Stochastic service network design: A deterministic primal heuristic

Stein W. Wallace
P.O. Box 2110 NO-6402 Molde, Norway
stein.w.wallace@himolde.no

Teodor Gabriel Crainic
Centre de recherche sur les transports, Université de Montréal
C.P. 6128, succ. Centre-ville, Montréal QC, Canada H3C 3J7
theo@crt.umontreal.ca

Michal Kaut, Arnt-Gunnar Lium
Michal.Kaut/Arnt.G.Lium@himolde.no

08 September-06

Abstract

We have earlier demonstrated that introducing stochastic demand into a service network design model produces solutions qualitatively different from those stemming from deterministic models. These structural differences are now used to create a primal deterministic heuristic for finding solutions to the stochastic program.

For the study reported in this paper, we use a version of a multi-period service network design model inspired by less-than-truckload motor carrier cases with repetitive schedules. Based on the knowledge of the demand distributions, we propose a two-phase approach, consisting of a construction followed by an improvement heuristic to build a feasible solution to the deterministic formulation. The construction and improvement operations follow from the major patterns observed in the previous study. In the talk, we will discuss the problem and the approach, present the primal heuristic, and analyse the available results.

Keywords: *Stochastic Programming, Service Network Design, Scheduling, Heuristic*

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 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 and modal applications. See Crainic, T.G. [2003, 2000], Crainic, T.G. and Kim, K.H. [2005] for recent reviews.

The service network literature and, at 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, in several papers, it is clearly stated that one works with forecasted demand, 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.

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?

In an article in Lium [2006] we have demonstrated that the service design patterns are indeed different when stochastics is taken explicitly into account. We have described some patterns that show up repeatedly in the stochastic cases, but which will never be produced by deterministic models. The patterns are simple. The purpose of this paper is to develop primal heuristics using these patterns to find good solutions to the stochastic service network design problem. A novelty of this approach is indeed this: The heuristic does not use any stochastic information: The approach is fully deterministic, and hence, at least in principle, applicable also to large (stochastic) cases.

The most important patterns that emerged from our analysis were

1. Hub-and-spoke structures emerged as hedging devises (and not only for volume-related reasons)
2. For each O-D pair (commodity) one should have many alternative paths from source to sink
3. A commodity should share paths with as many other commodities as possible.

The latter two properties imply that for a given commodity there should be many paths, each of which should be shared with one or more commodities. This way capacities can be shared: Whenever the demand for one commodity is low, the other commodities can use the available capacity.

In most service network design models, structures are enforced upon the solution before any optimization takes place at all. This is typically the case when requiring that the solution constitutes a hub-and-spoke network, or that free resources should be used to create buffers in the truck's schedules. In this sense this work fits into a well-known body of literature. We force the newly found patterns onto our solution. Our patterns cover both the standard volume-related arguments for consolidation as well as hedging against stochastic demand. It is worth noting, though, that our patterns come from solving cases without structures being enforced. They emerged from small stochastic cases being solved to optimality.

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. In particular, we will have the ex-post optimal solution in our collection. Next, 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, we refer to for example Kall and Wallace [1994] and Volume 10 of *Handbooks in Operations Research and Management Science*, edited by Ruszczyński and Shapiro.

Consolidation transportation carriers are 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 together 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 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.

Our model takes the form of a stochastic, fixed cost, capacitated, multicommodity network design formulation. Integer-valued decision variables are used to represent service selection decisions, while product-specific and scenario dependent variables capture the commodity flows. Fixed-costs are associated to the inclusion of services into the plan. Costs that vary with the intensity of service and commodity traffic are associated to 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.

References

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

Crainic, T.G. and Kim, K.H. Intermodal Transportation. In Barnhart, C. and Laporte, G., editor, *Transportation*, Handbooks in Operations Research and Management Science, page to appear. Elsevier, 2005.

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

P. Kall and S.W. Wallace. *Stochastic Programming.* Wiley, Chichester, 1994.

A.-G. Lium. *Stochastic Service Network Design.* PhD thesis, Molde University College, 2006.

A. Ruszczyński and A. Shapiro, editors. *Stochastic Programming*, volume 10 of *Handbooks in Operations Research and Management Science*. Elsevier, Amsterdam, 2003.