drawback of this type of organization is increased delays due to longer routes, increased time spent in terminals, and more complex planning processes and operations.

The term *service network design* refers to the process of selecting the services and schedules to operate. This process is often performed at the tactical planning level and the collection of services and schedules are known as the transportation, or load, plan. The objective is to provide the highest level of service possible and ensure customer satisfaction, while operating efficiently and profitably.

Service network design problems thus address two types of major decisions. The first concerns the choice of the service network, that is, the selection of the routes – origin and destination terminals, physical routes, and intermediate stops – on which services will be offered and the characteristics of each service, in particular its frequency or schedule. The second major type of decision is to determine the distribution of traffic, that is, the itineraries (routes) used to move the loads of each demand: services used, terminals passed through, operations performed in these terminals, etc. Operating rules specifying, for example, how loads and vehicles may be sorted and consolidated, are sometimes specified at particular terminals and become part of the service network (this is the case, in particular, for rail carriers). The service network thus specifies the movements through space and time of the vehicles and convoys considered, while itineraries move freight from origins to destinations and determine the volumes of commodities that flow on the services and through the terminals of the service network.

Several efforts have been directed towards the formulation of service network design models and static and time-dependent formulations have been proposed. The former assumes that demand does not vary during the planning period or that the distribution of departures is known (typically uniform) and only the service selection and frequencies are of interest. The time dimension of the service network is then implicitly considered through the definition of services and the inter-service operations at terminals. Time-dependent formulations include an explicit representation of movements in time and usually target the planning of schedules to support decisions related to when services are dispatched. This is usually achieved by representing the operations of the system over a certain number of time periods by using a space-time network. In such a structure, the representation of the physical network is replicated in each period. Starting from its origin in a given period, a service arrives (and leaves, in the case of intermediary stops) later at other terminals. Services thus generate temporal service links between different terminals at different time periods. Temporal links that connect two representations of the same terminal at two different time periods may represent the time required by terminal activities or the vehicles and freight waiting for the next departure. Additional arcs may be used to capture penalties for arriving too early or too late.

2.1 The initial deterministic model

The initial model takes the form of a deterministic, fixed cost, capacitated, multicommodity network design formulation. Integer-valued decision variables are used to represent service selection decisions, while continuous product-specific variables capture the commodity flows. Fixed costs are associated with the inclusion of services into the plan. Costs that vary with the intensity of service and commodity traffic are associated with the movements of commodities and services. The goal is to minimize the total system cost, or to maximize the net profit, under constraints enforcing demand, service, and operation rules and goals.

For the study reported in this paper, we build a version of a multi-period service network design model inspired by the less-than-truckload motor carrier case. Several simplifying assumptions are made:

- We consider a homogeneous fleet of capacitated vehicles and no restrictions on how many vehicles are used;
- The transport movements require one period, while terminal operations are instantaneous (within the period);
- Demand cannot be delivered later than the due date, but may arrive earlier;
- There is a (fixed) cost associated with operating a vehicle (service), but no cost is associated with moving freight; That is, truck movements cost the same whether they move loaded or empty;
- There are no costs associated with time delays or terminal operations;
- The plan is repeated periodically;

- No hub-and-spoke structure is assumed for the service network, i.e., all terminals and services have similar characteristics.

The resulting model is sufficiently simple to be solved by commercial off-the- shelf mixed-integer software for moderately sized problem instances, while retaining the main characteristics of service network design problems. It thus offers a good environment for an explicit study of the effects of uncertainty.

The space-time network is built by repeating the set of nodes (terminals) $\mathcal{N}$ in each of the periods $t = 0, \dots, T-1$. Each arc (i, j) represents either a service, if $i \neq j$, or a holding activity if $i = j$. A cost c_{ij} is associated to each arc (i, j), equal to the cost of driving a truck from terminal i to j if $i \neq j$, or to the cost of holding a truck at terminal i if $i = j$. For each commodity $k \in \mathcal{K}$, we define its demand $\delta(k)$, origin $o(k)$, destination $d(k)$, and the points in time $\sigma(k)$ and $\tau(k)$ when it becomes available at its origin and must be delivered (at the latest) at its destination, respectively. The truck capacity is denoted M (same units as for demand).

The decision variables then represent the commodity distribution decisions over the selected service network, and the frequencies of the selected services (a zero frequency indicating the service was not selected):

$Y_{ij}^t(k)$: Amount of commodity k going from terminal i at time t to terminal j at time $t+1$, for $\{i = o(k), t = \sigma(k), \forall j\} \cup \{j = d(k), t = \tau(k), \forall i\} \cup \{\sigma(k) < t < \tau(k), \forall i, j\}$;

X_{ij}^t : Number of trucks from node i at time t to in node j at time $t+1$, $\forall i, j, t$.

The basic deterministic service network design formulation may then be written as:

$$\min \sum_{i\in\mathcal{N}} \sum_{j\in\mathcal{N}} \sum_{t=0}^{T-1} c_{ij} X_{ij}^t \tag{1}$$

$$t = 0, \ldots, T-2: \quad \sum_{i\in\mathcal{N}} X_{ij}^t = \sum_{i\in\mathcal{N}} X_{ji}^{t+1} \quad \forall j \in \mathcal{N} \tag{2}$$

$$t = 0, \ldots, T-1: \quad X_{ij}^t \geq 0 \text{ and integer} \quad \forall i, j \in \mathcal{N} \tag{3}$$

$$(o(k), \sigma(k)): \quad \sum_{i\in\mathcal{N}} Y_{o(k)i}^{\sigma(k)}(k) = \delta(k) \qquad \forall k \in \mathcal{K} \tag{4}$$

$$(j, \sigma(k)+1): \quad Y_{o(k)j}^{\sigma(k)}(k) = \sum_{i\in\mathcal{N}} Y_{ji}^{\sigma(k)+1}(k) \qquad \forall j \in \mathcal{N}, \forall k \in \mathcal{K} \tag{5}$$

$$(1 + \sigma(k) < t < \tau(k) - 1): \quad \sum_{i\in\mathcal{N}} Y_{ij}^{t-1}(k) = \sum_{i\in\mathcal{N}} Y_{ji}^{t}(k) \qquad \forall j \in \mathcal{N}, \forall k \in \mathcal{K} \tag{6}$$

$$(j, \tau(k)-1): \quad \sum_{i\in\mathcal{N}} Y_{ij}^{\tau(k)-2}(k) = Y_{j,d(k)}^{\tau(k)-1}(k) \qquad \forall j \in \mathcal{N}, \forall k \in \mathcal{K} \tag{7}$$

$$(d(k), \tau(k)): \quad \sum_{i\in\mathcal{N}} Y_{i,d(k)}^{\tau(k)-1}(k) = \delta(k) \qquad \forall k \in \mathcal{K} \tag{8}$$

$$t = 0, \ldots, T-1: \quad \sum_{k\in\mathcal{K}} Y_{ij}^t(k) \leq M X_{ij}^t \qquad \forall i, j \in \mathcal{N} : i \neq j \tag{9}$$

$$t = 0, \ldots, T-1: \quad Y_{ij}^t(k) \geq 0 \qquad \forall i, j \in \mathcal{N}, \forall k \in \mathcal{K} \tag{10}$$

The objective function (1) minimizes the cost associated with vehicles moving between terminals, plus the cost of holding them at terminals. Constraints (2) enforce conservation of flow for trucks, which must be integer as indicated by constraints (3).

Equations (4)-(8) are conservation of flow constraints dealing with the flow of freight from origins to destinations at various time periods. Figure 1 provides a graphical overview of where each of these constraints apply in the space-time network. Constraints (4) are illustrated in the left-hand side of Figure 1, preserving the conservation of flow in the origin node for commodity k. Constraints (8) perform a similar function for the flow of freight at the destination node, as shown on the right-hand side of Figure 1. Equations (5) represent conservation-of-flow conditions at nodes one period after a commodity has left its origin node. The flow may come only from node $o(k)$ at time $\sigma(k)$, but may go to any node at time $\sigma(k)+1$. Similarly, equations (7) enforce conservation of flow at nodes one period before the flow arrives at its destination; The time periods where these two constraints apply are marked by (5) and (7) in Figure 1. Equations (6) are the general conservation-of-flow constraints valid from two time periods after a commodity has left its origin and up to two periods before arriving at its destination; It is denoted by (6) in the center of Figure 1. Relations (9) are the usual linking and vehicle capacity constraints. Note that commodities can be held at nodes without a truck being present (hence, $i \neq j$).

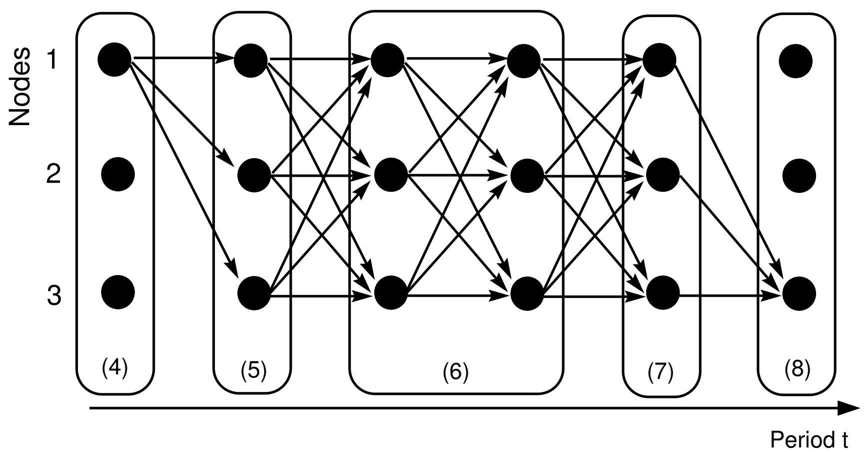

Figure 1: The relationship between constraint sets (4)-(8) and the underlying network.

2.2 Fleet capacity considerations

Most consolidation carriers experience demand variations, not only in overall quantity but also in how demand is distributed among the markets (origin-destination city pairs) served, and over time within a market. Consequently, carriers will plan for a demand pattern that represents the "regular" traffic (the demand one may reasonably expect to always be there) over an average or heavy (the average forecasted over a busy period) week. Such point forecasts of future demand are used in most service network design models to produce a plan where all demand is satisfied by using the company's own vehicles. The previous model is no exception.

Capacity does not always match demand, however. In the case of a low demand relative to the planned capacity, carriers tend to cancel services (delaying, eventually, some demand). When demand is higher than planned, carriers can act in a variety of ways: add capacity on a service, either using some of the company's own vehicles, if available, or by "renting" vehicles and/or drivers, refuse the load (a course of action very rarely chosen), delay some freight, possibly at a cost, and so on. Sometimes, such choices, and the associate cost-profit-quality of service trade-offs, are reflected in the tactical planning models. They should certainly be included when the demand variability is explicitly considered in the formulations.

We integrate this factor into our formulation by defining an *ad-hoc capacity increase*, denoted $Z(k),\ k \in \mathcal{K}$, as the amount of commodity k transported using a different vehicle from those of the fleet considered. The associated per-unit cost is b, set such that there is a reasonable balance between own transportation and taking the ad-hoc alternative.

We then update the basic formulation by replacing the objective function (1) with (11), which takes the cost of ad-hoc capacity increase into account. The constraints (4) and (8) are replaced by (12) and (13), respectively.

$$\min \sum_{i\in\mathcal{N}} \sum_{j\in\mathcal{N}} \sum_{t=0}^{T-1} c_{ij} X_{ij}^{t} + b \sum_{k\in K} Z(k) \tag{11}$$

$$(o(k),\sigma(k)) : \quad \sum_{i\in\mathcal{N}} Y_{o(k)i}^{\sigma(k)}(k) + Z(k) = \delta(k) \qquad \forall k \in \mathcal{K} \tag{12}$$

$$(d(k),\tau(k)) : \quad \sum_{i\in\mathcal{N}} Y_{i,d(k)}^{\tau(k)-1}(k) + Z(k) = \delta(k) \qquad \forall k \in \mathcal{K} \tag{13}$$

2.3 Circular schedules

In any multi-period formulation, one has to address the end-of-horizon effects. We can mitigate this problem by casting the previous model in a circular fashion. This implies that vehicles operate according to circular routes over the planning horizon considered, which is an approach to introduce fleet management issues into service network design models.

Assuming a T-period planning horizon, a circular notation means that the period following period t is $(t+1) \bmod T$, while the previous period becomes

$(t-1+T) \bmod T$. To improve readability, we present constraints according to the difference, **diff**, in number of time periods, between $\sigma(k)$ and $\tau(k)$. To keep the paper relatively short, the circular notation is introduced in the following section, together with the stochastic notation.

3 Stochastic Service Network Design

Let the deterministic demand δ be replaced by the random demand $\tilde{\delta}$, with density function $f(\delta)$. The model then becomes (the formulation assumes circular vehicle schedules, as indicated previously):

$$\min \sum_{i\in\mathcal{N}} \sum_{j\in\mathcal{N}} \sum_{t=0}^{T-1} c_{ij} X_{ij}^t + b \int Q(X,\delta) f(\delta) d\delta \tag{14}$$

$$\sum_{i\in\mathcal{N}} X_{ij}^t = \sum_{i\in\mathcal{N}} X_{ji}^{(t+1 \bmod T)} \quad t = 0,\dots,T-1, \forall j \in \mathcal{N} \tag{15}$$

$$X_{ij}^t \geq 0 \text{ and integer} \quad t = 0,\dots,T-1, \forall i,j \in \mathcal{N} \tag{16}$$

where

$$Q(X,\delta) = \min \sum_{k \in K} Z(k) \tag{17}$$

$$\sum_{i \in \mathcal{N}} Y_{o(k)i}^{\sigma(k)}(k) + Z(k) = \delta(k), \quad \forall k \in \mathcal{K} \tag{18}$$

$$(\text{diff} \geq 3): \quad Y_{o(k)j}^{\sigma(k)}(k) = \sum_{i \in \mathcal{N}} Y_{ji}^{(\sigma(k)+1) \bmod T}(k) \quad \forall j \in \mathcal{N}, \forall k \in \mathcal{K} \tag{19}$$

$$(\text{diff} \geq 4): \quad \sum_{i \in \mathcal{N}} Y_{ij}^{(t-1+T) \bmod T}(k) = \sum_{i \in \mathcal{N}} Y_{ji}^{t}(k)$$

$$t = 0, \ldots, T-1, \forall j \in \mathcal{N}, \forall k \in \mathcal{K} \tag{20}$$

$$(\text{diff} \geq 3): \quad \sum_{i \in \mathcal{N}} Y_{ij}^{(\tau(k)-2+T) \bmod T}(k) = Y_{j,d(k)}^{(\tau(k)-1+T) \bmod T}(k)$$

$$\forall j \in \mathcal{N}, \forall k \in \mathcal{K} \tag{21}$$

$$\sum_{i \in \mathcal{N}} Y_{i,d(k)}^{(\tau(k)-1+T) \bmod T}(k) + Z(k) = \delta(k), \quad \forall k \in \mathcal{K} \tag{22}$$

$$(\text{diff} = 2): \quad Y_{o(k),j}^{\sigma(k)}(k) = Y_{j,d(k)}^{(\sigma(k)+1) \bmod T}(k) \qquad \forall j \in \mathcal{N}, \forall k \in \mathcal{K} \tag{23}$$

$$\sum_{k \in \mathcal{K}} Y_{ij}^{t}(k) \leq M X_{ij}^{t}$$

$$\forall i, j \in \mathcal{N} : i \neq j, \ t = 0, \ldots, T-1 \tag{24}$$

$$Y_{ij}^{t}(k) \geq 0 \qquad \forall i, j \in \mathcal{N}, \forall k \in \mathcal{K}, t = 0, \ldots, T-1 \tag{25}$$

The objective function (14) minimizes the costs of this two-stage problem. The first stage determines how to operate the vehicles while the second stage addresses how to deal with ad-hoc capacity increase for given demand realization δ and first-stage decision X. Constraints (15) represent conservation of flow for vehicles, while (16) is a non-negativity and integrality constraint.

In the second stage, the objective function (17) minimizes the cost of commodities sent using ad-hoc capacity increase. Equations (18) state that the demand of a commodity at its origin is served either by the company's own vehicles or by some ad-hoc capacity increase. Equations (22) are the equivalent constraints for a commodity at its destination. Constraints (20) enforce the conservation of flow of commodities when there are four or more time periods between $\sigma(k)$ and $\tau(k)$. Similarly, equations (19) and (21) are conservation of flow constraints for commodity k one period after it becomes available at its origin and one period before it has to arrive at its destination. Constraints (23) represents the conservation of flow for commodity k when there are only 2 time periods between $\sigma(k)$ and $\tau(k)$. Finally, constraints (24) make sure that the vehicle capacity is not exceeded; (25) is a non-negativity constraint.

In most applications, stochastics is described by (large) data sets or by continuous distributions, as indicated in the model above. Stochastic programs, solved with exact methods, cannot handle that directly but need discrete distributions with limited cardinality. Hence, there is a need to pass from the underlying distribution (empirical, possibly) to a discrete one. This process has become known as *scenario generation*. The result is a set of scenarios (possible futures). In the multi-stage case, where randomness is revealed over time, these scenarios take the form of a tree, called the *scenario tree*. In our case, the model has only two stages, and hence the tree is just a "bush". However, we shall still talk about the scenario tree, even though it is a very shallow one.

The discretization process provides us with scenarios $s \in \mathcal{S}$, which can be used in the optimization process. A probability $p^s \geq 0$ is attached to each scenario, with $\sum p^s = 1$. As a result, the Y and the Z flow variables are indexed by s. A scenario is $|\mathcal{K}|$-dimensional, as it contains a demand $\delta^s(k)$ for each commodity k. The result of this process is to transform the stochastic formulation into a large-scale deterministic model, where the integral in the objective function of the stochastic formulation is replaced by a sum over all scenarios.

Thus, to reflect the move from a continuous to a discrete distribution, the objective function (14) is replaced with:

$$\min \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}} \sum_{t=0}^{T-1} c_{ij} X_{ij}^t + b \sum_{s \in \mathcal{S}} p^s Q(X, \delta^s) \tag{26}$$

The first stage then determines the "optimal" number of vehicles and associated schedules. The second stage describes the flow for each commodity k which can, and probably will, be different for the different scenarios s. Put in another way, the vehicle movements are the same for all scenarios, while the flow of freight is different for the various demand realizations. Obviously, both "stages" are solved simultaneously.

4 Experimental Setting

In Section 4.1 we look into how to deal with some of the challenges related to scenario generation. Section 4.2 provides an overview of the setup for our computational tests, with Section 4.2.1 providing an overview of the structure of the space-time network and Section 4.2.2 discussing the stochastic demand.

4.1 Building the Scenario Tree

Scenario generation is a very critical step in the setup of a stochastic program. On one hand, the scenario tree has to be small, otherwise the corresponding stochastic program cannot be solved to optimality with a reasonable computing effort. On the other hand, we must be sure that the scenario tree represents the underlying distribution reasonably well. We must make sure, in fact, that it is not the discretization procedure that drives the stochastic program, or, put differently, that the result of the stochastic program is not a random effect of a (random) scenario tree generation procedure. More to the point, sampling will eventually lead to a scenario tree which represents the underlying distributions arbitrarily well but may also results in a very large and (numerically) unsolvable stochastic program. On the other hand, sampling a small enough scenario tree for computational efficiency can yield a tree displaying uncontrollable proprieties and lead to very strange decisions. A general overview of scenario generation procedures can be found in Dupačová, Consigli, and Wallace (2001). It is worth noting that having good scenario trees is particularly important given our goal here: To study the structure of *optimal* solutions. We need to interpret our solutions relative to our formulations consisting of distributions and algebraic equations, and *not* relative to a scenario tree with potentially strange distributional properties caused by a faulty scenario generation procedure.

The scenario generation process we use is based on the method proposed by Høyland, Kaut, and Wallace (2003), which again is developed from Høyland and Wallace (2001). The idea is to construct a scenario tree with pre-specified properties. The authors show that, for their application (portfolio management), the necessary properties were four marginal moments (for each random variable) plus correlations. Which properties are important will depend on which model we are studying.

Most scenario generation procedures are themselves random. Hence, when run several times with the same data, they produce different scenario trees. By *in-sample stability* (see Kaut and Wallace 2007 for details) is understood that whichever of these trees is used in the optimization problem, the optimal objective function value is (approximately) the same. This is a necessary but not sufficient property of a satisfactory scenario generation procedure. By *out-of-sample stability* is understood that the true objective function values are also the same for those decisions obtained by the different scenario trees. Out-of-sample performance will normally be measured using some type of simulation.

When in- and out-of-sample stability are verified, we can run the scenario generation procedure only once and then solve the resulting program. The optimal objective function value will not depend on which scenario tree we ended up with. Similarly, the true expected objective value of this solution is the same as it would have been had we used any other of the possible trees. This emphasizes the importance of in- and out-of-sample stability. Note that we define stability relative to the objective function and not the solution. We do not consider it a problem that solutions are different if they have the same performance. This is not contrary to our goal of studying optimal solutions, as we look at structures of the solutions rather than the solutions themselves. And it turns out that different near-optimal solutions possess the same structural properties.

When using the method by Høyland, Kaut, and Wallace without any changes, we found in-sample instability. Analyzing the results showed that outcomes were generated outside the support of the random variables in some scenarios. This is not necessarily a problem (in the case of Kaut, Høyland, and Wallace it was not), but we suspected it might be of importance in a discrete model like ours, as even minute outcomes outside the support can potentially change the objective function value substantially due to the integrality requirements. Hence, we choose to enforce this property onto the tree.

We therefore changed the procedure slightly so that it generated outcomes within the support. This yielded the desired stability. In-sample stability was achieved using 50 scenarios (the difference between the highest and the lowest objective function values was 1.2%). Thus, in the present case, the first four marginal moments, correlations, and staying within the defined support were the necessary properties to achieve in-sample stability.

Out-of-sample stability was verified by *sampling* scenario trees $\mathcal{S}$ with 20000 scenarios from the same distributions used to construct scenario trees for the optimization. These were declared as the "truth". The evaluation was performed by fixing the integer variables (denoted $\hat{X}$) representing the vehicle schedules from the stochastic programs, and optimizing the flow of freight by solving the linear program (in fact, this is a collection of $|\mathcal{S}|$ separate linear programs):

$$\min \sum_{s \in \mathcal{S}} p^s Q(\hat{X}, \delta^s), \tag{27}$$

for the given $\hat{X}$ over the "true" tree $\mathcal{S}$. We then compared the values across the different schedules coming from our selection of scenario trees, and found that the difference between the highest and the lowest objective function values was 0.3% (in this comparison we add the vehicle costs $c\hat{X}$ to what we get from (27)). We find this satisfactory, and declare that we have out-of-sample stability.

4.2 Building the test problem instances

One of the challenges of creating good test instances is to avoid instances tailored to suit an algorithm or some desired results. To address the issue, we created 630 diverse schedules, all consisting of 12 commodities. Instances differ along a number of characteristics, most notably in vehicle operating costs, the association of random variables to commodities in the scenario tree, and correlations in demand for the various commodities. In the following subsections we describe how some of these characteristics were defined.

Notice that, for a better understanding of how demand uncertainty can affect the structure of service network design solutions, we have deliberately chosen to use very small test cases (10 - 16 commodities). There are two reasons for this. First, small dimensions allow us to solve to optimality instances of this $\mathcal{NP}$-hard and degenerate problem. Second, it is easier to visually study the results when dimensions are small, which contributes to developing insights.

4.2.1 Networks and costs

All our experiments use a time-space network that consists of six nodes repeated for seven periods. We always use complete networks with no hub-and-spoke structure. Costs associated with the use of vehicles (c_{ij}) are given by nine different matrices, making the instances diverse. The cost associated with ad-hoc capacity increase (b) is the same in all instances.

4.2.2 Demands

Very little work has been done on how dependence affects the solutions to stochastic optimization problems, with the notable exception of the analysis of correlations in portfolio management in finance (e.g., Chopra and Ziemba 1993). For service network design, Lium, Crainic and Wallace (2007) discuss the effects of demand correlations. We found that correlations did indeed matter. The behavior of the deterministic solution varied substantially depending on the underlying correlation structure (obviously overlooked by the deterministic model). The different stochastic cases (varying types of correlations) showed very different performance when tested with demand based on correlations different from what was originally used.

Correlations describe how the different demands vary relatively one to another, and the results show that a better understanding of these patterns can be crucial to generating robust schedules. We also saw that schedules corresponding to very high positive correlations were less robust when facing demands based on different correlation matrices, compared to schedules generated with zero or mixed correlated demands. The reason is that strong positive correlations push the problem instance close to the worst-case where robustness is of no concern in the construction of the schedule.

To account for these findings in our experiments, we used three different correlation settings: 1) all the commodities are positively correlated, with the correlation between each pair of commodities set at 0.4 or 0.7; 2) uncorrelated demand; 3) mix of (positive) demand correlations set at 0.5 and (negative) demand correlation set at -0.5.

To better demonstrate how uncertainty affects solutions, we use three different "levels" of uncertainty; high, low, and no uncertainty For instances with demand uncertainty, stochasticity is represented by the triangular distribution $\delta \sim$Tri(2,14,8) and $\delta \sim$Tri(5,11,8) for high and low uncertainty, respectively. In the deterministic case, the schedules are constructed using a mean of 8.

4.2.3 The problem instances

We constructed 90 deterministic problem instances using 10 demand matrices, each describing different O-D combinations for the commodities, and the 9 vehicle operating cost matrices.

We use the same approach to create instances for the stochastic problems, except that we introduce variability in demand. As indicated, we use 2 different "levels" of uncertainty, high and low, and 3 different types of correlation between commodities. This procedure yields 6 stochastic instances for each deterministic one, for a total of 540 stochastic problem instances.

5 Computational Results

To evaluate the effect of using stochastic models, we turn to a procedure similar to Monte Carlo simulation. We first solve each of the 90 deterministic and 540 stochastic problem instances. Since we know that we have in- and out-of- sample stability, we solve each case only once. Then, we generate three 1000-scenario trees that we declare to represent the "true" situations (actual realizations of demands). These trees are based on the same data as the problem instances: triangular distributions with $\delta \sim$Tri(2,14,8) and three different types of correlations between the random variables (uncorrelated, positive correlation of 0.7, and a mix of positively, 0.5, and negatively, -0.5, correlated demands).

For each service schedule, $\hat{X}$, we then solve the flow distribution problem (the linear program given by (27)) using the large 1000-scenario tree with the same correlation structure as was used to find $\hat{X}$.

We thus obtain the best utilization of the schedules selected by the stochastic service network design model (using a small tree), given the more appropriate distribution of demands contained in the "true" large tree. An advantage of this approach is that we avoid potential biases in the evaluation caused by a different evaluation procedure (e.g., a discrete-time simulator). So, in a sense, the evaluation is fair.

Statistical measures are then computed and are displayed in Table 1. A number in the "Stoch." ("Det.") line of Table 1 represents the average expected cost of using the 180 stochastic (the 90 deterministic) schedules, for a given correlation matrix, measured using the corresponding 1000-scenario tree. Half of the positively correlated schedules were based on a correlation of 0.4 and the other half on a correlation of 0.7. Deterministic schedules are based on average demand. The fifth column shows the average performance for each type of schedule.

Table 1: Average expected true cost of stochastic and deterministic problem instances

	Correlation type			
Schedules	Uncorrelated	Positive	Mix	Average Performance
Stoch.	4939.94	5110.8	4948.53	4999.76
Det.	5875.1	6235.6	5946.9	6019.18
Savings	15.92%	18.0%	16.8%	16.94%

The last row of Table 1 gives an indication of potential savings resulting from the integration of demand stochasticity into the service network design models when constructing schedules. Despite the limitations of our experiments, indications of average savings of 17% are interesting and shows that using deterministic formulations to solve stochastic problems may cause unnecessary costs.

6 The Impact of Stochasticity on Plans

The purpose of this section is to illustrate some of the structures we found by studying the solutions from Section 5. We shall use very simple examples in order to bring forward the major ideas. We shall see that, on one hand, the structures are somewhat simple while, on the other hand, deterministic models would not produce these structures. We shall also observe that the well-known structure of hub-and-spoke, in particular the idea of consolidation, will be part of the solution structures without being enforced. It is simply optimal, in light of randomness in demand, to use consolidation. Under certain reasonable conditions, consolidation takes place in a hub-and-spoke environment.

6.1 Reducing risk and capacity requirements through consolidation

Consolidation in LTL-trucking is normally thought of as a way to accommodate the fact that most loads are less than one truckload. In deterministic models, consolidation helps keep truck utilization up, and also facilitates the use of reasonably large trucks over long distances. However, there are also risk-related reasons for using consolidation. Let us look at a minor example where each load has an expected size equal to the truck capacity. In such a setting, consolidation will not be part of a deterministic solution.

Consider the example with two commodities in Figures 2 and 3. One commodity (represented by a dashed arrow) becomes available at node 1 at time 0 and has to be delivered to node 2 within three periods. The other commodity (represented by a dotted arrow) becomes available at node 3 at time 1 and has to be delivered to node 2 within one period. The solid arrows in the two figures show the optimal schedule for the two trucks that are used to transport the commodities from their origins to their destinations. We note that there will not be any consolidation in the deterministic solution (assuming just these two commodities) so that each commodity will use its own truck.

In the stochastic solution, we see both trucks following the same path. This is a different and, in the example, slightly more expensive schedule, compared to the deterministic one. However, it allows the two commodities to share the joint capacity of the two trucks from node 3 to node 2 at time 1. This is an example where commodity 1 is sent trough an intermediary breakbulk (node 3), before being consolidated with commodity 2 using the two trucks servicing the link from node 3 to node 2 at time 1. Having this kind of solution where two commodities share the common capacities of two trucks, makes expensive ad-hoc capacity increase less likely compared to if they had used one truck each as in the deterministic case.

Let us illustrate by using some numbers. Let trucks have capacity 2, and let the demand for each of the commodities be 0, 2 or 4, each with a probability of $\frac{1}{3}$. Hence, in total, there are nine possible scenarios, and if the demands are independent, each scenario has a probability of $\frac{1}{9}$.

If we use the deterministic solution shown in Figure 2, but let it face uncertainty, whatever demand that cannot be met will be handled by the ad-hoc capacity increase. Of the nine scenarios, five will result in ad-hoc capacity increase, and the expected amount of ad-hoc capacity increase is

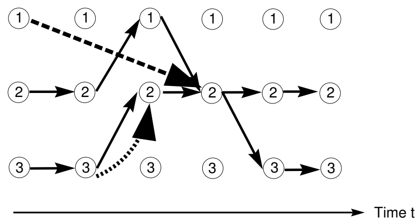

Figure 2: Deterministic solution, no consolidation, two trucks

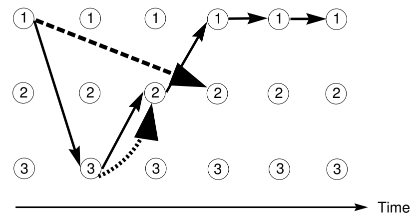

Figure 3: Stochastic solution, consolidation between nodes 3 and 2, two trucks

$\frac{4}{3}$. If we test the schedule in Figure 3, we note that the expected amount of ad-hoc capacity increase will be $\frac{8}{9}$, or $\frac{4}{9}$ lower than for the deterministic schedule. What has happened is that the two scenarios $(4, 0)$ and $(0, 4)$ can now be handled without ad-hoc capacity increase due to consolidation, that is, due to the ability to share transportation capacity.

What we observe here is closely related to what has been observed in other parts of the operations research literature. For example, in inventory theory, when the number of warehouses/inventories drops, safety stocks can be reduced without reducing the service levels. In finance, the risk in a portfolio of various financial instruments can be reduced by diversification, keeping the expected return the same.

To our knowledge, such use of consolidation as a means to hedge against uncertainty is a feature that has not been described in the service network design literature. There can be several reason for this, but since the literature almost exclusively refers to or uses deterministic models, hedging against uncertainty becomes irrelevant. Notice again, that consolidation is not a property we enforce on the solution, but a structure that emerges because it is good for the overall behavior of the schedules.

6.2 Flexibility through path sharing

In the above example there are only two commodities. Let us now pass to an example with four commodities. We shall see that not only is it good to consolidate, but also to have many paths available for each commodity. Also here, all observed aspects of consolidation will come from stochastics, and not from standard volume related arguments.

Figure 4 illustrates the example. We have four commodities becoming available at their respective origins at time zero to be transported to their destinations within three periods. The schedule in Figure 6 is based on stochastic demand, while for the schedule of Figure 5, demand equals the expected demand from the stochastic case. Otherwise, all parameters are the same.

The deterministic and stochastic schedules require the same number of trucks. In the deterministic case no ad-hoc capacity increase is used (when the demand is deterministic), while the stochastic case uses some ad-hoc capacity increase (less than 0.1% of the expected flow) to be able to transport all the commodities to their destinations. (However, this does not imply that the deterministic solution is better with respect to ad-hoc capacity increase. When the deterministic solution is subjected to random demand, a total of 1.17 % of the expected flow is moved with the help of the ad-hoc capacity.)

The most substantial difference between the two solutions is the number of paths connecting the various O-D pairs. In the deterministic case there are two O-D pairs that are connected with one path while the two others are connected by two paths. In the stochastic case each O-D pair is connected by at least two paths. The higher number of paths in the stochastic solution makes it easier to switch the flow of commodities from one path to another, if required because the first path is taken by another O-D pair. Such situations might occur when there is a surge in demand on a specific O-D pair that fully or partially can be routed on a different path. This larger number of paths in the stochastic solution gives more operational flexibility when routing commodities through the network. This is shown in Figure 8 where one of the commodities has three available paths, while the schedule based on a deterministic approach only offers one possible path, as shown in Figure 7. In most cases, this flexibility increases the capacity available to each commodity, without having to increase the total capacity in the network by adding more and/or larger trucks. This makes our stochastic solution perform better under uncertainty.

This result is in line with what we have observed in a variety of test cases. When uncertainty becomes an issue, our solutions habitually move away from “direct connections” between origins and destinations in deterministic cases, to more hub-and-spoke looking networks where the freight is being shipped trough intermediary terminals. This is not because it was “decided” a priori that freight between some O-D pairs should be handled this way, but because it turned out to be the best solution to deal with the uncertainty. Using our model has shown us that consolidation in hub-and-spoke networks takes

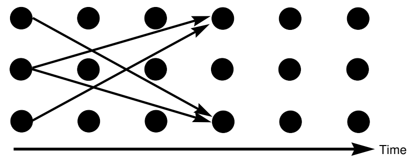

Figure 4: Origin and destination of commodities

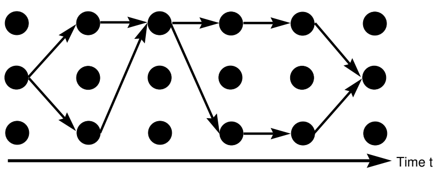

Figure 5: Schedule of trucks based on deterministic demand

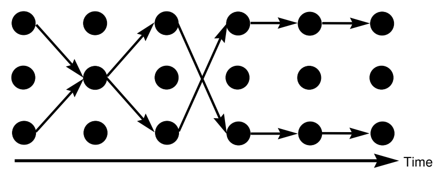

Figure 6: Schedule of trucks based on stochastic demand

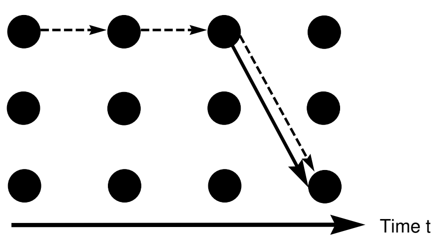

Figure 7: The flow of a commodity in a schedule based on a deterministic approach (represented by the dotted line)

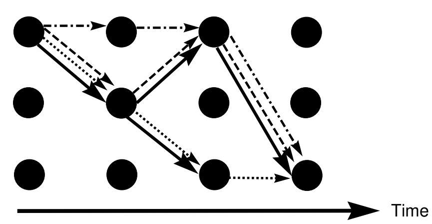

Figure 8: The flow of a commodity in a schedule based on a stochastic approach (represented by the non-solid lines)

place not necessarily due to economy of scale or other similar volume related reasons, but as a result of the need to hedge against uncertainty.

So, the conclusion so far is that a stochastic solution will try to provide many paths for each O-D pair, such that capacity is shared with other (different) O-D pairs on each of these paths. This provides optimal operational flexibility. It implies that as soon as some demand is low, others can immediately utilize the freed capacity if they are in need. A deterministic model would never produce such structures.

6.3 How correlations affect schedules

Look back at Figures 2 and 3. When we calculated the value of using the stochastic rather than the deterministic model, we assumed the demands were independent. Assume instead that the two demands are perfectly negatively correlated. In that case, the sum of the two demands will always be 4, and the expected ad-hoc capacity increase will be zero. But, if we use the solution from the deterministic model, the expected ad-hoc capacity increase will remain at $\frac{4}{3}$. What we observe is not surprising, but important. Negative correlations imply a chance to achieve hedging, but only if the two negatively correlated flows can be set up so as to share capacity. And the more negatively correlated they are, the more important the issue.

On the other hand, in the same example, if the two demands are perfectly positively correlated, the expected ad-hoc capacity increase will be the same for the deterministic and stochastic solution, since, in fact, consolidation will never take place.

More generally, as long as random variables are not perfectly positively correlated, there is something to be gained from flexible routing. The more we move towards perfectly negative correlations, the higher is the potential for hedging. As flexibility normally comes at a cost, there is a tradeoff between the cost of achieving operational flexibility and the expected gain from the investment. Of course, when there are many random variables, the relationships are more complex. But even so, we can conclude as follows: Operational flexibility is achieved by having many paths available for each O-D pair, such that each of these paths is shared by other O-D pairs. The more negative the correlations are, the more there potentially is to be gained from well structured schedules.

7 Conclusions and Perspectives

In the presented cases we show that uncertainty plays a very important role when constructing schedules in service network design, and that there exist certain characteristics in solutions based on uncertain demand that can be of practical value as they provide operational flexibility, making the schedules robust as seen by the customers. Solutions containing these characteristics will typically not be found when a deterministic approach is used, as flexibility does not even have meaning in a deterministic world. The future is "known" and no flexibility would therefore be needed. Our experiments show that solutions based on a stochastic approach can be structurally different from their deterministic counterparts. Such structural differences might vary from case to case, but there are two characteristics that seem to show up in most of the cases when dealing with uncertainty. One is the number of paths between the origin and destination, the other is that the more commodities share a link, the higher expected utilization the link achieves.

Correlation in demand between different commodities is clearly an issue (or should be) when planning for consolidation carriers. Our findings indicate that there are benefits to be obtained when commodities share paths/services. This is particularly important when negative correlations are present. The reason is that commodities that share some capacity increase the utilization of this resource. This result is very similar to how one can reduce risk in a financial portfolio by mixing negatively correlated financial instruments.

By neither enforcing nor encouraging any specific structure or design upon the schedules, we have seen that consolidation and hub-and-spoke systems offer better solutions when there are uncertainties in demand. Most OR literature advocate the use of consolidation as a mean to increase efficiency/higher frequencies/economy of scale, while we have demonstrated that uncertainty also favors the use of a hub-and-spoke systems as they provide hedging. Our findings indicate that making strategic decisions, such as how many hubs to use and where to locate them, and more tactical/operational activities such as creating schedules, should not be two separate decision processes, but – preferably – one.

There are especially two issues that look promising for future research. One is to use the knowledge we have gained so far in some heuristic method, so as to obtain good solutions without having to solve a stochastic program. Another issue is to study similar problems, such as more general network design problems.

Acknowledgments

We would like to thank Michal Kaut for modifying his scenario generation procedure to suite our needs. Stein W. Wallace was supported by grant 156315/530 from The Research Council of Norway. Partial funding for this research has been provided by the Natural sciences and Engineering Research Council of Canada (NSERC) through its Discovery Grants and Chairs and Faculty Support programs. While working on this project, Dr. Crainic was Adjunct Professor at the Molde University College, Molde, Norway, and member of CIRRELT, the Interuniversity research Centre on Network Enterprises, Logistics, and Transportation, Canada.

References

A.P. Armacost, C. Barnhart, and K.A. Ware. Composite Variable Formulations for Express Shipment Service Network Design. *Transportation Science*, 36(1):1–20, 2002.

C. Barnhart and R.R. Schneur. Network Design for Express Freight Service. *Operations Research*, 44(6):852–863, 1996.

K. Buedenbender, T. Grünert, and H.-J. Sebastian. A Hybrid Tabu Search/Branch and Bound Algorithm for the Direct Flight Network Design Problem. *Transportation Science*, 34(4):364–380, 2000.

Chopra, V.K. and Ziemba, W.T. The Effect of Errors in Mean, Variance and Co-Variance Estimates on Optimal Portfolio Choice. *Journal of Portfolio Management*, 19(2):6–11, 1993.

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

Christiansen, M., Fagerholt, K., and Ronen, D. Ship Routing and Scheduling: Status and Perspectives. *Transportation Science*, 38(1):1–18, 2004.

J.-F. Cordeau, P. Toth, and D. Vigo. A Survey of Optimization Models for Train Routing and Scheduling. *Transportation Science*, 32(4):380–404, 1998.

Crainic, T.G. Network Design in Freight Transportation. *European Journal of Operational Research*, 122(2):272–288, 2000.

Crainic, T.G. Long-Haul Freight Transportation. In Hall, R.W., editor, *Handbook of Transportation Science*, pages 451–516. Kluwer Academic Publishers, Norwell, MA, second edition, 2003.

Crainic, T.G., Ferland, J.-A., and Rousseau, J.-M. A Tactical Planning Model for Rail Freight Transportation. *Transportation Science*, 18(2):165–184, 1984.

Crainic, T.G. and Kim, K.H. Intermodal Transportation. In Barnhart, C. and Laporte, G., editors, *Transportation*, volume 14 of *Handbooks in Operations Research and Management Science*, pages 467–538. North-Holland, Amsterdam, 2007.

Crainic, T.G. and Rousseau, J.-M. Multicommodity, Multimode Freight Transportation: A General Modeling and Algorithmic Framework for the Service Network Design Problem. *Transportation Research Part B: Methodological*, 20:225–242, 1986.

Crainic, T.G. and Roy, J. O.R. Tools for Tactical Freight Transportation Planning. *European Journal of Operational Research*, 33(3):290–297, 1988.

Dupačová, J., Consigli, G., and Wallace, S.W. Scenarios for Multistage Stochastic Programs. *Annals of Operations Research*, 100:25–53, 2001. ISSN 0254-5330.

J.M. Farvolden and W.B. Powell. 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, 1991.

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

M.F. Gorman. An Application of Genetic and Tabu Searches to the Freight Railroad Operating Plan Problem. *Annals of Operations Research*, 78: 51–69, 1998a.

M.F. Gorman. Santa Fe Railway Uses an Operating-Plan Model to Improve Its Service Design. *Interfaces*, 28(4):1–12, 1998b.

T. Grünert and H.-J. Sebastian. Planning Models for Long-haul Operations of Postal and Express Shipment Companies. *European Journal of Operational Research*, 122:289–309, 2000.

T. Grünert, H.-J. Sebastian, and M. Thäerigen. The Design of a Letter-Mail Transportation Network by Intelligent Techniques. In R. Sprague, editor, *Proceedings Hawaii International Conference on System Sciences 32*, 1999.

A.E. Haghani. Formulation and Solution of Combined Train Routing and Makeup, and Empty Car Distribution Model. *Transportation Research B: Methodological*, 23B(6):433–452, 1989.

Higle, J.L. and Wallace, S.W. Sensitivity Analysis and Uncertainty in Linear Programming. *Interfaces*, 33:53–60, 2003.

Høyland, K., M. Kaut, and Wallace, S.W. A Heuristic for Moment-matching Scenario Generation. *Computational Optimization and Applications*, 24 (2-3):169–185, 2003. ISSN 0926-6003.

Høyland, K. and Wallace, S.W. Generating Scenario Trees for Multistage Decision Problems. *Management Science*, 47(2):295–307, 2001.

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

Michal Kaut and Stein W. Wallace. Evaluation of scenario-generation methods for stochastic programming. *Pacific Journal of Optimization*, to appear, 2007.

M.H. Keaton. Designing Optimal Railroad Operating Plans: Lagrangian Relaxation and Heuristic Approaches. *Transportation Research B: Methodological*, 23B(6):415–431, 1989.

M.H. Keaton. Service–Cost Tradeoffs for Carload Freight Traffic in the U.S. Rail Industry. *Transportation Research A: Policy and Practice*, 25A(6): 363–374, 1991.

M.H. Keaton. Designing Optimal Railroad Operating Plans: A Dual Adjustment Method for Implementing Lagrangian Relaxation. *Transportation Science*, 26:263–279, 1992.

B.W. Lamar, Y. Sheffi, and W.B. Powell. A Capacity Improvement Lower Bound for Fixed Charge Network Design Problems. *Operations Research*, 38(4):704–710, 1990.

Lium, A.-G., Crainic, T.G., and Wallace, S.W. Correlations in Stochastic Programming: A Case from Stochastic Service Network Design. *Asian Pacific Journal of Operational Research*, 2007. forthcoming.

H.N. Newton. *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., 1996.

H.N. Newton, C. Barnhart, and P.H. Vance. Constructing Railroad Blocking Plans to Minimize Handling Costs. *Transportation Science*, 32(4):330–345, 1998.

W.B. Powell. A Local Improvement Heuristic for the Design of Less-than-Truckload Motor Carrier Networks. *Transportation Science*, 20(4):246–357, 1986.

W.B. Powell and Y. Sheffi. The Load-Planning Problem of Motor Carriers: Problem Description and a Proposed Solution Approach. *Transportation Research A: Policy and Practice*, 17(6):471–480, 1983.

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

W.B. Powell and Y. Sheffi. Design and Implementation of an Interactive Optimization System for the Network Design in the Motor Carrier Industry. *Operations Research*, 37(1):12–29, 1989.

Roy, J. and Delorme, L. NETPLAN: A Network Optimization Model for Tactical Planning in the Less-than-Truckload Motor-Carrier Industry. *INFOR*, 27(1):22–35, 1989.

A. Ruszczyński and A. Sharpiro, editors. *Stochastic Programming*, volume 10 of *Handbooks in Operations Research and Management Science*. North-Holland, Amsterdam, 2003.

Smilowitz, K.R., Atamtürk, A, and Daganzo, C.F. Deferred Item and Vehicle Routing within Integrated Networks. *Transportation Research Part E: Logistics and Transportation*, 39:305–323, 2003.

S.W. Wallace. Decision Making Under Uncertainty: Is Sensitivity Analysis of Any Use? *Operations Research*, 48:20–25, 2000.