Transportation Resource Management

Sumit Kunnumkal
Indian School of Business, Gachibowli, Hyderabad, 500032, India
sumit_kunnumkal@isb.edu

Huseyin Topaloglu
School of Operations Research and Information Engineering,
Cornell University, Ithaca, New York 14853, USA
topaloglu@orie.cornell.edu

February 23, 2010

Abstract

In this chapter, we describe a variety of modeling and solution approaches for transportation resource management problems. The solution approaches that we propose either build on deterministic linear programming formulations, or formulate the problem as a dynamic program and use tractable approximations to the value functions. We describe two classes of methods to construct approximations to the value functions. The first class of methods relax certain constraints in the dynamic programming formulation of the problem by associating Lagrange multipliers with them, in which case, we can solve the relaxed dynamic program by concentrating on one location at a time. The second class of methods use a stochastic approximation idea along with sampled trajectories of the system to iteratively update and improve the value function approximations. After describing several solution approaches, we provide a flexible paradigm that is useful for modeling and communicating complex resource management problems. Our numerical experiments demonstrate the benefits from using models that explicitly address the randomness in resource management problems.

Many problems in transportation logistics involve managing a set of resources to satisfy the service requests that arrive randomly over time. For example, truckload carriers decide which load requests they should accept and to which locations they should reposition the empty vehicles. Railroad companies deploy the empty railcars to different stations in the network to serve the random customer demand in the most effective manner. Irrespective of the application setting, transportation resource management problems pose significant challenges. To begin with, these problems tend to be large in scale, spanning transportation networks with hundreds of locations and planning horizons with hundreds of time periods. Furthermore, the service requests arrive randomly, and we may not have full information about the future service requests. When making the decisions for the current time period, we need to keep a balance between maximizing the immediate profits and getting the vehicles into favorable positions to serve the potential future service requests. Finally, the decisions that we make for different locations and for different time periods interact in a nontrivial fashion. For example, the decision to serve a particular load from location i to j at time period t affects what other loads we can serve at the future time periods, since the vehicle that is used to serve the load at time period t becomes available at location j at a future time period.

In this chapter, we describe a variety of modeling and algorithmic approaches for transportation resource management problems under uncertainty. We begin with a problem that takes place over two time stages. At the first stage, we decide to which locations we should reposition the resources to serve the demands that arrive later. At the second stage, we observe the demand arrivals and decide which of these demands we should serve. The first algorithmic strategy that we discuss for the two-stage problem is based on a deterministic linear programming formulation. This formulation is simple to implement, but it does not explicitly address the randomness in the demand arrivals. Motivated by this observation, we move on to other algorithmic strategies that address the randomness in the demand arrivals by using a dynamic programming formulation. The dynamic programming formulation of the two-stage problem turns out to be computationally difficult as it involves a high dimensional state variable, and we resort to approximation strategies. In particular, the first two approximation strategies that we describe decompose the dynamic programming formulation by the locations and obtain good solutions by focusing on one location at a time. We also discuss a third approximation strategy that constructs separable value function approximations by using sampled trajectories of the system. After the discussion of two-stage problems, we move on to resource management problems that take place over multiple time periods. We give a dynamic programming formulation for a multi-stage resource management problem and demonstrate that the algorithmic concepts that we develop for two-stage problems can be extended to multi-stage problems without too much difficulty.

In addition to providing algorithmic tools for two-stage and multi-stage problems, we also dwell on modeling issues. Most transportation resource management problems involve complex resources. For example, a driver resource in a driver scheduling application may require keeping track of its inbound location, time to reach the inbound location, duty time within the shift, days away from home, vehicle type and home domicile. It would be cumbersome to model such complex entities through traditional mathematical programming tools, and we provide a flexible paradigm that is useful for modeling and communicating complex resource management problems. We conclude the chapter with a series of numerical experiments that demonstrate the benefits from using models that explicitly capture the randomness in the system.

Transportation resource management models have a long history dating back to the earliest days of linear programming; see Dantzig and Fulkerson (1954), Ferguson and Dantzig (1955), White and Bomberault (1969) and White (1972). The early models formulate the problem over a state time network. The nodes of the state time network represent the supply of resources at different locations and at different time periods, whereas the arcs represent the movements of the resources between different locations. These models tend to ignore the randomness in the problem and simply work with the expected values of the demand arrivals. A subclass of the deterministic models is the so called myopic models, which ignore the unknown future altogether and only work with what is known with certainty; see Powell (1988), Powell (1996) and Erera, Karacik and Savelsbergh (2008). This approach is justified by the fact that many customers call in advance, and a large portion of the future demands is actually known with certainty. Whether they ignore the uncertainty in the future altogether or incorporate it through expectations of the demand arrivals, these kinds of deterministic models can yield tractable optimization problems and can be relatively simple to implement. As a result, they form the backbone of many commercial systems; see Hane, Barnhart, Johnson, Marsten, Nemhauser and Sigismondi (1995), Holmberg, Joborn and Lundgren (1998) and Gorman (2001).

The majority of the models that we discuss in this chapter fall under the category of stochastic models, which try to capture the random demand arrivals by formulating the problem as a dynamic program and using value functions to assess the impact of the current decisions on the future evolution of the system. The dynamic programming formulations of practical resource management problems generally involve high dimensional state variables, and this high dimensionality makes it impossible to compute the value functions by using standard Markov decision process tools. Therefore, most of the effort is concentrated around approximating the value functions in a tractable fashion; see Frantzeskakis and Powell (1990), Crainic, Gendreau and Dejax (1993), Carvalho and Powell (2000), Godfrey and Powell (2002*a*), Godfrey and Powell (2002*b*), Kleywegt, Nori and Savelsbergh (2002), Adelman (2004), Topaloglu and Powell (2006), Adelman (2007) and Schenk and Klabjan (2008). The overview of transportation resource management models that we give in this chapter is rather compact, and it mostly evolves around methods that approximate the value functions. We refer the reader to Powell, Jaillet and Odoni (1995), Powell and Topaloglu (2005) and Barnhart and Laporte (2006) for detailed reviews of deterministic and stochastic models. We also emphasize that our goal in this chapter is to give an overview of transportation resource management models in the literature, rather than to make original contributions. At the end of each section in this chapter, we give pointers to the relevant literature and acknowledge the original contributors.

The chapter is organized as follows. In Section 1, we formulate a two-stage resource management problem as a dynamic program. This formulation has a high dimensional state variable and we discuss a variety of algorithmic approaches to obtain tractable approximations. In Section 2, we extend the ideas that we discuss in Section 1 to multi-stage problems. In Section 3, we give a flexible paradigm for modeling complex resource management problems. In Section 4, we present numerical experiments.

1 Two-Stage Problems

In this section, we consider a resource management problem that takes place over a planning horizon of two time periods. Our goal is to illustrate the fundamental algorithmic concepts by using a relatively simple problem setting. It turns out that all of the development in this section can be extended to multi-stage problems without too much difficulty.

We have a set of resources that can be used over a planning horizon of two time periods to serve the demands that arrive randomly into the system. At the beginning of the first time period, there is a certain number of resources at different locations. The decisions that we make at the first time period involve repositioning the resources between different locations so that the resources end up at favorable locations to serve the demands that arrive at the second time period. At the beginning of the second time period, we observe the demand arrivals. The decisions that we make at the second time period involve serving the demands by using the resources. The objective is to maximize the total expected profit over the two time periods, which is the difference between the total cost of repositioning the resources and the total expected revenue from serving the demands.

The set of locations in the transportation network is $\mathcal{L}$. At the beginning of the first time period, we have s_i resources available at location i, and $\{s_i : i \in \mathcal{L}\}$ is a part of the problem data. We use x_{ij} to denote the number of resources that we reposition from location i to j at the first time period. The cost of repositioning a resource from location i to j is c_{ij}. We let r_i be the number of resources available at location i at the beginning of the second time period. Naturally, $\{r_i : i \in \mathcal{L}\}$ is determined by the repositioning decisions that we make at the first time period. The set of possible demand types is $\mathcal{K}$. We use y_i^k to denote the number of demands of type k that we serve by using a resource at location i at the second time period. The revenue from serving a demand of type k by using a resource at location i is p_i^k. If it is not feasible to serve a demand of type k by using a resource at location i, then we capture this infeasibility by letting $p_i^k = -\infty$. We use the random variable D^k to represent the number of demands of type k that arrive at the second time period.

Letting $r = \{r_i : i \in \mathcal{L}\}$ and $D = \{D^k : k \in \mathcal{K}\}$ to respectively capture the number of resources at different locations and the demand arrivals at the second time period, we can find the total revenue at the second time period by computing the value function $V(r, D)$ through the problem

$$\begin{aligned}
V(r, D) = \max \quad & \sum_{i \in \mathcal{L}} \sum_{k \in \mathcal{K}} p_i^k \, y_i^k && && \text{(1a)} \\
\text{subject to} \quad & \sum_{k \in \mathcal{K}} y_i^k \le r_i && \forall\, i \in \mathcal{L} && \text{(1b)} \\
& \sum_{i \in \mathcal{L}} y_i^k \le D^k && \forall\, k \in \mathcal{K} && \text{(1c)} \\
& y_i^k \in \mathbb{Z}_+ && \forall\, i \in \mathcal{L},\ k \in \mathcal{K}. && \text{(1d)}
\end{aligned}$$

In the problem above, constraints (1b) ensure that the number of resources that we use to serve the demands does not violate the resource availabilities at different locations, whereas constraints (1c) ensure that the number of demands that we serve does not exceed the demand arrivals. If the number of resources at different locations is given by the vector r, then the total expected revenue that we generate at the second time period is captured by $\bar{V}(r) = \mathbb{E}\{V(r, D)\}$. Therefore, we can maximize the total expected profit over the two time periods by solving the problem

$$Z = \max \quad -\sum_{i\in\mathcal{L}} \sum_{j\in\mathcal{L}} c_{ij}\, x_{ij} + \bar{V}(r) \tag{2a}$$

$$\text{subject to} \quad \sum_{j\in\mathcal{L}} x_{ij} = s_i \qquad \forall\, i \in \mathcal{L} \tag{2b}$$

$$\sum_{i\in\mathcal{L}} x_{ij} - r_j = 0 \qquad \forall\, j \in \mathcal{L} \tag{2c}$$

$$x_{ij}, r_j \in \mathbb{Z}_+ \qquad \forall\, i, j \in \mathcal{L}, \tag{2d}$$

where constraints (2b) ensure that the number of resources that we reposition does not violate the resource availabilities at different locations, and constraints (2c) compute the number of resources at different locations at the beginning of the second time period.

Conceptually, we can compute $V(r, D)$ for all possible values of the vector r and demand arrivals D by solving problem (1). In this case, we can take expectations to compute $\bar{V}(r) = \mathbb{E}\{V(r, D)\}$ for all possible values of the vector r and use this value function in problem (2) to find the optimal repositioning decisions. Practically, however, the number of possible values of the vector r and demand arrivals D can be so large that this approach becomes computationally intractable. Therefore, we primarily focus our attention to developing methods to approximate the value function $\bar{V}(\cdot)$.

1.1 Deterministic Linear Program

A common approach for problems that take place under uncertainty is to assume that all random variables take on their expected values and to solve a deterministic approximation. In our problem setting, if we assume that all demand arrivals take on their expected values and we can utilize fractional number of resources, then such a deterministic approximation can be written as

$$Z^{DLP} = \max \quad -\sum_{i\in\mathcal{L}} \sum_{j\in\mathcal{L}} c_{ij}\, x_{ij} + \sum_{i\in\mathcal{L}} \sum_{k\in\mathcal{K}} p_i^k\, y_i^k \tag{3a}$$

$$\text{subject to} \quad \sum_{j\in\mathcal{L}} x_{ij} = s_i \qquad \forall\, i \in \mathcal{L} \tag{3b}$$

$$-\sum_{i\in\mathcal{L}} x_{ij} + \sum_{k\in\mathcal{K}} y_j^k \le 0 \qquad \forall\, j \in \mathcal{L} \tag{3c}$$

$$\sum_{i\in\mathcal{L}} y_i^k \le \mathbb{E}\{D^k\} \qquad \forall\, k \in \mathcal{K} \tag{3d}$$

$$x_{ij}, y_j^k \in \mathbb{R}_+ \qquad \forall\, i, j \in \mathcal{L},\ k \in \mathcal{K}. \tag{3e}$$

Constraints (3b) in the problem above are analogous to constraints (2b) in problem (2), whereas constraints (3d) are analogous to constraints (1c) in problem (1). In constraints (3d), a portion of the demands $\{D^k : k \in \mathcal{K}\}$ may be known at the beginning of the first time period, in which case, these demands simply become deterministic quantities. Constraints (3c) ensure that the number of resources that we use to serve the demands at the second time period does not violate the number of resources that we reposition at the first time period. There are two uses of problem (3). First, this problem can be used to make resource repositioning decisions. In particular, letting $\{\hat{x}_{ij} : i, j \in \mathcal{L}\}$, $\{\hat{y}_i^k : i \in \mathcal{L},\ k \in \mathcal{K}\}$ be an optimal solution to problem (3), we can use $\{\hat{x}_{ij} : i, j \in \mathcal{L}\}$ as an approximation to the optimal solution to problem (2). The solution $\{\hat{x}_{ij} : i, j \in \mathcal{L}\}$ may utilize fractional number of resources. If this is a concern, then we can impose integrality requirements in problem (3).

The second use of problem (3) is that the optimal objective value of this problem provides an upper bound on the optimal total expected profit Z. To see the reason, we let $\{\hat{x}_{ij}(D) : i, j \in \mathcal{L}\}$, $\{\hat{y}_i^k(D) : i \in \mathcal{L},\ k \in \mathcal{K}\}$ be an optimal solution and $Z^{DLP}(D)$ be the optimal objective value of problem (3) when we replace the right side of constraints (3d) with the demand arrivals $D = \{D^k : k \in \mathcal{K}\}$. In this case, $Z^{DLP}(D)$ corresponds to the perfect hindsight total profit that is computed after the demand arrivals are observed to be D. Therefore, $Z^{DLP}(D)$ is an upper bound on the total profit that is obtained by any approach when the demand arrivals are observed to be D. Taking expectations, we observe that $\mathbb{E}\{Z^{DLP}(D)\}$ is an upper bound on the total expected profit that is obtained by any approach. Since the total expected profit obtained by the optimal policy is Z, we obtain

$$Z \le \mathbb{E}\{Z^{DLP}(D)\} \le Z^{DLP}(\mathbb{E}\{D\}) = Z^{DLP},$$

where the second inequality follows from Jensen's inequality and the equality follows from the definition of $Z^{DLP}(\mathbb{E}\{D\})$; see Bertsimas and Popescu (2003). This kind of an upper bound on the optimal total expected profit can be useful when assessing the optimality gap of suboptimal approximation strategies, such as the ones that we develop in this chapter. An interesting observation from the chain of inequalities above is that $\mathbb{E}\{Z^{DLP}(D)\}$ also provides an upper bound on Z and this upper bound is tighter than the one provided by Z^{DLP}. However, to compute $\mathbb{E}\{Z^{DLP}(D)\}$, we need to compute $Z^{DLP}(D)$ for all possible values of the demand arrivals D and take an expectation. Taking this expectation can be computationally intractable when the number of possible values of the demand arrivals is large, but we can resort to Monte Carlo samples to approximate the expectation $\mathbb{E}\{Z^{DLP}(D)\}$.

Deterministic approximations similar to the one in problem (3) form the core of many practical transportation models; see Gorman (2001). However, such approximations may not always yield satisfactory results as they completely ignore the uncertainty in the demand arrivals. Our goal is to present models that address the uncertainty in the demand arrivals more carefully.

1.2 Separable Value Functions

In this section, we describe a special case where $\bar{V}(r) = \mathbb{E}\{V(r, D)\}$ can be computed in a tractable fashion. It turns out that we can build on the results that we obtain for this special case to develop approximation strategies for more general cases.

We consider the special case where for each demand type k, there is a single location i_k such that only the resources at location i_k can serve a demand of type k. In other words, we have $p_i^k = -\infty$ for all $i \in \mathcal{L} \setminus \{i_k\}$, which implies that $y_i^k = 0$ for all $i \in \mathcal{L} \setminus \{i_k\}$ in an optimal solution to problem (1). In this case, it is easy to see that we can replace the constraints $\sum_{i \in \mathcal{L}} y_i^k \le D^k$ for all $k \in \mathcal{K}$ in problem

(1) with the constraints $y_i^k \leq D^k$ for all $i \in \mathcal{L}$, $k \in \mathcal{K}$. If we do replace the constraints $\sum_{i \in \mathcal{L}} y_i^k \leq D^k$ for all $k \in \mathcal{K}$ in problem (1) with the constraints $y_i^k \leq D^k$ for all $i \in \mathcal{L}$, $k \in \mathcal{K}$, then all of the constraints in problem (1) decompose by the locations, and we can obtain an optimal solution to problem (1) by concentrating on one location at a time. In particular, the subproblem for location i is of the form

$$
\begin{aligned}
V_i(r_i, D) = \max \quad & \sum_{k \in \mathcal{K}} p_i^k \, y_i^k && \text{(4a)} \\
\text{subject to} \quad & \sum_{k \in \mathcal{K}} y_i^k \leq r_i && \text{(4b)} \\
& y_i^k \leq D^k \quad \forall\, k \in \mathcal{K} && \text{(4c)} \\
& y_i^k \in \mathbb{Z}_+ \quad \forall\, k \in \mathcal{K}, && \text{(4d)}
\end{aligned}
$$

and we have $V(r, D) = \sum_{i \in \mathcal{L}} V_i(r_i, D)$. An important observation is that problem (4) is a knapsack problem. The capacity of the knapsack is r_i. There are $|\mathcal{K}|$ items, each item occupies one unit of space, and we can put at most D^k units of item k into the knapsack. Therefore, it is possible to solve problem (4) by sorting the objective function coefficients of the items and putting the items into the knapsack starting from the item with the largest objective function coefficient.

We need the value function $\bar{V}(r) = \sum_{i \in \mathcal{L}} \mathbb{E}\{V_i(r_i, D)\}$ to compute the optimal repositioning decisions through problem (2). It turns out that taking this expectation is tractable due to the knapsack structure of problem (4). We assume without loss of generality that the set of demand types is $\mathcal{K} = \{1, 2, \ldots, K\}$, and the demand types are ordered such that $p_i^1 \geq p_i^2 \geq \ldots \geq p_i^K$. Noting the knapsack interpretation of problem (4), we let $\psi_i^k(q)$ be the probability that the kth item uses the qth unit of available space in the knapsack for location i. In this case, the expected revenue that we generate from the qth unit of available space becomes $\sum_{k \in \mathcal{K}} \psi_i^k(q) \, p_i^k$, and the expectation of the optimal objective value of problem (4) can be computed as

$$
\mathbb{E}\{V_i(r_i, D)\} = \sum_{q=1}^{r_i} \sum_{k \in \mathcal{K}} \psi_i^k(q) \, p_i^k.
$$

To compute $\psi_i^k(q)$, we observe that for the kth item to use the qth unit of available space in the knapsack, the total amount of space used by the first $k-1$ items should be strictly smaller than q, the total amount of space used by the first k items should be greater than or equal to q, and item k should satisfy $i_k = i$, since otherwise we have $p_i^k = -\infty$ and it is never desirable to put the kth item into the knapsack. Therefore, we have $\psi_i^k(q) = \mathbb{P}\{D^1 + \ldots + D^{k-1} < q \leq D^1 + \ldots + D^k\}$ if $i_k = i$, and $\psi_i^k(q) = 0$ otherwise. This implies that we can compute $\psi_i^k(q)$ in a tractable fashion as long as we can compute the convolutions of the distributions of $\{D^k : k \in \mathcal{K}\}$, which is the case when the demand random variables have, for example, independent Poisson distributions.

The approach outlined above is proposed by Cheung (1994). Powell and Cheung (1994*a*), Powell and Cheung (1994*b*), Cheung and Powell (1996*a*) and Cheung and Powell (1996*b*) apply this approach to a variety of settings, including dynamic fleet management and inventory distribution. An interesting aspect of this approach is that it not only allows us to compute $\bar{V}(r) = \mathbb{E}\{V(r, D)\}$ in a special case, but also allows us to construct approximations when the special case does not necessarily hold. This is the topic of the next section.

1.3 Decomposition by Demand Relaxation

In this section, we demonstrate how we can build on the special case that we discuss in Section 1.2 to construct tractable approximation strategies for more general cases. The fundamental idea is to relax constraints (1c) in problem (1) by associating the nonnegative Lagrange multipliers $\lambda = \{\lambda^k : k \in \mathcal{K}\}$ with them. In this case, the relaxed version of problem (1) becomes

$$
\begin{aligned}
V^{DDR}(r, D, \lambda) = \max \quad & \sum_{i \in \mathcal{L}} \sum_{k \in \mathcal{K}} [p_i^k - \lambda^k]\, y_i^k + \sum_{k \in \mathcal{K}} \lambda^k\, D^k && && \text{(5a)} \\
\text{subject to} \quad & \sum_{k \in \mathcal{K}} y_i^k \leq r_i && \forall\, i \in \mathcal{L} && \text{(5b)} \\
& y_i^k \leq D^k && \forall\, i \in \mathcal{L},\ k \in \mathcal{K} && \text{(5c)} \\
& y_i^k \in \mathbb{Z}_+ && \forall\, i \in \mathcal{L},\ k \in \mathcal{K}. && \text{(5d)}
\end{aligned}
$$

Constraints (5c) would be redundant in problem (1), but we add them to the problem above to tighten the relaxation. We use the argument λ in the value function $V^{DDR}(r, D, \lambda)$ to emphasize that the optimal objective value of problem (5) depends on the Lagrange multipliers that we associate with the constraints $\sum_{i \in \mathcal{L}} y_i^k \leq D^k$ for all $k \in \mathcal{K}$. The superscript DDR in the value function stands for decomposition by demand relaxation.

Noting that all of the constraints in problem (5) decompose by the locations, we can obtain an optimal solution to this problem by concentrating on one location at a time. In this case, the subproblem for location i has the same form as problem (4) with the exception that the objective function of the subproblem for location i now reads $\sum_{k \in \mathcal{K}} [p_i^k - \lambda^k]\, y_i^k$ instead of $\sum_{k \in \mathcal{K}} p_i^k\, y_i^k$. This implies that we can compute $\mathbb{E}\{V^{DDR}(r, D, \lambda)\}$ in a tractable fashion by using the approach that we describe in the previous section. Furthermore, it is possible to show that the demand relaxation strategy provides an upper bound on the optimal total expected profit. To see the reason, we let $\{\hat{y}_i^k : i \in \mathcal{L},\ k \in \mathcal{K}\}$ be an optimal solution to problem (1) and note that

$$
V(r, D) = \sum_{i \in \mathcal{L}} \sum_{k \in \mathcal{K}} p_i^k\, \hat{y}_i^k \leq \sum_{i \in \mathcal{L}} \sum_{k \in \mathcal{K}} [p_i^k - \lambda^k]\, \hat{y}_i^k + \sum_{k \in \mathcal{K}} \lambda^k\, D^k \leq V^{DDR}(r, D, \lambda). \tag{6}
$$

The first inequality follows from the fact the Lagrange multipliers are nonnegative and we have $\sum_{i \in \mathcal{L}} \hat{y}_i^k \leq D^k$ for all $k \in \mathcal{K}$ by constraints (1c) in problem (1). The second inequality follows from the fact that $\{\hat{y}_i^k : i \in \mathcal{L},\ k \in \mathcal{K}\}$ is a feasible but not necessarily an optimal solution to problem (5). Taking expectations, we obtain $\bar{V}(r) = \mathbb{E}\{V(r, D)\} \leq \mathbb{E}\{V^{DDR}(r, D, \lambda)\}$. Therefore, if we replace $\bar{V}(r)$ in the objective function of problem (2) with $\mathbb{E}\{V^{DDR}(r, D, \lambda)\}$ and solve this problem, then the resulting objective function value $Z^{DDR}(\lambda)$ provides an upper bound on the optimal total expected profit Z, as long as the Lagrange multipliers are nonnegative.

There are several approaches that we can use to choose a good set of Lagrange multipliers. To begin with, we note that $Z^{DDR}(\lambda) \geq Z$ whenever we have $\lambda \geq 0$. Therefore, we can solve the problem $\min_{\lambda \geq 0} Z^{DDR}(\lambda)$ to obtain the tightest possible upper bound on the optimal total expected profit and use the optimal solution to this problem as the Lagrange multipliers. Cheung and Powell (1996*b*) show that $Z^{DDR}(\lambda)$ is a convex function of λ so that the problem $\min_{\lambda \geq 0} Z^{DDR}(\lambda)$ can be solved by using standard subgradient optimization. A possible shortcoming of this approach is that solving the problem $\min_{\lambda \geq 0} Z^{DDR}(\lambda)$ can be time consuming. Kunnumkal and Topaloglu (2008*b*) propose another approach for choosing the Lagrange multipliers that uses the dual solution to problem (3). In particular, they note that constraints (3d) in problem (3) are analogous to constraints (1c) in problem (1). In this case, letting $\hat{\lambda} = \{\hat{\lambda}^k : k \in \mathcal{K}\}$ be the optimal values of the dual variables associated with constraints (3d) in problem (3), they choose the Lagrange multipliers as $\hat{\lambda}$. Since $\hat{\lambda} \geq 0$ by dual feasibility to problem (3), we have $Z^{DDR}(\hat{\lambda}) \geq Z$. Furthermore, Kunnumkal and Topaloglu (2008*b*) show that $Z^{DDR}(\hat{\lambda}) \leq Z^{DLP}$, and we obtain

$$Z \leq \min_{\lambda \geq 0} Z^{DDR}(\lambda) \leq Z^{DDR}(\hat{\lambda}) \leq Z^{DLP}.$$

Therefore, the demand relaxation strategy provides an upper bound on the optimal total expected profit, and this upper bound is tighter than the one provided by the deterministic approximation.

The demand relaxation strategy is proposed by Cheung (1994), but the idea of relaxing complicating constraints in stochastic optimization problems frequently appears in the literature. Ruszczynski (2003) gives an overview of this idea within the framework of general stochastic programs. Adelman and Mersereau (2008) introduce the term "weakly coupled dynamic program" to refer to a stochastic optimization problem which would decompose if a few linking constraints did not exist. They study a dual framework for such stochastic optimization problems. Cheung and Powell (1996*a*) and Cheung and Powell (1996*b*) provide applications of the demand relaxation strategy to dynamic fleet management and inventory distribution problems. Similar relaxation ideas are used by Castanon (1997), Yost and Washburn (2000) and Bertsimas and Mersereau (2007) in partially observable Markov decision processes, by Federgruen and Zipkin (1984), Topaloglu and Kunnumkal (2006) and Kunnumkal and Topaloglu (2008*a*) in inventory distribution and by Erdelyi and Topaloglu (2009*b*), Kunnumkal and Topaloglu (2009*b*) and Topaloglu (2009*b*) in airline network revenue management.

1.4 Decomposition by Supply Relaxation

In this section, we construct a tractable approximation strategy by using a relaxation idea that is conceptually similar to the one in the previous section, but the specific details of this relaxation idea are quite different. We begin by choosing an arbitrary location i and relax constraints (1b) in problem (1) for all other locations by associating the nonnegative Lagrange multipliers $\mu = \{\mu_j : j \in \mathcal{L} \setminus \{i\}\}$ with them. In this case, the relaxed version of problem (1) becomes

$$V_i^{DSR}(r, D, \mu) = \max \quad \sum_{k \in \mathcal{K}} p_i^k \, y_i^k + \sum_{j \in \mathcal{L} \setminus \{i\}} \sum_{k \in \mathcal{K}} [p_j^k - \mu_j] \, y_j^k + \sum_{j \in \mathcal{L} \setminus \{i\}} \mu_j \, r_j \tag{7a}$$

$$\text{subject to} \quad \sum_{k \in \mathcal{K}} y_i^k \leq r_i \tag{7b}$$

$$y_i^k + \sum_{j \in \mathcal{L} \setminus \{i\}} y_j^k \leq D^k \qquad \forall \, k \in \mathcal{K} \tag{7c}$$

$$y_i^k, y_j^k \in \mathbb{Z}_+ \qquad \forall \, j \in \mathcal{L} \setminus \{i\}, \; k \in \mathcal{K}. \tag{7d}$$

The subscript i in the value function $V_i^{DSR}(r, D, \mu)$ emphasizes that the optimal objective value of the problem above depends on the choice of location i, and the superscript DSR stands for decomposition by supply relaxation.

It is possible to show that problem (7) has the same structure as problem (4), in which case, we can compute $\mathbb{E}\{V_i^{DSR}(r, D, \mu)\}$ by using the approach that we describe in Section 1.2. To see the equivalence between problems (4) and (7), we first observe that the decision variables $\{y_j^k : j \in \mathcal{L} \setminus \{i\},\ k \in \mathcal{K}\}$ appear only in constraints (7c) in problem (7). This implies that we can replace the decision variables $\{y_j^k : j \in \mathcal{L} \setminus \{i\}\}$ with a single decision variable, and the objective function coefficient of this decision variable would be $\max_{j \in \mathcal{L} \setminus \{i\}} [p_j^k - \mu_j]$. Therefore, letting $\rho^k = \max_{j \in \mathcal{L} \setminus \{i\}} [p_j^k - \mu_j]$ for notational brevity and replacing the decision variables $\{y_j^k : j \in \mathcal{L} \setminus \{i\}\}$ with the single decision variable z^k, problem (7) can be written as

$$
\begin{aligned}
\max \quad & \sum_{k \in \mathcal{K}} p_i^k \, y_i^k + \sum_{k \in \mathcal{K}} \rho^k \, z^k + \sum_{j \in \mathcal{L} \setminus \{i\}} \mu_j \, r_j && \text{(8a)} \\
\text{subject to} \quad & \sum_{k \in \mathcal{K}} y_i^k \le r_i && \text{(8b)} \\
& y_i^k + z^k \le D^k \qquad \forall\, k \in \mathcal{K} && \text{(8c)} \\
& y_i^k, z^k \in \mathbb{Z}_+ \qquad \forall\, k \in \mathcal{K}. && \text{(8d)}
\end{aligned}
$$

We assume without loss of generality that $\rho^k \ge 0$ for all $k \in \mathcal{K}$. If there exists some $k \in \mathcal{K}$ such that $\rho^k < 0$, then the corresponding decision variable z^k takes value zero in an optimal solution, and this decision variable can be dropped from problem (8) without changing the optimal objective value. Since we have $\rho^k \ge 0$ for all $k \in \mathcal{K}$, the decision variables $\{z^k : k \in \mathcal{K}\}$ take their largest possible values in an optimal solution, and the largest possible values of these decision variables are $\{D^k - y_i^k : k \in \mathcal{K}\}$. Therefore, we can replace the decision variables $\{z^k : k \in \mathcal{K}\}$ in the problem above with $\{D^k - y_i^k : k \in \mathcal{K}\}$ and obtain the equivalent problem

$$
\begin{aligned}
\max \quad & \sum_{k \in \mathcal{K}} [p_i^k - \rho^k] \, y_i^k + \sum_{k \in \mathcal{K}} \rho^k \, D^k + \sum_{j \in \mathcal{L} \setminus \{i\}} \mu_j \, r_j && \text{(9a)} \\
\text{subject to} \quad & \sum_{k \in \mathcal{K}} y_i^k \le r_i && \text{(9b)} \\
& y_i^k \le D^k \qquad \forall\, k \in \mathcal{K} && \text{(9c)} \\
& y_i^k \in \mathbb{Z}_+ \qquad \forall\, k \in \mathcal{K}. && \text{(9d)}
\end{aligned}
$$

Ignoring the constant terms $\sum_{k \in \mathcal{K}} \rho^k \, D^k$ and $\sum_{j \in \mathcal{L} \setminus \{i\}} \mu_j \, r_j$, problem (9) has the same structure as problem (4). By using an argument similar to the one in (6), we can show that $V(r, D) \le V_i^{DSR}(r, D, \mu)$ as long as the Lagrange multipliers are nonnegative; see Zhang and Adelman (2009).

Since the choice of location i is arbitrary, we have $V(r, D) \le V_i^{DSR}(r, D, \mu)$ for all $i \in \mathcal{L}$, and it is not clear which one of $\{V_i^{DSR}(r, D, \mu) : i \in \mathcal{L}\}$ we should use as an approximation to $V(r, D)$. We resolve this ambiguity by averaging over all $i \in \mathcal{L}$ and using $\frac{1}{|\mathcal{L}|} \sum_{i \in \mathcal{L}} V_i^{DSR}(r, D, \mu)$ as an approximation to $V(r, D)$. Noting that $V_i^{DSR}(r, D, \mu)$ is an upper bound on $V(r, D)$ for all $i \in \mathcal{L}$, $\frac{1}{|\mathcal{L}|} \sum_{i \in \mathcal{L}} V_i^{DSR}(r, D, \mu)$ is also an upper bound on $V(r, D)$. In this case, if we replace $\bar{V}(r)$ in the objective function of problem

(2) with $\frac{1}{|\mathcal{L}|}\sum_{i\in\mathcal{L}}\mathbb{E}\{V_i^{DSR}(r,D,\mu)\}$ and solve this problem, then the resulting objective function value $Z^{DSR}(\mu)$ provides an upper bound on the optimal total expected profit Z, as long as we have $\mu \geq 0$.

The methods that we use to choose a good set of Lagrange multipliers are similar to those in Section 1.3. In particular, we can solve the problem $\min_{\mu\geq 0} Z^{DSR}(\mu)$ to obtain the tightest possible upper bound on Z and use the optimal solution to this problem as the Lagrange multipliers. It is possible to show that $Z^{DSR}(\mu)$ is a convex function of μ so that we can solve the problem $\min_{\mu\geq 0} Z^{DSR}(\mu)$ by using standard subgradient optimization. Another approach for choosing a good set of Lagrange multipliers is based on the dual solution of problem (3). We note that constraints (1b) in problem (1) capture the resource availabilities at different locations at the beginning of the second time period. Similarly, constraints (3c) in problem (3) capture the resource availabilities at different locations. Therefore, letting $\{\hat{\mu}_i : i \in \mathcal{L}\}$ be the optimal values of the dual variables associated with constraints (3c) in problem (3), we can choose the Lagrange multipliers as $\{\hat{\mu}_j : j \in \mathcal{L}\setminus\{i\}\}$.

The supply relaxation strategy has its roots in the network revenue management literature. In this literature, it is customary to relax the capacity constraints for all of the flight legs except for one and to obtain approximations to the value function by concentrating on one flight leg at a time; see Talluri and van Ryzin (2005). Zhang and Adelman (2009) establish that the supply relaxation strategy provides an upper bound on the value function. Kunnumkal and Topaloglu (2008*c*) and Kunnumkal and Topaloglu (2009*a*) provide alternative proofs for the same result. Erdelyi and Topaloglu (2009*a*) use the supply relaxation idea to solve an overbooking problem. Topaloglu (2009*a*) extends the supply relaxation strategy to general stochastic programs and provides applications in dynamic fleet management.

1.5 Separable Projective Approximations

In this section, we describe an iterative approach that uses sampled realizations of the demand arrivals to construct a separable approximation to the value function. The fundamental idea is to approximate the value function $\bar{V}(r)$ with a separable function of the form

$$V^{SPA}(r) = \sum_{i\in\mathcal{L}} V_i^{SPA}(r_i),$$

where the approximations $\{V_i^{SPA}(\cdot) : i \in \mathcal{L}\}$ are single dimensional, piecewise linear and concave functions with points of nondifferentiability being a subset of integers. The superscript SPA stands for separable projective approximations. The concavity of the approximation $V_i^{SPA}(\cdot)$ is intended to capture the intuitive expectation that the marginal benefit from an additional resource at a location should be decreasing as we have more resources available at the location. Noting that $V_i^{SPA}(\cdot)$ is a piecewise linear function with points of nondifferentiability being a subset of integers, if we have a total of Q available resources, then we can characterize the approximation $V_i^{SPA}(\cdot)$ with a sequence of Q slopes $\{v_i(q) : q = 0,\dots,Q-1\}$, where $v_i(q)$ is the slope of $V_i^{SPA}(\cdot)$ over the interval $[q, q+1]$. In other words, we have $v_i(q) = V_i^{SPA}(q+1) - V_i^{SPA}(q)$.

We use an iterative approach to construct the approximations $\{V_i^{SPA}(\cdot) : i \in \mathcal{L}\}$. In particular, letting $\{V_i^{SPA,n}(\cdot) : i \in \mathcal{L}\}$ be the approximations at iteration n, we begin by solving problem (2)

Step 1. Initialize the iteration counter by letting $n = 1$. Initialize the approximations $\{V_i^{SPA,1}(\cdot) : i \in \mathcal{L}\}$ to arbitrary piecewise linear and concave functions with points of nondifferentiability being a subset of integers.

Step 2. Solve problem (2) after replacing $\bar{V}(r)$ in the objective function with the value function approximation $\sum_{i \in \mathcal{L}} V_i^{SPA,n}(r_i)$. Let $\{\hat{x}_{ij}^n : i, j \in \mathcal{L}\}$, $\{\hat{r}_i^n : i \in \mathcal{L}\}$ be an optimal solution.

Step 3. Sample a realization of the demand arrivals and denote this sample by $\{\hat{D}^{k,n} : k \in \mathcal{K}\}$. Solve problem (1) after replacing the right side of constraints (1b) with $\hat{r}^n = \{\hat{r}_i^n : i \in \mathcal{L}\}$ and the right side of constraints (1c) with $\hat{D}^n = \{\hat{D}^{k,n} : k \in \mathcal{K}\}$.

Step 4. Use the solution to problem (1) to update the approximations $\{V_i^{SPA,n}(\cdot) : i \in \mathcal{L}\}$ and obtain the approximations $\{V_i^{SPA,n+1}(\cdot) : i \in \mathcal{L}\}$ that is used at the next iteration. Increase n by one and go to Step 2.

Figure 1: An iterative approach to construct a separable value function approximation.

after replacing $\bar{V}(r)$ in the objective function with $\sum_{i \in \mathcal{L}} V_i^{SPA,n}(r_i)$. This provides an optimal solution $\{\hat{x}_{ij}^n : i, j \in \mathcal{L}\}$, $\{\hat{r}_i^n : i \in \mathcal{L}\}$. We then sample a realization of the demand arrivals, which we denote by $\{\hat{D}^{k,n} : k \in \mathcal{K}\}$. We solve problem (1) after replacing the right side of constraints (1b) with $\hat{r}^n = \{\hat{r}_i^n : i \in \mathcal{L}\}$ and the right side of constraints (1c) with $\hat{D}^n = \{\hat{D}^{k,n} : k \in \mathcal{K}\}$. The challenge is to use the solution to problem (1) to improve the approximations $\{V_i^{SPA,n}(\cdot) : i \in \mathcal{L}\}$. We give a description of the general approach in Figure 1. In Step 1, we initialize the approximations $\{V_i^{SPA,n}(\cdot) : i \in \mathcal{L}\}$ to arbitrary piecewise linear and concave functions with points of nondifferentiability being a subset of integers. In Step 2, we solve problem (2) after replacing $\bar{V}(r)$ in the objective function with the value function approximation $\sum_{i \in \mathcal{L}} V_i^{SPA,n}(r_i)$. In Step 3, we sample a realization of the demand arrivals and solve problem (1) by using the sampled demand arrivals and the resource availabilities obtained from problem (2). In Step 4, we update the approximations $\{V_i^{SPA,n}(\cdot) : i \in \mathcal{L}\}$, and this provides the approximations $\{V_i^{SPA,n+1}(\cdot) : i \in \mathcal{L}\}$ that we use at the next iteration.

The approach that we use to update the approximations is based on a stochastic approximation idea, but we need to be careful to preserve the concavity of the value function approximation. Using e_i to denote the $|\mathcal{L}|$ dimensional unit vector with a one in the element corresponding to $i \in \mathcal{L}$, we let $\hat{\vartheta}_i^n = V(\hat{r}^n + e_i, \hat{D}^n) - V(\hat{r}^n, \hat{D}^n)$. Therefore, $\hat{\vartheta}_i^n$ captures the marginal benefit from an additional resource at location i in problem (1) when the number of resources at different locations is given by $\hat{r}^n = \{\hat{r}_i^n : i \in \mathcal{L}\}$ and the realization of the demand arrivals is given by $\hat{D}^n = \{\hat{D}^{k,n} : k \in \mathcal{K}\}$. In this case, letting $V_i^{SPA,n}(\cdot)$ be the piecewise linear and concave function characterized by the sequence of slopes $\{v_i^n(q) : q = 0, \dots, Q-1\}$, we update these slopes as

$$\gamma_i^n(q) = \begin{cases} [1 - \alpha^n]\, v_i^n(q) + \alpha^n\, \hat{\vartheta}_i^n & \text{if } q = \hat{r}_i^n \\ v_i^n(q) & \text{otherwise,} \end{cases}$$

where $\alpha^n \in [0,1]$ is the step size parameter at iteration n. Since we may not have $\gamma_i^n(0) \geq \gamma_i^n(1) \geq \ldots \geq \gamma_i^n(Q-1)$, the piecewise linear function characterized by the sequence of slopes $\{\gamma_i^n(q) : q = 0, \ldots, Q-1\}$ is not necessarily concave. To obtain a concave approximation, we solve the problem

$$\begin{aligned} \min \quad & \sum_{q=0}^{Q-1} [w(q) - \gamma_i^n(q)]^2 \\ \text{subject to} \quad & w(q) \geq w(q+1) \qquad \forall\, q = 0, \ldots, Q-2. \end{aligned}$$

Letting $\{\hat{w}(q) : q = 0, \ldots, Q-1\}$ be an optimal solution to the problem above, we use the sequence of slopes $\{\hat{w}(q) : q = 0, \ldots, Q-1\}$ to characterize the approximation $V_i^{SPA,n+1}(\cdot)$. Since we have $\hat{w}(0) \geq \hat{w}(1) \geq \ldots \geq \hat{w}(Q-1)$, the approximation $V_i^{SPA,n+1}(\cdot)$ is indeed concave. Intuitively speaking, the quadratic program above projects the piecewise linear function characterized by the sequence of slopes $\{\gamma_i^n(q) : q = 0, \ldots, Q-1\}$ onto the set of piecewise linear and concave functions with points of nondifferentiability being a subset of integers. Throughout this chapter, we use $V_i^{SPA,n+1}(\cdot) = \mathcal{U}(V_i^{SPA,n}(\cdot), \hat{r}_i^n, \hat{\vartheta}_i^n, \alpha^n)$ to succinctly capture the updating procedure.

The separable projective approximations are proposed by Powell, Ruszczynski and Topaloglu (2004). The authors develop several convergence results for the updating procedure that we describe above, and these convergence results have close connections with the stochastic approximation theory. Topaloglu and Powell (2006) and Topaloglu (2006) use the separable projective approximations for dynamic fleet management problems. Godfrey and Powell (2001) use separable approximations to the value functions in a product distribution problem, but the method they use to update the approximations is different from the one that we describe above. Godfrey and Powell (2002*a*) and Godfrey and Powell (2002*b*) apply the approach in Godfrey and Powell (2001) to dynamic fleet management problems. Cheung and Powell (2000) and Topaloglu and Powell (2003) propose alternative methods to construct separable and concave value function approximations. Both of these methods are rooted in the stochastic approximation theory as well. Cheung and Chen (1998) apply the method that is developed in Cheung and Powell (2000) to an empty container allocation problem. Stochastic approximation methods date back to Robbins and Monro (1951). Kushner and Clark (1978), Benveniste, Metivier and Priouret (1991), Bertsekas and Tsitsiklis (1996) and Kushner and Yin (1997) provide detailed coverage of the stochastic approximation theory.

2 Multi-Stage Problems

In this section, we demonstrate how the ideas that we discuss in Section 1 can be extended to more complex resource management problems that take place over multiple time periods. We consider a setting where we have a set of resources that can be used to serve the demands that arrive randomly over time. At each time period, we observe the realization of the demand arrivals, and we need to decide which demands we should serve and to which locations we should reposition the unused resources. The objective is to maximize the total expected profit over a finite planning horizon.

The problem takes place over the planning horizon $\mathcal{T} = \{1, \ldots, \tau\}$. For notational brevity, we assume that the travel times between all pairs of locations are one time period. We continue using $\mathcal{L}$

to denote the set of locations in the transportation network. To capture the decisions, we let x_{ijt} be the number of resources that we reposition from location i to j at time period t and y^k_{it} be the number of demands of type k that we serve by using a resource at location i at time period t. The cost of repositioning one resource from location i to j is c_{ij}. If we serve a demand of type k by using a resource at location i, then we generate a revenue of p^k_i, and the resource ends up at location δ^k at the next time period. Similar to the convention that we use in Section 1, if it is infeasible to serve a demand of type k by using a resource at location i, then we assume that $p^k_i = -\infty$. We use the random variable D^k_t to represent the number of demands of type k that arrive at time period t. We assume that the demand arrivals at different time periods are independent. We let r_{it} be the number of resources available at location i at the beginning of time period t. We note that $\{r_{i1} : i \in \mathcal{L}\}$ is a part of the problem data, whereas $\{r_{it} : i \in \mathcal{L}\}$ for $t \in \mathcal{T} \setminus \{1\}$ is determined by the decisions at time period $t-1$.

We can formulate the problem as a dynamic program by using $r_t = \{r_{it} : i \in \mathcal{L}\}$ as the state variable at time period t. Letting $\mathbf{1}(\cdot)$ be the indicator function, and using $x_t = \{x_{ijt} : i, j \in \mathcal{L}\}$ and $y_t = \{y^k_{it} : i \in \mathcal{L},\ k \in \mathcal{K}\}$ to capture the decisions and $D_t = \{D^k_t : k \in \mathcal{K}\}$ to capture the demand arrivals at time period t, the set of feasible decisions and state transitions are given by

$$\mathcal{X}(r_t, D_t) = \Big\{ (x_t, y_t, r_{t+1}) : \sum_{j \in \mathcal{L}} x_{ijt} + \sum_{k \in \mathcal{K}} y^k_{it} = r_{it} \qquad \forall\, i \in \mathcal{L} \tag{10a}$$

$$\sum_{i \in \mathcal{L}} x_{ijt} + \sum_{i \in \mathcal{L}} \sum_{k \in \mathcal{K}} \mathbf{1}(\delta^k = j)\, y^k_{it} - r_{j,t+1} = 0 \quad \forall\, j \in \mathcal{L} \tag{10b}$$

$$\sum_{i \in \mathcal{L}} y^k_{it} \le D^k_t \qquad \forall\, k \in \mathcal{K} \tag{10c}$$

$$x_{ijt}, y^k_{it}, r_{j,t+1} \in \mathbb{Z}_+ \qquad \forall\, i, j \in \mathcal{L},\ k \in \mathcal{K} \Big\}. \tag{10d}$$

Constraints (10a) ensure that the number of resources that we reposition between different locations and use to serve the demands does not violate the resource availabilities. Constraints (10b) compute the number of resources available at different locations at the beginning of the next time period. Constraints (10c) ensure that the number of demands that we serve does not exceed the demand arrivals. We can find the optimal policy by computing the value functions $\{\bar{V}_t(\cdot) : t \in \mathcal{T}\}$ through the optimality equation

$$\bar{V}_t(r_t) = \mathbb{E}\Big\{ \max_{(x_t, y_t, r_{t+1}) \in \mathcal{X}(r_t, D_t)} - \sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{L}} c_{ij}\, x_{ijt} + \sum_{i \in \mathcal{L}} \sum_{k \in \mathcal{K}} p^k_i\, y^k_{it} + \bar{V}_{t+1}(r_{t+1}) \Big\}$$

with the boundary condition that $\bar{V}_{\tau+1}(\cdot) = 0$; see Puterman (1994). The optimality equation above involves a high dimensional state variable, which makes it difficult to compute the value functions for all possible values of the vector r_t. This difficulty is further compounded by the fact that the decisions (x_t, y_t) in the problem above are high dimensional vectors, and we take an expectation that involves the high dimensional demand arrivals D_t. Powell (2007) uses the term "three curses of dimensionality" to refer to the difficulty due to the sizes of the state, decision and demand vectors.

It turns out that all of the ideas in Section 1 can be extended to multi-stage problems. For economy of space, we only show extensions of the deterministic linear program and the separable projective approximations. Cheung and Powell (1996*b*) extend the demand relaxation strategy and Topaloglu (2009*a*) extends the supply relaxation strategy to multi-stage problems.

2.1 Deterministic Linear Program

Assuming that all demand arrivals take on their expected values, we can approximate the total expected profit over the planning horizon by using the optimal objective value of the problem

$$
\begin{aligned}
\max \quad & -\sum_{t\in\mathcal{T}}\sum_{i\in\mathcal{L}}\sum_{j\in\mathcal{L}} c_{ij}\, x_{ijt} + \sum_{t\in\mathcal{T}}\sum_{i\in\mathcal{L}}\sum_{k\in\mathcal{K}} p_i^k\, y_{it}^k && && \text{(11a)}\\
\text{subject to} \quad & \sum_{j\in\mathcal{L}} x_{ijt} + \sum_{k\in\mathcal{K}} y_{it}^k - r_{it} = 0 && \forall\, i\in\mathcal{L},\ t\in\mathcal{T} && \text{(11b)}\\
& \sum_{i\in\mathcal{L}} x_{ijt} + \sum_{i\in\mathcal{L}}\sum_{k\in\mathcal{K}} \mathbf{1}(\delta^k = j)\, y_{it}^k - r_{j,t+1} = 0 && \forall\, j\in\mathcal{L},\ t\in\mathcal{T}\setminus\{\tau\} && \text{(11c)}\\
& \sum_{i\in\mathcal{L}} y_{it}^k \le \mathbb{E}\{D_t^k\} && \forall\, k\in\mathcal{K},\ t\in\mathcal{T} && \text{(11d)}\\
& x_{ijt}, y_{it}^k, r_{it} \in \mathbb{R}_+ && \forall\, i,j\in\mathcal{L},\ k\in\mathcal{K},\ t\in\mathcal{T}. && \text{(11e)}
\end{aligned}
$$

Constraints (11b), (11c) and (11d) in problem (11) are respectively analogous to constraints (3b), (3c) and (3d) in problem (3). In the problem above, our understanding is that the values of the decision variables $\{r_{i1} : i \in \mathcal{L}\}$ are fixed by the number of resources that we have at different locations at the beginning of the planning horizon. By following the same hindsight approach that we use in Section 1.1, it is possible to show that the optimal objective value of problem (11) provides an upper bound on the optimal total expected profit. Such an upper bound becomes useful when we want to assess the optimality gaps of suboptimal approximation strategies.

We need minor modifications in problem (11) when we use this problem to make the resource repositioning and demand service decisions dynamically over time. In particular, if we are at time period t and the demand arrivals at the current time period are given by $\{D_t^k : k \in \mathcal{K}\}$, then we replace the right side of constraints (11d) for time period t with the current realization of the demand arrivals $\{D_t^k : k \in \mathcal{K}\}$. For the other time periods, we keep using the expected values of the demand arrivals. In addition, we fix the values of the decision variables $\{r_{it} : i \in \mathcal{L}\}$ to reflect the number of resources that we have at different locations at the current time period. Finally, we replace the set of time periods $\mathcal{T}$ with $\{t, \ldots, \tau\}$ so that we only concentrate on the time periods in the future. Solving problem (11) after these modifications ensures that the decisions that we obtain for time period t are feasible for the current state of the resources and the current realization of the demand arrivals.

2.2 Separable Projective Approximations

The idea behind the separable projective approximations is to construct approximations to the value functions by using sampled trajectories of the system. In particular, assuming that we have a set of approximations $\{\hat{V}_t(\cdot) : t \in \mathcal{T}\}$ to the value functions $\{\bar{V}_t(\cdot) : t \in \mathcal{T}\}$, if the number of resources at different locations is given by r_t and the realization of the demand arrivals is given by D_t, then we can make the decisions at time period t by solving the problem

$$
\max_{(x_t, y_t, r_{t+1}) \in \mathcal{X}(r_t, D_t)} -\sum_{i\in\mathcal{L}}\sum_{j\in\mathcal{L}} c_{ij}\, x_{ijt} + \sum_{i\in\mathcal{L}}\sum_{k\in\mathcal{K}} p_i^k\, y_{it}^k + \hat{V}_{t+1}(r_{t+1}).
$$

Therefore, repeated solutions of the problem above for the successive time periods in the planning horizon would allow us to simulate the trajectory of the system under a particular realization of the demand arrivals. The challenge is to use the information that we obtain from the simulated trajectories of the system to update and improve the value function approximations.

We use the separable functions $\{V_t^{SPA}(\cdot) : t \in \mathcal{T}\}$ as our value function approximations. Similar to the development in Section 1.5, each one of these value function approximations is of the form $V_t^{SPA}(r_t) = \sum_{i \in \mathcal{L}} V_{it}^{SPA}(r_{it})$, where $V_{it}^{SPA}(\cdot)$ is a single dimensional, piecewise linear and concave function with points of nondifferentiability being a subset of integers. Letting Q be the total number of available resources, we can characterize $V_{it}^{SPA}(\cdot)$ by the sequence of slopes $\{v_{it}(q) : q = 0, \ldots, Q-1\}$, where $v_{it}(q)$ is the slope of $V_{it}^{SPA}(\cdot)$ over the interval $[q, q+1]$.

The updating procedure in Section 1.5, together with the idea of simulating the trajectories of the system, immediately provides an iterative method that can be used to construct separable value function approximations. We describe this method in Figure 2. In Step 1, we initialize the value function approximations. In Step 2, we sample a realization of the demand arrivals over the whole planning horizon. In Step 3, we simulate the evolution of the system under the current demand arrivals and the current value function approximations. In Step 4, we update the value function approximations to obtain the value function approximations that we use at the next iteration. The updating procedure $\mathcal{U}(V_{it}^{SPA,n}(\cdot), \hat{r}_{it}^n, \hat{\vartheta}_{it}^n, \alpha^n)$ in this step is identical to the one that we describe in Section 1.5.

3 A Flexible Modeling Framework

The resource management problems that we have considered thus far are relatively simplistic. In particular, the only distinguishing characteristic of a resource is its location, whereas the only distinguishing characteristic of a demand is its type. Our motivation for working with somewhat simplistic resource management problems is that such problems allow us to demonstrate the crucial algorithmic issues more clearly, but practical resource management problems can easily take on much more complex characteristics. For example, a driver resource in a driver scheduling application can be described by its inbound location, time to reach the inbound location, duty time within the shift, days away from home, vehicle type and home domicile. Conceptually, we can model such a complex resource by using multiple indices in our decision and state variables to capture the different attributes of the resource, but this approach would clearly be too cumbersome. In this section, we describe a flexible paradigm that is useful for modeling complex resource management problems.

3.1 Resources

To represent the resources, we let $\mathcal{A}$ be the attribute space of the resources. Roughly speaking, the attribute space represents the set of all possible states of a particular resource. We refer to each element of the attribute space as an attribute vector. We let r_{at} be the number of resources with attribute vector a at time period t. For example, if we assume that the set of locations is $\mathcal{L}$, the set of vehicle types is $\mathcal{V}$ and the travel time between any pair of locations is bounded by $\bar{T}$, then the attribute space for

Step 1. Initialize the iteration counter by letting $n = 1$. Initialize the approximations $\{V_{it}^{SPA,1}(\cdot) : i \in \mathcal{L}\}$ to arbitrary piecewise linear and concave functions with points of nondifferentiability being a subset of integers.

Step 2. Sample a realization of the demand arrivals over the whole planning horizon and denote this sample by $\{\hat{D}_t^{k,n} : k \in \mathcal{K},\ t \in \mathcal{T}\}$.

Step 3. Simulate the evolution of the system under the current demand arrivals and the current value function approximations.

Step 3.a. Initialize the time by letting $t = 1$. Initialize the state of the resources $\hat{r}_1^n = \{\hat{r}_{i1}^n : i \in \mathcal{L}\}$ to reflect the initial state.

Step 3.b. Solve the problem

$$\tilde{V}_t(\hat{r}_t^n, \hat{D}_t^n) = \max_{(x_t, y_t, r_{t+1}) \in \mathcal{X}(\hat{r}_t^n, \hat{D}_t^n)} -\sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{L}} c_{ij}\, x_{ijt} + \sum_{i \in \mathcal{L}} \sum_{k \in \mathcal{K}} p_i^k\, y_{it}^k + \hat{V}_{t+1}^{SPA,n}(r_{t+1})$$

to obtain the optimal solution $(\hat{x}_t^n, \hat{y}_t^n, \hat{r}_{t+1}^n)$. Compute and store $\hat{\vartheta}_{it}^n = \tilde{V}_t(\hat{r}_t^n + e_i, \hat{D}_t^n) - \tilde{V}_t(\hat{r}_t^n, \hat{D}_t^n)$ for all $i \in \mathcal{L}$.

Step 3.c. Increase t by one. If $t < \tau$, then go to Step 3.b.

Step 4. Using α^n to denote the step size parameter at iteration n, update the value function approximations by letting $V_{it}^{SPA,n+1}(\cdot) = \mathcal{U}(V_{it}^{SPA,n}(\cdot), \hat{r}_{it}^n, \hat{\vartheta}_{it}^n, \alpha^n)$ for all $i \in \mathcal{L}$, $t \in \mathcal{T}$. Increase n by one and go to Step 2.

Figure 2: An iterative approach to construct separable value function approximations when there are multiple time periods in the planning horizon.

the vehicles in a simple dynamic fleet management application is $\mathcal{A} = \mathcal{L} \times \mathcal{V} \times \{0, 1, \dots, \bar{T}\}$. A vehicle with the attribute vector $a = (a_1, a_2, a_3) = (\text{New York, Large}, 4) \in \mathcal{A}$ is a vehicle of type large that is inbound to New York and will reach New York in four time periods. We can access the number of resources with attribute vector a at time period t by referring to r_{at}. This implies that we can lump the resources with the same attribute vector together and treat them as indistinguishable.

We use a similar approach to represent the demands. We let $\mathcal{B}$ be the attribute space of the demands and D_{bt} be the number of demands with attribute vector b at time period t. For example, if the set of expedition options is $\mathcal{E}$, then the attribute space for the demands may be $\mathcal{B} = \mathcal{L} \times \mathcal{L} \times \mathcal{E}$. A demand with the attribute vector $b = (b_1, b_2, b_3) = (\text{New York, Delaware, Express}) \in \mathcal{B}$ is an express demand with origin location New York and destination location Delaware. It may be not possible to serve an express demand from New York to Delaware by using a vehicle in New Jersey, but if the demand is not express, then it may be the case that a vehicle in New Jersey can serve this demand. Therefore, the characteristics of the demand may dictate the feasible resource assignments.

3.2 Decisions

We use two classes of decisions. The class of decisions $\mathcal{D}^R$ captures the repositioning decisions, whereas the class of decisions $\mathcal{D}^S$ captures the demand service decisions. We let $\mathcal{D} = \mathcal{D}^R \cup \mathcal{D}^S$ to capture the set of all decision classes and refer to each element of $\mathcal{D}$ as a decision. For example, we may have $\mathcal{D}^R = \mathcal{L} \times \mathcal{L}$, and the decision $d = (d_1, d_2) = (\text{New York, Delaware}) \in \mathcal{D}^R$ may represent the decision to reposition a resource from New York to Delaware. Similarly, we may have $\mathcal{D}^S = \mathcal{B}$, and the decision $d = (b_1, b_2, b_3) \in \mathcal{D}^S = \mathcal{B}$ may represent the decision to serve a demand with the attribute vector b. It is reasonable to assume that for each decision $d \in \mathcal{D}^S$, there exists an attribute vector $b_d \in \mathcal{B}$ such that the decision d corresponds to serving a demand with attribute vector b_d.

We use x_{adt} to denote the number of resources with attribute vector a that we modify by using decision d at time period t. If we modify a resource with attribute vector a by using decision d at time period t, then we generate a profit contribution of c_{adt}. If it is not feasible to apply decision d on a resource with attribute vector a, then we capture this situation by letting $c_{adt} = -\infty$. In this case, we can capture the decisions at time period t by $x_t = \{x_{adt} : a \in \mathcal{A},\ d \in \mathcal{D}\}$, the state of the resources by $r_t = \{r_{at} : a \in \mathcal{A}\}$ and the state of the demands by $D_t = \{D_{bt} : b \in \mathcal{B}\}$.

3.3 Optimization Problem

The set of feasible decisions at time period t is given by

$$\mathcal{X}(r_t, D_t) = \Big\{ x_t : \sum_{d \in \mathcal{D}} x_{adt} = r_{at} \qquad \forall\, a \in \mathcal{A} \tag{12a}$$

$$\sum_{a \in \mathcal{A}} x_{adt} \leq D_{b_d,t} \quad \forall\, d \in \mathcal{D}^S \tag{12b}$$

$$x_{adt} \in \mathbb{Z}_+ \qquad \forall\, a \in \mathcal{A},\ d \in \mathcal{D} \Big\}, \tag{12c}$$

where we use the fact that for each decision $d \in \mathcal{D}^S$, there exists an attribute vector $b_d \in \mathcal{B}$ such that the decision d corresponds to serving a demand with attribute vector b_d. We note that constraints (12a) and (12b) are respectively analogous to constraints (10a) and (10c). We capture the outcome of applying a decision d on a resource with attribute vector a by defining

$$\delta_{a'}(a,d) \begin{cases} 1 & \text{if applying the decision } d \text{ on a resource with attribute vector } a \text{ transforms the resource} \\ & \quad \text{into a resource with attribute vector } a' \\ 0 & \text{otherwise.} \end{cases}$$

The precise definition of $\delta_{a'}(a,d)$ is naturally dependent on the physics of the problem and how the system reacts to decisions. Using the definition above, we can write the dynamics of the resources as

$$r_{a,t+1} = \sum_{a' \in \mathcal{A}} \sum_{d \in \mathcal{D}} \delta_a(a',d)\, x_{a'dt} \qquad \forall\, a \in \mathcal{A},$$

which are analogous to constraints (10b). The implicit assumption in the definition of $\delta_{a'}(a,d)$ is that the current attribute vector of the resource and the decision that we apply on the resource deterministically give the future attribute vector of the resource. This is indeed the case in many resource management applications, but there may be settings with random travel times and equipment failures where this assumption is not necessarily satisfied.

A Markovian policy π is a collection of decision functions $\{X_t^\pi(\cdot) : t \in \mathcal{T}\}$ such that each $X_t^\pi(\cdot)$ maps the state of the system (r_t, D_t) to a decision vector $x_t \in \mathcal{X}(r_t, D_t)$. Our goal is to find a policy that maximizes the total expected profit over the planning horizon. In other words, using $C_t(x_t) = \sum_{a \in \mathcal{A}} \sum_{d \in \mathcal{D}} c_{adt}\, x_{adt}$ to denote the total profit associated with the decisions x_t at time period t and Π to denote the set of all Markovian policies, we want to solve the problem

$$\sup_{\pi \in \Pi} \mathbb{E}\left\{ \sum_{t \in \mathcal{T}} C_t(X_t^\pi(r_t, D_t)) \right\}.$$

The formulation above provides a compact representation of the resource management problem that we are interested in, but it does not help from an algorithmic perspective. In particular, it is impossible to make a direct search over a reasonably broad class of policies Π. Therefore, the modeling framework that we describe in this section is useful for communicating the problem, but the actual solution requires a concrete algorithmic approach such as the ones described in Sections 1 and 2.

The modeling framework that we describe in this section is due to Shapiro (1999) and Powell, Shapiro and Simão (2001). These works are intended to capture a broad class of resource management problems arising from a variety of application settings, and their presentation is significantly larger in scope than ours. Powell (2007) uses this modeling framework to come up with a sound formulation for general dynamic programs, with a particular emphasis on the evolution of information over time. Shapiro and Powell (2006) demonstrate how to use the modeling framework within the context of several algorithmic approaches to decompose large scale resource management problems. Powell, Shapiro and Simão (2002), Powell and Topaloglu (2003), Powell, Wu, Simão and Whisman (2004), Powell and Topaloglu (2005), Schenk and Klabjan (2008) and Simao, Day, George, Gifford, Nienow and Powell (2009) apply the modeling framework to dynamic fleet management, military airlift operations, express package routing and driver scheduling problems. We note that the presentation in this chapter uses demands as the only source of uncertainty. The driver scheduling application in Simao et al. (2009) is particularly interesting as it models a variety of sources of uncertainty, such as random travel times, equipment failures and driver availability.

4 Numerical Illustrations

In this section, we numerically compare the performances of the different approximation strategies that we describe in Sections 1 and 2. Our goal here is to give a feel for the relative performances of the different approximation strategies, rather than to present a full scale experimental study.

4.1 Numerical Results for Two-Stage Problems

In this section, we provide numerical results for two-stage problems. The test problems that we use in our experimental setup have the same structure as the problem that we describe in Section 1. We generate ten locations uniformly over a 100×100 region. Using $d(i,j)$ to denote the Euclidean distance between locations i and j, the cost of repositioning a resource from location i to j is given by $d(i,j)$. We have ten demand types and each demand type is associated with a particular location. The revenue from using a resource at location i to serve a demand type that is associated with location j is given by $200 - \rho \times d(i,j)$, where ρ is a parameter that we vary. As ρ gets larger, it becomes less profitable to serve a demand type that is associated with a particular location by using a resource at another location. In our computational experiments, we vary the coefficient of variation of the demand random variables and the parameter ρ. In particular, letting v be the coefficient of variation of the demand random variables, we label our test problems by the pair $(v, \rho) \in \{0.25, 0.5, 0.75, 1.0\} \times \{0.25, 0.5, 0.75, 1.0\}$. We note that the numerical results that we present in this section do not appear in the previous literature.

We test the performances of the approximation strategies described in Sections 1.1, 1.3, 1.4 and 1.5, which we respectively refer to as DLP, DDR, DSR and SPA. DLP immediately provides a set of first stage decisions $\{\hat{x}_{ij} : i, j \in \mathcal{L}\}$. On the other hand, DDR, DSR and SPA provide value function approximations that are all separable functions of the form $\sum_{i \in \mathcal{L}} \hat{V}_i(\cdot)$. To obtain a set of first stage decisions $\{\hat{x}_{ij} : i, j \in \mathcal{L}\}$, DDR, DSR and SPA replace $\bar{V}(r)$ in the objective function of problem (2) with $\sum_{i \in \mathcal{L}} \hat{V}_i(r_i)$ and solve this problem. For DDR, DSR and SPA, it is possible to show that each one of $\{\hat{V}_i(\cdot) : i \in \mathcal{L}\}$ is a piecewise linear and concave function with points of nondifferentiability being a subset of integers. In this case, we can formulate problem (2) with the value function approximation $\sum_{i \in \mathcal{L}} \hat{V}_i(\cdot)$ as a minimum cost network flow problem; see Nemhauser and Wolsey (1988). A set of first stage decisions allow us to compute the cost that we incur at the first stage. In addition, we can compute the number of resources at different locations at the beginning of the second stage by $\hat{r}_j = \sum_{i \in \mathcal{L}} \hat{x}_{ij}$ for all $j \in \mathcal{L}$. To check the quality of the first stage decisions, we replace the right side of constraints (1b) in problem (1) with $\{\hat{r}_i : i \in \mathcal{L}\}$ and solve this problem for many realizations of demand arrivals. In this way, we approximate the total expected revenue in the second stage. The total expected profit is the difference between the total cost and the total expected revenue in the two stages.

Problem	Upper Bound			DDR vs.		Total Expected Profit				SPA vs.		
(v, ρ)	DLP	DDR	DSR	DLP	DSR	DLP	DDR	DSR	SPA	DLP	DDR	DSR
$(0.25, 0.25)$	18,292	17,628	18,187	3.76	3.17	16,782	16,881	16,897	16,908	0.74	0.16	0.06
$(0.25, 0.5)$	18,292	17,444	18,179	4.86	4.21	16,621	16,846	16,820	16,794	1.03	-0.31	-0.16
$(0.25, 0.75)$	18,292	17,287	18,167	5.81	5.09	16,478	16,689	16,750	16,765	1.71	0.45	0.09
$(0.25, 1.0)$	18,292	17,164	18,193	6.57	6.00	16,339	16,581	16,717	16,724	2.30	0.85	0.04
$(0.5, 0.25)$	18,275	16,744	18,035	9.15	7.71	15,213	15,551	15,519	15,553	2.18	0.01	0.21
$(0.5, 0.5)$	18,275	16,307	18,004	12.07	10.40	14,890	15,345	15,311	15,321	2.81	-0.16	0.07
$(0.5, 0.75)$	18,275	15,932	17,972	14.70	12.81	14,605	15,020	15,127	15,147	3.58	0.84	0.13
$(0.5, 1.0)$	18,275	15,627	17,949	16.94	14.86	14,328	14,789	14,949	14,971	4.29	1.21	0.14
$(0.75, 0.25)$	18,285	15,672	17,856	16.67	13.94	13,713	14,335	14,269	14,329	4.30	-0.04	0.42
$(0.75, 0.5)$	18,285	15,070	17,812	21.33	18.19	13,158	13,848	13,924	13,924	5.50	0.55	0.00
$(0.75, 0.75)$	18,285	14,527	17,772	25.87	22.34	12,627	13,501	13,574	13,576	6.99	0.55	0.01
$(0.75, 1.0)$	18,285	14,048	17,747	30.16	26.33	12,117	13,207	13,302	13,306	8.94	0.74	0.03
$(1.0, 0.25)$	18,260	13,868	17,484	31.67	26.07	11,990	12,695	12,578	12,714	5.69	0.14	1.07
$(1.0, 0.5)$	18,260	13,084	17,484	39.56	33.63	11,228	12,093	12,153	12,159	7.65	0.54	0.05
$(1.0, 0.75)$	18,260	12,349	17,484	47.86	41.58	10,481	11,583	11,632	11,637	9.94	0.47	0.04
$(1.0, 1.0)$	18,260	11,736	17,484	55.59	48.98	9,764	11,077	11,109	11,143	12.37	0.59	0.30

Table 1: Numerical results for two-stage problems.

Table 1 summarizes our numerical results. The first column in this table shows the problem characteristics by using the pair (v, ρ). In Sections 1.1, 1.3 and 1.4, we show that DLP, DDR and DSR provide upper bounds on the optimal total expected profit. The second, third and fourth columns in Table 1 respectively show the upper bounds obtained by these three strategies. The fifth and sixth columns show the percent gap between the upper bounds obtained by DDR and the remaining two strategies. DDR provides the tightest upper bounds and we use it as a reference point. The seventh, eighth, ninth and tenth columns respectively show the total expected profits obtained by DLP, DDR, DSR and SPA. Finally, the eleventh, twelfth and thirteenth columns show the percent gaps between the total expected profits obtained by SPA and the remaining three strategies. SPA generally obtains the highest total expected profits and we use it as a reference point.

Our results indicate that the gaps between the upper bounds obtained by DDR and the other two strategies can be quite significant. We observe that the gaps between the upper bounds get larger as the coefficient of variation of the demand random variables increases and as it becomes more costly to serve a demand type from a location that is not associated with the demand type. We use the coefficient of variation as a proxy for the level of uncertainty in the problem, and the problem clearly becomes more difficult as the level of uncertainty increases. Furthermore, the substitution possibilities become more scarce as it becomes more costly to serve a demand type from a location that is not associated with the demand type. Therefore, we also expect the problem to be more difficult as ρ gets larger. These observations imply that the upper bounds provided by DDR are especially attractive for the test problems that we expect to be more difficult. The total expected profits obtained by DDR, DSR and SPA are quite close to each other, but SPA provides a small consistent improvement over DDR and DSR. It is interesting that although the upper bounds obtained by DDR are significantly tighter than those obtained by DSR, the performances of DDR and DSR are quite similar. As a matter of fact, DSR appears to provide slightly better performance than DDR. For both DDR and DSR, we choose the Lagrange multipliers by using the dual solution to problem (3). Similar to our observations for the upper bounds, the gaps between the total expected profits obtained by DLP and the remaining

Problem	Tot. Exp. Prof. DLP	SPA	SPA vs. DLP
Base	155,378	166,913	6.91
$\lvert\mathcal{L}\rvert = 10$	190,681	198,391	3.89
$\lvert\mathcal{L}\rvert = 40$	162,132	172,021	5.75
$\sum_{i\in\mathcal{L}} r_{i1} = 100$	88,766	98,083	9.50
$\sum_{i\in\mathcal{L}} r_{i1} = 400$	250,448	263,775	5.05
$c = 1.6$	220,115	232,262	5.23
$c = 8.0$	113,671	121,617	6.53

Table 2: Numerical results for multi-stage problems.

three strategies get larger as the coefficient of variation of the demand random variables increases or as it becomes more costly to serve a demand type from a location that is not associated with the demand type. Therefore, using a model that addresses the uncertainty in the system becomes especially beneficial for the test problems that we expect to be more difficult.

4.2 Numerical Results for Multi-Stage Problems

In this section, we provide numerical results for multi-stage test problems. The results that we provide are taken from Topaloglu and Powell (2006). In our experimental setup, we generate one base test problem and modify its different attributes to come up with test problems with different characteristics. In the base test problem, we have 20 locations. The problem takes place over a planning horizon of 60 time periods. The cost of repositioning a resource from location i to j is given by $c \times d(i,j)$, and we use $c = 4$. For each demand type k, there is an origin location o^k and destination location δ^k, and a demand of type k can only be served by a resource at location o^k. The revenue from serving a demand of type k is $p \times d(o^k, \delta^k)$, and we use $p = 5$. We have a total of 200 resources at our disposal.

Table 2 summarizes our numerical results. The first column in this table shows the problem characteristics that we modify to obtain different test problems. The first test problem corresponds to the base case. The second and third test problems have different number of locations in the transportation network. The fourth and fifth test problems have different number of resources. Finally, the sixth and seventh test problems have different repositioning costs. The second and third columns in Table 2 show the total expected profits obtained by DLP and SPA. The fourth column shows the percent gap between the total expected profits obtained by DLP and SPA. Our numerical experiments in the previous section indicate that SPA performs at least as well as DDR and DSR. Therefore, we only compare the performance of SPA with that of DLP.

Our computational results indicate that the performance gap between DLP and SPA can be on the order of four to ten percent. One trend worth mentioning is that the performance gaps increase as we have fewer resources at our disposal. In particular, the performance gap between SPA and DLP increases from 5.05% to 6.91% as the number of resources decreases from 400 to 200. Similarly, the performance gap between SPA and DLP increases from 6.91% to 9.50% as the number of resources decreases from 200 to 100. We consistently observed this trend in all of our numerical experiments. As the resources become more scarce, we need to allocate the resources more carefully, and SPA appears to do a better job of allocating scarce resources.

5 Conclusions

In this chapter, we describe a variety of modeling and solution approaches for transportation resource management problems. The solution approaches that we propose either build on deterministic linear programming formulations, or formulate the problem as a dynamic program and use tractable approximations to the value functions. We use two classes of methods to construct value function approximations. The first class of methods relax certain constraints in the dynamic programming formulation of the problem by associating Lagrange multipliers with them. In this case, we can solve the relaxed dynamic program by concentrating on one location at a time. The second class of methods use a stochastic approximation idea along with sampled trajectories of the system to iteratively update and improve the value function approximations. Our numerical experiments indicate that the models that explicitly address the randomness in the demand arrivals can provide significant benefits.

References

Adelman, D. (2004), 'A price-directed approach to stochastic inventory routing', *Operations Research* **52**(4), 499–514.

Adelman, D. (2007), 'Price-directed control of a closed logistics queueing network', *Operations Research* **55**(6), 1022–1038.

Adelman, D. and Mersereau, A. J. (2008), 'Relaxations of weakly coupled stochastic dynamic programs', *Operations Research* **56**(3), 712–727.

Barnhart, C. and Laporte, G., eds (2006), *Handbooks in Operations Research and Management Science, Volume on Transportation*, North Holland, Amsterdam.

Benveniste, A., Metivier, M. and Priouret, P. (1991), *Adaptive Algorithms and Stochastic Approximations*, Springer.

Bertsekas, D. P. and Tsitsiklis, J. N. (1996), *Neuro-Dynamic Programming*, Athena Scientific, Belmont, MA.

Bertsimas, D. and Mersereau, A. J. (2007), 'A learning approach for interactive marketing to a customer segment', *Operations Reseach* **55**(6), 1120–1135.

Bertsimas, D. and Popescu, I. (2003), 'Revenue management in a dynamic network environment', *Transportation Science* **37**, 257–277.

Carvalho, T. A. and Powell, W. B. (2000), 'A multiplier adjustment method for dynamic resource allocation problems', *Transportation Science* **34**, 150–164.

Castanon, D. A. (1997), Approximate dynamic programming for sensor management, *in* 'Proceedings of the 36th Conference on Decision & Control', San Diego, CA.

Cheung, R. (1994), Dynamic Networks with Random Arc Capacities, with Application to the Stochastic Dynamic Vehicle Allocation Problem, PhD thesis, Princeton University, Princeton, NJ.

Cheung, R. and Chen, C. (1998), 'A two-stage stochastic network model and solution methods for the dynamic empty container allocation problem', *Transportation Science* **32**, 142–162.

Cheung, R. K. M. and Powell, W. B. (1996*a*), 'Models and algorithms for distribution problems with uncertain demands', *Transportation Science* **30**(1), 43–59.

Cheung, R. K. M. and Powell, W. B. (2000), 'SHAPE: A stochastic hybrid approximation procedure for two-stage stochastic programs', *Operations Research* **48**(1), 73–79.

Cheung, R. K. and Powell, W. B. (1996*b*), 'An algorithm for multistage dynamic networks with random arc capacities, with an application to dynamic fleet management', *Operations Research* **44**(6), 951–963.

Crainic, T., Gendreau, M. and Dejax, P. (1993), 'Dynamic and stochastic models for the allocation of empty containers', *Operations Research* **41**, 102–126.

Dantzig, G. and Fulkerson, D. (1954), 'Minimizing the number of tankers to meet a fixed schedule', *Naval Research Logistics Quarterly* **1**, 217–222.

Erdelyi, A. and Topaloglu, H. (2009*a*), 'A dynamic programming decomposition method for making overbooking decisions over an airline network', *INFORMS Journal on Computing* (to appear).

Erdelyi, A. and Topaloglu, H. (2009*b*), Using decomposition methods to solve pricing problems in network revenue management, Technical report, Cornell University, School of Operations Research and Information Engineering.
Available at `http://legacy.orie.cornell.edu/~huseyin/publications/publications.html`.

Erera, A., Karacik, B. and Savelsbergh, M. (2008), 'A dynamic driver management scheme for less-than-truckload carriers', *Computers and Operations Research* **35**, 3397–3411.

Federgruen, A. and Zipkin, P. (1984), 'A combined vehicle routing and inventory allocation problem', *Operations Research* **32**, 1019–1037.

Ferguson, A. and Dantzig, G. B. (1955), 'The problem of routing aircraft - A mathematical solution', *Aeronautical Engineering Review* **14**, 51–55.

Frantzeskakis, L. and Powell, W. B. (1990), 'A successive linear approximation procedure for stochastic dynamic vehicle allocation problems', *Transportation Science* **24**(1), 40–57.

Godfrey, G. A. and Powell, W. B. (2001), 'An adaptive, distribution-free approximation for the newsvendor problem with censored demands, with applications to inventory and distribution problems', *Management Science* **47**(8), 1101–1112.

Godfrey, G. A. and Powell, W. B. (2002*a*), 'An adaptive, dynamic programming algorithm for stochastic resource allocation problems I: Single period travel times', *Transportation Science* **36**(1), 21–39.

Godfrey, G. A. and Powell, W. B. (2002*b*), 'An adaptive, dynamic programming algorithm for stochastic resource allocation problems II: Multi-period travel times', *Transportation Science* **36**(1), 40–54.

Gorman, M. (2001), 'Intermodal pricing model creates a network pricing perspective at BNSF', *Interfaces* **31**(4), 37–49.

Hane, C., Barnhart, C., Johnson, E., Marsten, R., Nemhauser, G. and Sigismondi, G. (1995), 'The fleet assignment problem: Solving a large-scale integer program', *Math. Prog.* **70**, 211–232.

Holmberg, K., Joborn, M. and Lundgren, J. T. (1998), 'Improved empty freight car distribution', *Transportation Science* **32**, 163–173.

Kleywegt, A. J., Nori, V. S. and Savelsbergh, M. W. P. (2002), 'The stochastic inventory routing problem with direct deliveries', *Transportation Science* **36**(1), 94–118.

Kunnumkal, S. and Topaloglu, H. (2008*a*), 'A duality-based relaxation and decomposition approach for inventory distribution systems', *Naval Research Logistics Quarterly* **55**(7), 612–631.

Kunnumkal, S. and Topaloglu, H. (2008*b*), Linear programming based decomposition methods for inventory distribution systems, Technical report, Cornell University, School of Operations Research and Information Engineering.
Available at `http://legacy.orie.cornell.edu/∼huseyin/publications/publications.html`.

Kunnumkal, S. and Topaloglu, H. (2008*c*), 'A refined deterministic linear program for the network revenue management problem with customer choice behavior', *Naval Research Logistics Quarterly* **55**(6), 563–580.

Kunnumkal, S. and Topaloglu, H. (2009*a*), 'Computing time-dependent bid prices in network revenue management problems', *Transportation Science* (to appear).

Kunnumkal, S. and Topaloglu, H. (2009*b*), 'A new dynamic programming decomposition method for the network revenue management problem with customer choice behavior', *Production and Operations Management* (to appear).

Kushner, H. J. and Clark, D. S. (1978), *Stochastic Approximation Methods for Constrained and Unconstrained Systems*, Springer-Verlang, Berlin.

Kushner, H. J. and Yin, G. G. (1997), *Stochastic Approximation Algorithms and Applications*, Springer-Verlag.

Nemhauser, G. and Wolsey, L. (1988), *Integer and Combinatorial Optimization*, John Wiley & Sons, Inc., Chichester.

Powell, W. B. (1988), A comparative review of alternative algorithms for the dynamic vehicle allocation problem, *in* B. Golden and A. Assad, eds, 'Vehicle Routing: Methods and Studies', North Holland, Amsterdam, pp. 249–292.

Powell, W. B. (1996), 'A stochastic formulation of the dynamic assignment problem, with an application to truckload motor carriers', *Transportation Science* **30**(3), 195–219.

Powell, W. B. (2007), *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, John Wiley & Sons, Hoboken, NJ.

Powell, W. B. and Cheung, R. (1994*a*), 'A network recourse decomposition method for dynamic networks with random arc capacities', *Networks* **24**, 369–384.

Powell, W. B. and Cheung, R. K. (1994*b*), 'Stochastic programs over trees with random arc capacities', *Networks* **24**, 161–175.

Powell, W. B., Jaillet, P. and Odoni, A. (1995), Stochastic and dynamic networks and routing, *in* C. Monma, T. Magnanti and M. Ball, eds, '*Handbook in Operations Research and Management Science*, Volume on *Networks*', North Holland, Amsterdam, pp. 141–295.

Powell, W. B., Ruszczynski, A. and Topaloglu, H. (2004), 'Learning algorithms for separable approximations of stochastic optimization problems', *Mathematics of Operations Research* **29**(4), 814–836.

Powell, W. B., Shapiro, J. A. and Simão, H. P. (2001), A representational paradigm for dynamic resource transformation problems, *in* R. F. C. Coullard and J. H. Owens, eds, 'Annals of Operations Research', J.C. Baltzer AG, pp. 231–279.

Powell, W. B., Shapiro, J. A. and Simão, H. P. (2002), 'An adaptive dynamic programming algorithm for the heterogeneous resource allocation problem', *Transportation Science* **36**(2), 231–249.

Powell, W. B. and Topaloglu, H. (2003), Stochastic programming in transportation and logistics, *in* A. Ruszczynski and A. Shapiro, eds, '*Handbook in Operations Research and Management Science*, Volume on *Stochastic Programming*', North Holland, Amsterdam.

Powell, W. B. and Topaloglu, H. (2005), Fleet management, *in* S. Wallace and W. Ziemba, eds, '*in Applications of Stochastic Programming, Math Programming Society SIAM Series in Optimization*'.

Powell, W. B., Wu, T., Simão, H. P. and Whisman, A. (2004), 'Using low dimensional patterns in optimizing simulators: An illustration for the military airlift problem', *Mathematical and Computer Modeling* **29**, 657–675.

Puterman, M. L. (1994), *Markov Decision Processes*, John Wiley and Sons, Inc., New York.

Robbins, H. and Monro, S. (1951), 'A stochastic approximation method', *Annals of Math. Stat.* **22**, 400–407.

Ruszczynski, A. (2003), Decomposition methods, *in* A. Ruszczynski and A. Shapiro, eds, '*Handbook in Operations Research and Management Science*, Volume on *Stochastic Programming*', North Holland, Amsterdam.

Schenk, L. and Klabjan, D. (2008), 'Intra market optimization for express package carriers', *Transportation Science* **42**, 530–545.

Shapiro, J. A. (1999), A Framework for Representing and Solving Dynamic Resource Transformation Problems, PhD thesis, Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ.

Shapiro, J. A. and Powell, W. B. (2006), 'A metastrategy for large-scale resource management based on informational decomposition', *INFORMS Journal on Computing* **18**(1).

Simao, H. P., Day, J., George, A. P., Gifford, T., Nienow, J. and Powell, W. B. (2009), 'An approximate dynamic programming algorithm for large-scale fleet management: A case application', *Transportation Science* **43**(2), 178–197.

Talluri, K. T. and van Ryzin, G. J. (2005), *The Theory and Practice of Revenue Management*, Springer, New York, NY.

Topaloglu, H. (2006), 'A parallelizable dynamic fleet management model with random travel times', *European Journal of Operational Research* **175**, 782–805.

Topaloglu, H. (2009*a*), 'A tighter variant of Jensen's lower bound for stochastic programs and separable approximations to recourse functions', *European Journal of Operational Research* **199**, 315–322.

Topaloglu, H. (2009*b*), 'Using Lagrangian relaxation to compute capacity-dependent bid-prices in network revenue management', *Operations Research* **57**(3), 637–649.

Topaloglu, H. and Kunnumkal, S. (2006), 'Approximate dynamic programming methods for an inventory allocation problem under uncertainty', *Naval Research Logistics Quarterly* **53**, 822–841.

Topaloglu, H. and Powell, W. B. (2003), 'An algorithm for approximating piecewise linear functions from sample gradients', *Operations Research Letters* **31**, 66–76.

Topaloglu, H. and Powell, W. B. (2006), 'Dynamic programming approximations for stochastic, time-staged integer multicommodity flow problems', *INFORMS Journal on Computing* **18**(1), 31–42.

White, W. (1972), 'Dynamic transshipment networks: An algorithm and its application to the distribution of empty containers', *Networks* **2**(3), 211–236.

White, W. and Bomberault, A. (1969), 'A network algorithm for empty freight car allocation', *IBM Systems Journal* **8**(2), 147–171.

Yost, K. A. and Washburn, A. R. (2000), 'The LP/POMDP marriage: Optimization with imperfect information', *Naval Research Logistics Quarterly* **47**, 607–619.

Zhang, D. and Adelman, D. (2009), 'An approximate dynamic programming approach to network revenue management with customer choice', *Transportation Science* **42**(3), 381–394.