A Successive Linear Approximation Procedure for Stochastic, Dynamic Vehicle Allocation Problems

LINOS F. FRANTZESKAKIS and WARREN B. POWELL

Department of Civil Engineering and Operations Research, Princeton University, Princeton, New Jersey 08544

The Stochastic Dynamic Vehicle Allocation problem involves managing a fleet of vehicles over time in an uncertain demand environment to maximize expected total profits. The problem is formulated as a Stochastic Programming problem. A new heuristic algorithm is developed and is contrasted to various deterministic approximations. The paper presents computational results that were obtained by employing a Rolling Horizon Procedure to simulate the operation of the truckload carrier. Results indicate the superiority of the new algorithm over other approaches tested.

INTRODUCTION

The vehicle allocation problem arises when a common carrier must manage a fleet of vehicles to maximize profits over a planning horizon. Demand materializes as shippers call the carrier requesting a vehicle to be available in a specific location on a specific day to carry something to a given destination. We refer to the above request as a *load*, or *loaded movement*; note that the carrier may decline a load if it is considered unprofitable or the vehicle supply available cannot accommodate it. Each load may be served by a single vehicle and the vehicle is dedicated to that load; no vehicle can be shared by loads. Truckload motor carriers, container companies, rental car agencies and (in some aspects) railroads fall into this category. Time is discretized into intervals, typically one day, and on each day, the operator must either assign each vehicle to a requested loaded movement or move it empty to another region to pick up a requested or expected loaded movement, or to hold it in the same region until the next day. The dynamic element of the problem is obviously very important: decisions made on one day influence directly the vehicle supplies of regions in the future and thus the decisions to be made at that time. The use of the term "Dynamic Vehicle Allocation Problem" (DVA) underlines the importance of future consequences of decisions made at any time concerning the management of the vehicle fleet.

The major difficulty of the DVA problem lies in the fact that the demands are typically uncertain and the level of uncertainty increases further into the future. For example, in the truckload market (which is the application of the DVA that motivates our research) it is unusual for shippers to book a load more than 2 or 3 days in advance and typically the carrier at the beginning of the day knows only about 40% of the loads that will be carried or requested on that day, and about 10% or less of the loads to be carried on the next day. Consequently, the motor carrier must be able to estimate future demands and to make decisions that anticipate their impacts on future periods. Travel times between regions may also be treated as random, thus causing the supply of vehicles in a region at a time period to be stochastic, even after the dispatch decisions have been made.

A variety of formulations of the DVA problem have been proposed to handle both uncertainty and the infinite horizon nature of the problem. DEJAX and CRAINIC,[13] as well as BOOKBINDER and SETHI,[8] review models of empty vehicle repositioning and dynamic transportation problems. POWELL[23] specifically reviews alternative formulations of the DVA and contrasts methods for handling uncertainty. Earlier papers include MISRA[20] and BAKER,[1] which deal with the static deterministic case for repositioning empty freight cars, COOPER and LEBLANC[9] which handles the static stochastic case and WHITE and BOMBERAULT[40] which formulates the deterministic dynamic case as a linear transshipment network. The models developed to deal with both the dynamic and the stochastic nature of the problem are usually application-specific, since certain modeling simplifications

Transportation Science
Vol. 24, No. 1, February 1990

0041-1655/90/2401-0040 $01.25

appropriate for each application are necessary to reduce the complexity of the problem.

Papers by JORDAN[16] and JORDAN and TURNQUIST[17] show the first effort to incorporate the uncertainty of demands and vehicle supplies into the model and result in a network optimization model for the allocation of empty freight cars in the railroad carrier's case. The model is a nonlinear optimization problem with linear decomposable constraints, where the objective is to maximize total expected profits. Since its concavity was not proven by the authors, the model is considered a heuristic. After the model is solved for a specific day, the first period decisions are implemented. As time passes, the additional information that becomes available is used in the construction of the model for the next day. An important feature of the model is that the solution of the model for the next day may revise decisions made on previous days, by rerouting empty cars in transit. The model is solved using a FRANK–WOLFE algorithm[14] and it is computationally very efficient.

The truckload carrier case, which is in the focus of our interest, is characterized by a high level of competition among the carriers and a high degree of uncertainty of the future demands. The models developed for the stochastic DVA are primarily heuristic and face the difficulty in the choice of the decision variables to describe the shippers' recommended actions. The choice usually leads to unreasonable assumptions in the shippers' behavior that limit the applicability of those models. In POWELL et al.[21] the decision variables are the fractions of the vehicle supply of a region at a time period to be dispatched empty or loaded to another region. The model, thus, requires a certain fraction of the vehicle supply of a region at a time period to be sent to a destination, where this fraction is statistically independent of the demand for loads for that destination. Also, at the same time that empty vehicles are sent to one destination because of unavailability of loads, the vehicles assigned to another destination may be insufficient to handle the actual demand. An alternative model by POWELL[22] requires that if vehicles assigned to a destination are not used due to insufficient demand, they are not sent empty to that destination (as the previous model did), but instead they are held in the same region. Again, no substitutability of vehicles for destinations facing excess demand is allowed. Another theoretical weak point of these models is that no consideration is given to the truncation of the planning horizon, which is to say that the end effects are ignored. These models result in constrained nonlinear optimization formulations, where concavity is guaranteed only in [21], and are solved by the application of the Frank–Wolfe algorithm. The optimal solutions are not integer, however, which limits their use as real-time vehicle routing tools.

The primary goal of this research is to introduce an alternative heuristic approach as well as to review and contrast alternative modeling and solution approaches. Particular attention is given to the handling of forecasting uncertainties. The algorithms tested cover deterministic models and stochastic dynamic models. Their performance is tested on a particular realistic dataset. The presentation is organized as follows. In Section 1 we formulate the problem mathematically and Section 2 presents the algorithm. Section 3 describes briefly representatives of static and dynamic deterministic formulations. Section 4 provides the details of the rolling horizon procedure used to simulate the operation of the carrier over time and provide results on the models' performance. Section 5 presents the results of a series of numerical experiments contrasting the new procedure to deterministic static and dynamic formulations. Finally, Section 6 concludes the paper.

1. PROBLEM FORMULATION

THIS SECTION presents a stochastic programming formulation of the DVA problem with a planning horizon of N periods. Let us restate the problem: at the present time period and with the available vehicle allocation, the carrier must decide which loads to accept or refuse and how many vehicles to relocate or hold over (in order to obtain a more favorable future vehicle allocation) to maximize the total expected profits over a planning horizon of N periods in an environment of independent random future demands with a known distribution. The carrier's fleet has a certain number of vehicles and each vehicle can accommodate any loaded movement; there is no compatibility problem between vehicles and loads.

In order to gain insight into a problem which is easily obscured by a variety of tedious details, we make certain simplifying modeling assumptions which were applied to all the models to be presented. The continuum of the United States is divided into 60 discrete regions and let $\mathbf{R} = \{1, \ldots, R\}$ denote the set of regions, while R refers, according to the context, to either the Rth region or the cardinality of the set $\mathbf{R}$.

Listed below are the assumptions that were made in the development of the models:

(A1) The travel times between all regions are equal to an integer number of time periods and each time period is one day.

(A2) The carrier at the beginning of the day knows all the loads that will be called in to be picked up that day; in addition, it does not know any of the loads that will be called in to be picked up on the following days. That allows the decision maker (the carrier, in our case) to solve the model once per day.

(A3) Loads that are not picked up on the first day are lost. Thus, the performance of the models can be evaluated on a day by day basis and the necessary bookkeeping is reduced.

Some of these assumptions are not particularly realistic, but the resulting problem is still quite complex. Furthermore, the analysis still provides insights into the solution of the original problem. It should be emphasized here that assumptions (A1) and (A2) are not necessary modeling restrictions and can be easily relaxed; they are made only to simplify the presentation of the problem and provide a clearer picture of the optimal dispatch strategies of the carrier. In addition, two more simplifications were made to ease the mathematical presentation:

(A4) All vehicles are available for the first time at time period $t = 1$, and

(A5) Travel times between all regions are uniformly equal to one time period, either for a loaded or empty move.

It should be emphasized that these assumptions are not actually used in the construction of the models whose performance evaluation is presented in Section 5.

We begin the presentation of the formulation of the DVA problem by giving the necessary notation. In this paper we will use the convention that the first time period is $t = 1$, and, when a time period is mentioned, it is implied that we refer to the beginning of that time period. Revenues from carrying a load are treated as negative costs, in order to formulate the problem as a cost minimization problem (as is more common in this area of research). Additionally, since each vehicle carries one load, we will use the terms "flow of vehicles" and "number of vehicles" without distinction. Let us denote:

$x_{ij}(t)$ = number of trucks moving loaded from region i to region j, departing from i in the beginning of period t, $t = 1, \ldots, N$,
$= 0$, for $t \leq 0$,

$y_{ij}(t)$ = number of trucks moving empty from region i to region j, departing from i in the beginning of period t, $t = 1, \ldots, N$,
$= 0$, for $t \leq 0$,

$\Phi_{ij}(t)$ = random variable denoting the number of loads that will be called from i to j to be picked up at time t, $t = 2, \ldots, N$,

r_{ij} = average contribution (revenue minus direct operating cost) for pulling a load from i to j.

c_{ij} = cost of moving empty from i to j,

$L_{ij}(1)$ = actual number of loads known at time $t = 1$ to be available moving from i to j at the first time period,

$T_i(t)$ = number of trucks becoming available for the first time in region i at time t,
$= 0$, for $t = 2, \ldots, N$ (by our simplification A4),

$S_i(t)$ = supply of trucks at i on day t, $t = 1, \ldots, n$
$= \sum_{k \in R} [x_{ki}(t-1) + y_{ki}(t-1)] + T_i(t)$, (by our simplification A5).

In the model formulations that are presented in the remainder of the paper, we follow the notation that:

- —bold letters define a matrix or a vector with known or random elements, for example, $\mathbf{r} = \{r_{ij}, \forall i = 1, \ldots, R, \forall j = 1, \ldots, R\}$,
- —Greek capital letters are reserved for random variables,
- —Greek small letters denote a realization of the random variable,
- —a bar on a random variable denotes its expected value,
- —vectors are implied to be column vectors,
- —$\mathbf{A} * \mathbf{B}$ denotes matrix or vector multiplication,
- —$\mathbf{1}$ represents a column vector of size R with elements equal to 1,
- —$\mathbf{A}^{\mathrm{T}}$ denotes the transpose of vector or matrix $\mathbf{A}$,
- —inequality constraints in vector or matrix form imply inequalities in each individual element of the vectors or matrices.

Let us now define the following optimization problem, for each time period $t = 2, \ldots, N-1$:

$$\psi^t(\mathbf{S}(t), \boldsymbol{\phi}(t)) = \min_{\{\mathbf{x}(t), \mathbf{y}(t)\}} [-\mathbf{r}^{\mathrm{T}}\mathbf{x}(t) + \mathbf{c}^{\mathrm{T}}\mathbf{y}(t) + \mathrm{E}_{\boldsymbol{\Phi}(t+1)}[\Psi^{t+1}(\mathbf{S}(t+1), \boldsymbol{\Phi}(t+1))]] \tag{1}$$

subject to:
$$\mathbf{x}(t) \leq \boldsymbol{\phi}(t) \tag{1a}$$
$$[\mathbf{x}(t) + \mathbf{y}(t)] * (\mathbf{1}) = \mathbf{S}(t) \tag{1b}$$
$$[\mathbf{x}(t) + \mathbf{y}(t)]^{\mathrm{T}} * (\mathbf{1}) = \mathbf{S}(t+1) \tag{1c}$$

Note that we will use $\psi^t(\mathbf{S}(t), \boldsymbol{\phi}(t))$ to denote the realization of the random variable $\Psi^t(\mathbf{S}(t), \boldsymbol{\Phi}(t))$ when the link capacities $\boldsymbol{\Phi}(t) = \boldsymbol{\phi}(t)$ for period t. This choice of notation, instead of the more conventional $\psi^t(\mathbf{S}(t), \boldsymbol{\Phi}(t) \mid \boldsymbol{\Phi}(t) = \boldsymbol{\phi}(t))$, was made to ease the presentation.

Let us additionally define for the last time period:

$$\psi^N(\mathbf{S}(N), \boldsymbol{\phi}(N)) = \min_{\{x(N), y(N)\}}[-\mathbf{r}^T\mathbf{x}(N) + \mathbf{c}^T\mathbf{y}(N)] \quad (2)$$

$$\text{subject to:} \quad \mathbf{x}(N) \leqslant \boldsymbol{\phi}(N) \quad (2a)$$

$$[\mathbf{x}(N) + \mathbf{y}(N)] * (1) = \mathbf{S}(N). \quad (2b)$$

The problem of maximizing total profits over the N time periods horizon can now be formulated as the following N-stage Stochastic Program:

$$\psi^1(\mathbf{S}(1)) = \min_{\{x(1), y(1)\}}[-\mathbf{r}^T\mathbf{x}(1) + \mathbf{c}^T\mathbf{y}(1) + \mathrm{E}_{\mathbf{\Phi}(2)}[\Psi^2(\mathbf{S}(2), \mathbf{\Phi}(2))]] \quad (3)$$

$$\text{subject to:} \quad \mathbf{x}(1) \leqslant \mathbf{L}(1) \quad (3a)$$

$$[\mathbf{x}(1) + \mathbf{y}(1)] * (1) = \mathbf{S}(1) \quad (3b)$$

$$[\mathbf{x}(1) + \mathbf{y}(1)]^T * (1) = \mathbf{S}(2). \quad (3c)$$

Admittedly, in the above formulation we used the small Greek symbol $\psi^1(\mathbf{S}(1))$ for the deterministic quantity of the total expected profits over the N-period planning horizon, given the initial vehicle allocation $\mathbf{S}(1)$. This abuse of notation was made only to expose the similarity of the problems for different time periods.

In these formulations, the first set of constraints ensures that no loaded movements take place in excess of the demand for such movements. The second set of constraints implies that the total number of vehicles dispatched out of region i at time period t does not exceed the vehicle supply available in i at t. Since the option of holding a vehicle in region i at time t until the next time period is available in the formulation (and represented by $y_{ii}(t)$), these constraints are formulated as equality rather than inequality constraints. The third set of constraints in formulations (1) and (3) are flow conservation constraints.

Problems with uncertain objective or constraint (as in our case) coefficients fall in the realm of Stochastic Programming. The complexity of those problems necessitates the truncation of their planning horizon to a fixed number N of time periods, although this may introduce deviations from the optimal solution of the infinite horizon problem. When such a problem considers the uncertainty of the future up to N time periods, it is classified as an N-stage Stochastic Linear Program with recourse.

The most general case is the one of fixed recourse, where uncertainty is restricted to the right-hand side of the constraints and the constraint set has the general form $Ax = b$. Network recourse arises when the problem, conditioned on a specific realization of the random variables, is a pure network; that is, A is an arc-node incidence matrix. The simplest case of recourse is the case of simple recourse, where the problem, conditioned on a specific realization of the random variables, involves just the incurrence of a penalty whenever the first time period decisions anticipated different values of the random variables than their actual realizations. When the realization is higher than the estimated value of the random variable, a cost is incurred per unit of discrepancy; this cost is called "underage cost." In the reverse case, when the realization was overestimated, an "overage cost" per unit of discrepancy is incurred.

Most of the research in Stochastic Programming, since it was introduced by BEALE[2] and DANTZIG,[10] has focused on 2-stage linear programs with random variables described by discrete distributions and for the case of fixed recourse with the uncertainty restricted to the right-hand side vector of the constraints. DANTZIG and MANDANSKY[12] solve the dual of this problem by using DANTZIG–WOLFE decomposition,[11] while KALL[18] and STRAZICKY[26] present another dual method based on basis factorization techniques. Important work on the problem of fixed recourse is presented in WETS[35–38] and WALKUP and WETS.[28,29] The L-shaped Decomposition, introduced by VAN SLYKE and WETS,[27] is the first method that solves the primal problem and is based on the use of cutting planes (outer linearization or Bender's decomposition). WOLLMER,[41] finally, deals with the problem of fixed recourse when the first stage variables are 0–1 integer variables, by using Bender's Decomposition and an implicit enumeration scheme.

Methods for handling simple recourse include algorithms by WETS[39] and ZIEMBA,[42] who also introduces the convex simplex method. For an application of Wets' algorithms for the simple recourse problem, see WALLACE and BREKKE.[34] For the particular case of a Transportation problem with uncertain demands, which is a special case of the simple recourse problem, the Forest Iteration Method has been introduced by QI.[24]

The problem of network recourse has been extensively examined by WALLACE,[30] who first proved that in a transportation problem the constraint set decomposes into polyhedral cones. WALLACE[31] also introduces a method for solving these problems by using Shurr complements modified for the network case and applied these basis factorization and partitioning techniques in [32] and [33].

Multistage Stochastic Linear Problems have been treated primarily for the case of linear objective function and linear constraints and for the case where the randomness is restricted to the right-hand sides of the constraints. BEALE, FOREST and TAYLOR[3] deals with the case where the uncertainty is restricted to the total inventory constraints of the type (1b). BEALE, DANTZIG and WATSON[4] solves it by assuming a suitable functional for the *value function* ψ^t (such as a quadratic function) and fit the problem into the dynamic programming framework. BIRGE[5, 6] uses decomposition and partitioning techniques to solve the problem. LOUVEAUX[19] for the case of a stochastic integer program uses the concept of block separability (which may be inherent in the structure of some problems) and develops efficient solution techniques. Even if such a separability is not present in the structure of a problem, procedures have been developed to force it onto the problem. These procedures are basically bounding procedures and are based on the ray approximation procedure in BIRGE and WETS.[7] It uses the sublinearity property of the recourse function to obtain a separable function that dominates the recourse function. A different solution approach to the Multistage problem is the Scenario Analysis approach, where the problem is solved for a limited number of realizations (scenarios) and conclusions are drawn from their optimal solutions. Such an approach is appropriate when no probability distributions can be assumed for the random variables of the problem. ROCKAFELLAR and WETS[25] introduced the Scenario Aggregation Algorithm, which is a rigorous procedure for combining the optimal solutions for the alternative scenarios to obtain a general decision policy. The deterministic equivalent program to the multi-stage network problem increases with the planning horizon, and the cardinality of the discretization of the random variables. Thus, approximation techniques have been introduced by Wets[38] and Birge and Wets.[7]

In the context of the Dynamic Vehicle Allocation Problem, the uncertainty involved is due to the random demands on the links rather than on the nodes of a network. Thus the random coefficients are the right-hand sides in the constraints that imply that the flow on a link cannot exceed the demand for that link. Another characteristic of the problem is that the 2-stage problem reduces to a pure network (see Powell[23]). Additionally, the flow of vehicles between a pair of regions does not solely depend on the demand between those regions and the vehicle availability in the origin region, but also, through the network conservation constraints, it is jointly dependent on the realizations of the demand between other pairs of regions. Thus, it is a Stochastic Linear problem with the network recourse. In the effort to solve such a complicated problem, many attempts have been made to relax the network recourse condition and treat the problem in a different level of difficulty; some of them are described next.

The simple recourse strategy corresponds to the case where, when the demand between a pair of regions falls short of the fraction of the vehicles assigned to it, the excess vehicles move empty anyway. Thus, the overage cost of this problem equals the cost of making that empty move, while the underage cost is the lost revenue of that move. The model introduced by Powell et al.,[21] actually deals with that problem.

Null recourse, a term introduced by Powell,[23] implies that neither a corrective action is taken when the values of the random variables are known, nor any penalties are incurred for not anticipating their realized value. In the context of the DVA problem, this corresponds to the case where, when the demand between a pair of regions falls short of the fraction of vehicles assigned to it, the excess vehicles are held in the origin region (incurring no cost) rather than dispatched empty. The model introduced by Powell[22] actually corresponds to this case.

The strong dynamic nature of the DVA problem can only be accommodated by a multistage Stochastic Program. The decomposition and partitioning methods of Birge[6] and Louveaux,[19] though, are developed for the case of fixed recourse and if applied to the case of network recourse would eliminate the desirable network structure of the constraint set. Thus, our heuristic approach attempts to maintain the network structure, allowing the use of efficient network optimization techniques.

In the next section we develop a heuristic algorithm for the DVA problem, after which Section 3 discusses alternative formulations of the DVA problem (or, equivalently, relaxations of the N-stage Stochastic Linear Program with network recourse).

2. THE SUCCESSIVE LINEAR APPROXIMATION PROCEDURE

IN THIS section we develop a new heuristic algorithm, which we will refer to as Successive Linear Approximations Procedure (SLAP). First, we present the intuitive motivation behind the basic idea of the algorithm. Then, in Section 2.2 we introduce the recursion which is central to the procedure, while Section 2.3 discusses in detail the linear approximations needed.

2.1. Motivation

It is common sense that the total expected profits in a region j at a time period depends on the number of available vehicles there. Let us assume that the available vehicles are given an arbitrary index number and we decide that the vehicle with index number 1 will take the best available load, the 2nd will take the second best, and so forth. Under this arbitrary dispatch logic, it is obvious that the expected profits of the kth vehicle will be higher than those of the $(k + 1)$th one, since they both compete for the same loads and the first one has dispatch priority over the second. Thus, the total expected profits in that region does not grow linearly with the number of available vehicles; the function is concave. To get into the cost minimizing context that Stochastic Programming requires, we have that (changing the signs) the total expected future costs in a region at a time period (which is the recourse function) is a convex function of the vehicle supply, as proven later (see Proposition 1).

Let us consider now the last period of an N-stage Stochastic Programming problem. For every region j, the future after the beginning of period N is summarized by a function that has the convexity property discussed above. Let us now try to devise a dispatch policy for region i at time $N-1$, that is a rule on where a vehicle should be dispatched when the demands for period $N-1$ are known. Let us assume that this dispatch policy will be used for all vehicles available in that region at that time. If a vehicle is dispatched empty or loaded to region j, the expected cost of that decision is equal to c_{ij} or $-r_{ij}$, respectively, plus the expected cost incurred by the vehicle once in a region j at time N. That second component is actually the marginal cost of having one more vehicle in region j at time N. Although the recourse function for that region at time N is known, that marginal cost (which is the slope of that function) cannot be determined unless the "existing" vehicle supply is known (i.e., the point where the slope should be taken). In Figure 1a, for example, that marginal cost is known only if we know how many vehicles the other regions sent to region j. But this supply depends on the dispatch of all the other regions at time $N-1$, making the problem very complicated. This is essentially what the concept of network recourse implies: dispatch decisions in different regions cannot be made individually, because they are interrelated through the future vehicle supplies that they induce.

Suppose now that the convex function describing the total expected cost in region j at time N is replaced by a linear one (which implies a constant marginal cost of a vehicle), as shown in Figure 1b. Then the expected cost of a dispatch decision in region i at time $N-1$ does not depend on the vehicle supply of region j at time N and thus does not depend on the dispatch decisions of the other regions at time $N-1$ either. So the problem in period $N-1$ is decomposed by region.

The choice of the linear approximation to be used in the place of the recourse function of period j at time N is crucial. Consider, for example, a model that would substitute the function by the slope of the function at point 0, θ^0. Let us examine a simple example with two regions and with known vehicle supplies, as depicted in Figure 2a. Suppose (for the sake of the argument) that the cost to move empty or loaded from one region to another is c and $-r$ respectively. Let us also suppose that the demands between any pair of regions happens to be uniformly equal to zero for period $N-1$ so that all vehicles are moving empty or are held over. Note that the expected recourse functions depicted level off when the supply exceeds a certain limit. This implies that excess vehicles do not cost anything; that is the cost of holding a vehicle until the next dispatch instant is zero. Let us assume that $j=1$ is the region that has the biggest slope, θ_1^0, of the recourse function at point 0 and F be the fleet size. Then, it is obvious that, if $c+\theta_1^0<\theta_2^0$, region 1 will accumulate at time N the entire fleet size

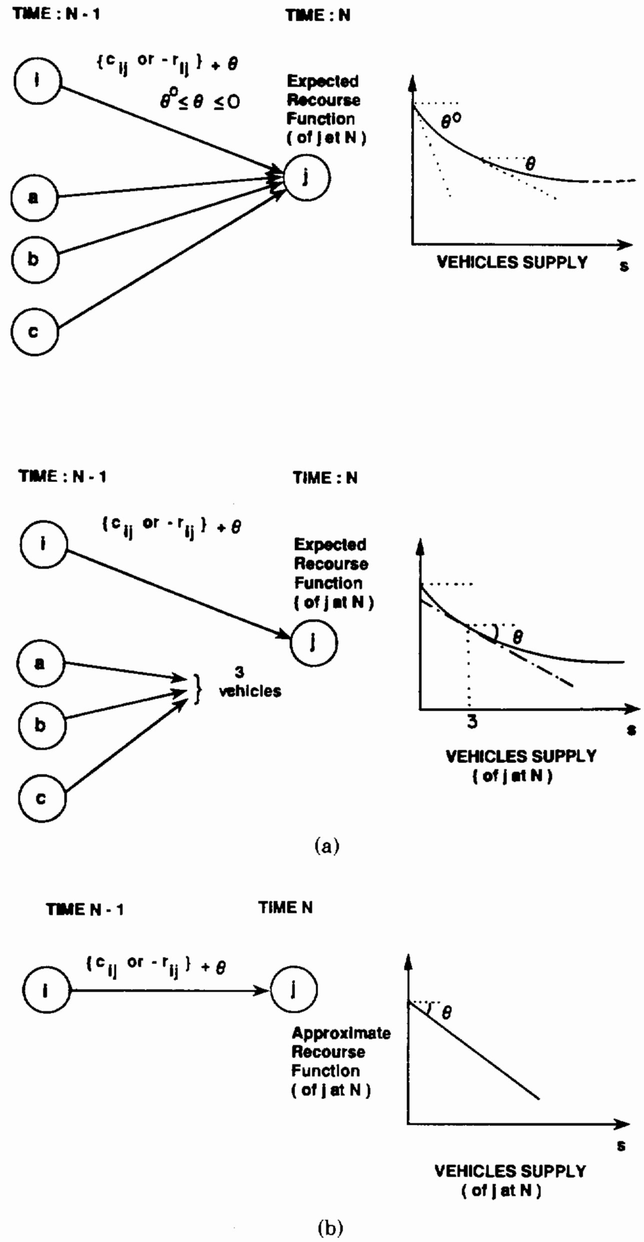

Fig. 1. Decomposition of the Network Recourse when a linear approximation is used.

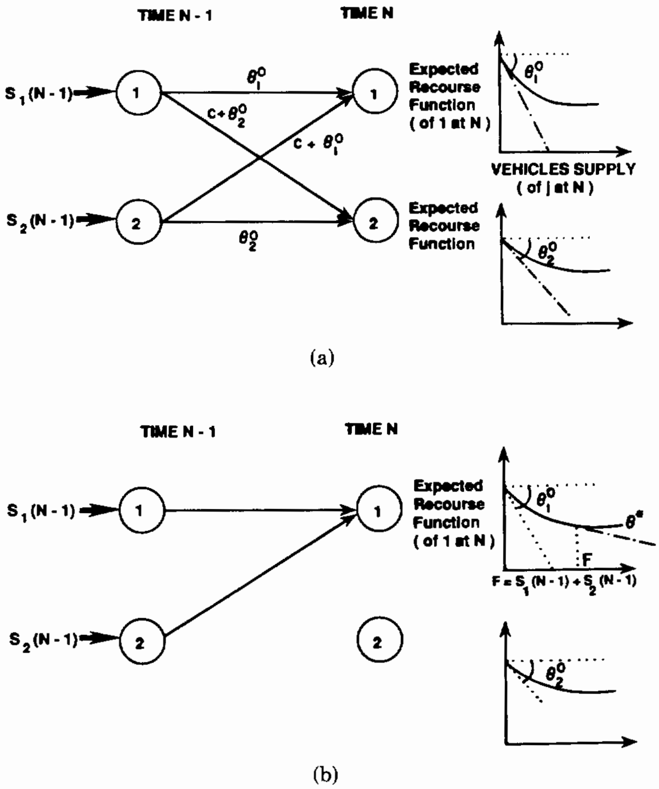

Fig. 2. Problems associated with the choice of the linear approximation. (a) Resulting network of proposed model. (b) Optimal flows of the above network.

F, as shown in Figure 2b. This makes the marginal cost of a vehicle in 1 at time N equal to the slope θ^*, which may be much higher than our initial estimate θ_1^0, resulting naturally in a very poor vehicle allocation time N.

The intention of the Successive Linear Approximations procedure is the substitution of the expectation of the recourse of a region for a period t by a linear approximation in a way that avoids the above mentioned complication. Section 2.2 describes the above line of thought for the time period t, shows the impact of that substitution when the optimization problem of period $t-1$ is considered and how the multistage problem decomposes in that sense in solving successive 2-stage problems. Thus, the SLAP procedure is essentially a backward recursion combined with a linear approximation of a convex function at each step of the recursion. The last Section 2.3 discusses the issues related to the choice of the linear approximation.

2.2. A Step in the SLAP Algorithm

In this section, the step of the SLAP algorithm involving the $(t-1)$th time period is described. We assume that from the step involving period t we have for every region i the expected recourse function $\bar{\Psi}_i^t(\mathbf{S}(t), \mathbf{\Phi}(t))$.

The convexity of $\bar{\Psi}_i^t(S_i(t) = s, \mathbf{\Phi}(t))$ with respect to s (which is proven in Section 2.3) implies that we can form a tangential linear approximation of $\bar{\Psi}_i^t(S_i(t) = s, \mathbf{\Phi}(t))$ as follows:

$$\bar{\Psi}_i^t(S_i(t) = s, \mathbf{\Phi}(t)) \triangleq \hat{\theta}_i(t)s + b_i(t) \quad \forall i = 1, \ldots, R \tag{4}$$

where $\hat{\theta}_i(t)$ is an approximate "average" or marginal expected cost of a vehicle in region i at time t, and $b_i(t)$ is a constant term depending on the region i and time t. Note that the constant $b_i(t)$ term does not need to be determined from the optimization point of view, as will become apparent later. More details on how this linear approximation is fitted are given in Section 2.3.

Consider next the objective function of the optimization problem (1) for period $t-1$:

$$\psi^{t-1}(\mathbf{S}(t-1), \boldsymbol{\phi}(t-1)) \tag{5}$$

$$= \min_{\{x(t-1), y(t-1)\}} [-\mathbf{r}^{\mathrm{T}}\mathbf{x}(t-1) + \mathbf{c}^{\mathrm{T}}\mathbf{y}(t-1) + \bar{\Psi}^t(\mathbf{S}(t), \mathbf{\Phi}(t))] \tag{5a}$$

$$= \min_{\{x(t-1), y(t-1)\}} \left[-\mathbf{r}^{\mathrm{T}}\mathbf{x}(t-1) + \mathbf{c}^{\mathrm{T}}\mathbf{y}(t-1) + \sum_{j\in R} \bar{\Psi}_j^t(\mathbf{S}_j(t) = s, \mathbf{\Phi}(t))\right] \tag{5b}$$

$$\triangleq \min_{\{x(t-1), y(t-1)\}} \left[-\mathbf{r}^{\mathrm{T}}\mathbf{x}(t-1) + \mathbf{c}^{\mathrm{T}}\mathbf{y}(t-1) + \sum_{j\in R} [\hat{\theta}_j(t)\mathbf{S}_j(t) + b_j(t)]\right] \tag{5c}$$

$$= \min_{\{x(t-1), y(t-1)\}} \sum_{i\in R}\sum_{j\in R} [-r_{ij}x_{ij}(t-1) + c_{ij}y_{ij}(t-1)] + \sum_{j\in R}\left[\hat{\theta}_j(t) \sum_{i\in R} (x_{ij}(t-1) + y_{ij}(t-1)) + b_j(t)\right]. \tag{5d}$$

From (5a) to (5b) we used the decomposition of the function $\bar{\Psi}^t$ by region and then in (5c) substituted the region contribution $\bar{\Psi}_i^t$ with its linear approximation. From (5c) to (5d) we used the flow conservation constraints (1c) for $t = t-1$.

If we ignore the constant terms of the linear approximations from the perspective of finding the optimal solution, we obtain that:

$$\psi^{t-1}(\mathbf{S}(t-1), \boldsymbol{\phi}(t-1)) \triangleq \min_{\{x(t-1), y(t-1)\}} \sum_{i\in R}\sum_{j\in R} ([-r_{ij} + \hat{\theta}_j(t))x_{ij}(t-1) + (c_{ij} + \hat{\theta}_j(t))y_{ij}(t-1)] \tag{6}$$

subject to:

$$\mathbf{x}(t-1) \leq \boldsymbol{\phi}(t-1) \tag{6a}$$

$$[\mathbf{x}(t-1) + \mathbf{y}(t-1)]^{\mathrm{T}} * (1) = \mathbf{S}(t). \tag{6b}$$

This problem clearly decomposes by region i, since the objective function is separable and there are no binding constraints.

In order to understand better the structure of the optimal solution of the above problem, let us view it in the context of the individual vehicles rather than the flows between regions. Since a vehicle in region i at time $t-1$ can move loaded or empty to any other region or can be held in i until time t, there are $2R$ possible dispatch decisions for it, which we will call *options*. The $2R$ dispatch options for a vehicle in region i at time $t-1$ are ranked in a decreasing order according to their direct contribution to total costs plus the average (negative) value of a truck at the destination of the option at time t. That is, $-r_{ij} + \hat{\theta}_j(t)$ for loaded movements to region j, and $c_{ij} + \hat{\theta}_j(t)$ for empty movements to that region. It is thus possible for an empty move to be ranked higher than a loaded move. Then define for period $t-1$:

g_{in}^{t-1} = the contribution of the nth ranked possible movement, for $n = 1, \ldots, 2R$
$= -r_{ij} + \hat{\theta}_j(t)$ if the nth ranked option is to move loaded to region j,
$= c_{ij} + \hat{\theta}_j(t)$ if the nth ranked option is to move empty to region j.

All vehicles in region i at time $t-1$ have the same dispatch options, but the destination of each of them depends on the realizations, $\phi_{ij}(t-1)$, of the random demands, $\mathbf{\Phi}_{ij}(t-1)$, out of region i at time $t-1$. Let us assume that the vehicles in region i at time $t-1$ have an index number k, $k = 1, \ldots, S_i(t-1)$ and that vehicles with higher index number have higher priority in the dispatch process. Thus, for example, the kth vehicle will take, under any demand realizations $\phi_{ij}(t-1)$, an option at least as profitable as the option that the $(k+1)$th vehicle. So, let us introduce the random variable:

$\Delta_{ik}^{t-1}(\mathbf{\Phi}) = n$, if the kth vehicle in region i at time $t-1$ is dispatched on the nth option, $n = 1, \ldots, 2R$, when the vector of demands is $\mathbf{\Phi}$.

d_{ikn}^{t-1} = the probability that the kth vehicle in region i at time $t-1$ will be dispatched on the nth ranked possible movement.
$= \mathrm{Prob}[\Delta_{ik}^{t-1}(\mathbf{\Phi}) = n]$.

Let us additionally define:

$\mathbf{\Theta}_{ik}^{t-1}(\mathbf{\Phi})$ = random variable denoting the contribution of the kth vehicle in region i at time $t-1$, when the vector of demands is $\mathbf{\Phi}(t-1)$.

The solution to the optimization problem (6) is obvious. Each vehicle k in region i at time $t-1$ is dispatched to the highest ranked available dispatch option n. Therefore, at optimality and for a given realization of the random demands $\boldsymbol{\phi}(t-1)$, the actual contribution of any vehicle will be equal to the contribution of the highest ranked option n that this vehicle can take given $\boldsymbol{\phi}(t-1)$. Thus:

$$\theta_{ik}^{t-1}(\boldsymbol{\phi}) = g_{in}^{t-1}, \tag{7}$$

where n is the index number of the highest ranked available option.

The objective value of the optimization problem for region i at time $t-1$ given the vector of demands $\boldsymbol{\phi}(t-1)$ is then:

$$\psi_i^{t-1}(S_i(t-1), \boldsymbol{\phi}(t-1)) = \textstyle\sum_{k=1}^{S_i(t-1)} \theta_{ik}^{t-1}(\boldsymbol{\phi}) \tag{8}$$

Finally, combining the optimal solutions of the subproblems, we have for problem (2) as a whole:

$$\begin{aligned}\psi^{t-1}&(\mathbf{S}(t-1), \boldsymbol{\phi}(t-1)) \\ &= \textstyle\sum_{i\in R} \psi_i^{t-1}(S_i(t-1), \boldsymbol{\phi}(t-1)) \\ &= \textstyle\sum_{i\in R} [\sum_{k=1}^{S_i(t-1)} \theta_{ik}^{t-1}(\boldsymbol{\phi})]\end{aligned} \tag{9}$$

Introducing the optimization problem in the context of random demands, we have:

$$\begin{aligned}\bar{\Psi}^{t-1}&(\mathbf{S}(t-1), \mathbf{\Phi}(t-1)) \\ &= \mathrm{E}_{\mathbf{\Phi}(t-1)}[\Psi^{t-1}(\mathbf{S}(t-1), \mathbf{\Phi}(t-1))] \\ &= \textstyle\sum_{i\in R} \bar{\Psi}_i^{t-1}(S_i(t-1), \mathbf{\Phi}(t-1))\end{aligned} \tag{10}$$

where

$$\begin{aligned}\bar{\Psi}_i^t&(S_i(t-1), \mathbf{\Phi}(t-1)) \\ &= \mathrm{E}_{\mathbf{\Phi}(t-1)}[\Psi_i^{t-1}(S_i(t-1), \mathbf{\Phi}(t-1))] \\ &= \mathrm{E}_{\mathbf{\Phi}(t-1)}[\textstyle\sum_{k=1}^{S_i(t-1)} \theta_{ik}^{t-1}(\mathbf{\Phi})] = \sum_{k=1}^{S_i(t-1)} \bar{\mathbf{\Theta}}_{ik}^{t-1}\end{aligned} \tag{11}$$

where, by definition:

$$\begin{aligned}\bar{\mathbf{\Theta}}_{ik}^{t-1} &= \mathrm{E}_{\mathbf{\Phi}(t-1)}[\theta_{ik}^{t-1}(\mathbf{\Phi})] \\ &= \textstyle\sum_{n=1}^{2R} g_{in}^{t-1} d_{ikn}^{t-1} \quad \forall i = 1, \ldots, R, k\end{aligned} \tag{12}$$

The dispatch probabilities for the period $t-1$ can be derived as described in the "Appendix," as long as distributions of the random variables are known.

The above exhibits how, once the expected recourse functions of period t are known, the expected recourse functions for period $t-1$ can be constructed through the use of a linear approximation. Thus a backward recursion is employed, which starts from period $t = N$ where, since the Nth period is the last to be considered, $\hat{\theta}_j(N+1)$ is equal to zero for all regions j. When we reach the problem for $t = 1$, we have that:

$$\psi^1(\mathbf{S}(1)) \triangleq \min_{\{x(1),y(1)\}} \sum_{i \in R} \sum_{j \in R} [(-r_{ij} + \hat{\theta}_j(2))x_{ij}(1) + (c_{ij} + \hat{\theta}_j(2))y_{ij}(1)] \quad (13)$$

subject to: $\mathbf{x}(1) \leq \mathbf{L}(N - 1)$

$$[\mathbf{x}(1) + \mathbf{y}(1)] * 1 = \mathbf{S}(1)$$

Note that as the above described procedure goes backward in time, the individual optimization problems are all solved trivially. What is required is the calculation of the dispatch probabilities (as shown in the "Appendix") and the evaluation of the parameters of the linear approximations. To satisfy that last requirement we only need to decide on the estimated vehicle supply, to give us the point where the linear approximation will be fitted. Finally, an easy network optimization problem $\psi^1(\mathbf{S}(1))$ has to be solved.

The resulting first-period network is depicted in Figure 3. Notice that each region i is now represented in all periods from $t = 1$ until $t = \max_l(1 + t_l)$, of all loads l that have region i as their destination, where t_l denotes the travel time needed for the completion of load l. The multiple representation of the unique super sink node SS is only for reasons of clarity of the resulting network.

2.3. Linear Approximations

We begin this section by proving the convexity of the expectation of the recourse function for the problem in region i at period t.

Proposition 1. *$\bar{\Psi}_i^t(S_i(t), \mathbf{\Phi}(t))$ is a convex function of $S_i(t)$.*

Proof. As shown in Appendix A of [15], for any region i and time period t, $\bar{\mathbf{\Theta}}_{ik}^t$ is an increasing function of k. This is a result of diminishing returns of the available trucks in region i at time period t. The slope of $\bar{\Psi}_i^t(S_i(t), \mathbf{\Phi}(t))$, which is the contribution of region i to the expected costs of that period $\bar{\Psi}^t(\mathbf{S}(t), \mathbf{\Phi}(t))$, with respect to supply $S_i(t) = k$ is equal to $\bar{\mathbf{\Theta}}_{ik}^t$. Since $\bar{\mathbf{\Theta}}_{ik}^t$ is a decreasing function, $\bar{\Psi}_i^t(S_i(t), \mathbf{\Phi}(t))$ is a convex function of the available supply of vehicles $S_i(t)$. (Q.E.D.)

It would be reasonable for the linear approximation of the type (4) for the total expected contribution (to costs) of region i at time period t:

$$\bar{\Psi}_i^t(S_i(t) = s, \mathbf{\Phi}(t)) \triangleq \hat{\theta}_i(t)s + b_i(t) \qquad \forall i = 1, \ldots, R$$

to be fitted around an approximation of the expected vehicle supply of region i at period t, $m_i(t) \triangleq \mathrm{E}[S_i(t)]$.

One source of information for estimating $m_i(t)$ in advance is the forecasted frequencies of outbound and inbound loaded movements of region i on a weekday corresponding to day t. These frequencies were typical of the operation of an actual carrier. The total outbound flow of region i, though, depends not only on the vehicle supply of i at that day, but also on the demand for loads out of i; similarly for the total inbound flow. Our approach is to take a weighted average of the total inbound and outbound flow of the

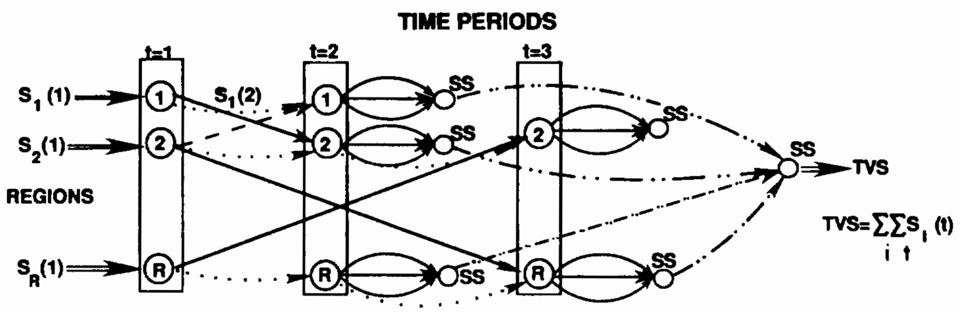

LEGEND

NETWORK LINKS (OUT OF REGION I AT TIME t)									
LOAD (DESTINATION J)		EMPTY MOVE (DESTINATION J)		HOLDING		STOCHASTIC LINK FOR VEHICLE K		DUMMYLINKS FOR SS	
COST	UPPER BOUNDS	COST	UPPER BOUNDS	COST	UPPER BOUNDS	COST	UPPER BOUNDS	COST	UPPER BOUNDS
r_{ij}	1	$-C_{ij}$	M	O	M	θ_{ik}^t	1	O	M

Fig. 3. Resulting network of the SLAP procedure.

region at that day:

$$m_i(t) = q \sum_{k \in R} \bar{\Phi}_{ki}(t-1) + (1-q) \sum_{j \in R} \bar{\Phi}_{ij}(t) \quad (14)$$

Regions with high outbound and low inbound flow are expected to attract empty vehicles from other regions to satisfy their vehicle deficit. Thus, one would expect that, since the historical empty flows are not available, the estimate of $m_i(t)$ should depend heavily on the outbound flow; thus, a value of q approaching zero would be more appropriate for these regions. Alternatively, regions with low outbound and high inbound flow are not expected to attract any empty vehicles. Thus, the total inbound flow is a good estimation of $m_i(t)$ and a value of q approaching 1 would be appropriate for these regions. In our experiments we used uniformly the value 0.5 for the parameter q.

Alternatively,

$$m_i(t) = \max[\sum_{k \in R} \bar{\Phi}_{ki}(t-1), \sum_{j \in R} \bar{\Phi}_{ij}(t)] \quad (14a)$$

may be used as an extreme case.

It is important to emphasize that the critical step for relaxing the network recourse structure in model SLAP is one that uses the linear approximation (4) instead of the nonlinear expected contribution function $\bar{\Psi}_i^t(s, \Phi(t))$ of region i at time t when the vehicle supply is s. This is equivalent to stating that the value of a truck in region i at time t when we consider the problem for the previous stage is treated as constant, and not decreasing with the supply s of that region at that time. The choice of this constant average value of a truck in i at t is very important.

In our procedure, an estimate for the expected value $m_i(t)$ is first obtained by (14). Then a choice for the average value of a truck in region i at time t may be the value of the $(m_i(t))$th vehicle in that region at that time:

$$\hat{\theta}_i(t) = \theta^t_{i,m_i(t)} \quad (15)$$

This choice is depicted in Figure 4. Note that the expected recourse function is piecewise linear so this choice actually selects the slope of the appropriate linear piece.

Alternatively, a distribution $\bar{p}_i^k(t) = P(S_i(t) = k)$ may be assumed. For example, it may be reasonable to assume that $S_i(t)$ can be approximately described using a Poisson distribution with mean $m_i(t)$. Then:

$$\hat{\theta}_i(t) = \sum_{k=1}^{\infty} \bar{p}_i^k(t)\theta_{ik}(t) \quad (16)$$

The use of an average value of a truck in a region at time t essentially allows the model to decouple the dispatch decisions of different regions at time $t-1$. When the dispatch options out of region i at time $t-1$ are reviewed, the average value of a truck in each destination region j of an option is required at time t.

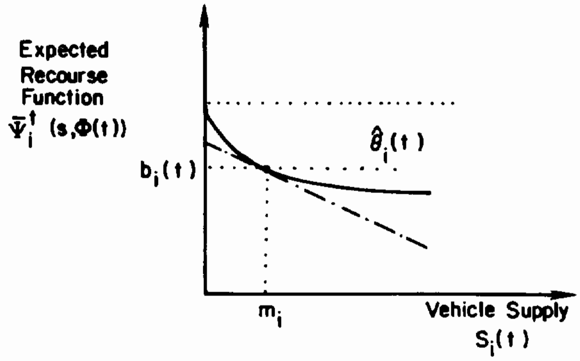

Fig. 4. Choice of the linear approximation for the SLAP procedure.

The model assumes that the dispatch decisions of the other regions k at time $t-1$ (which, naturally, affect the vehicle supply of j at t and hence the average value of a truck there) are sufficiently described by their expected value (i.e., the historical frequencies). This defines a new type of recourse, which we call Nodal Recourse, and justifies the use of a linear approximation around the approximate average vehicle supply $m_i(t)$.

Three variations of the SLAP approach were developed. In the first two, the estimates $m_i(t)$ are computed by (14) for a value of the parameter α equal to 0.5 uniformly. Model SLAP1 used the linear approximation of the type (15), while model SLAP2 used the approximation in (16) (with the assumption of a Poisson distribution). Finally, the third variation, SLAP3, used (14a) to obtain estimates for the vehicle supplies and a linear approximation of the type (15).

3. ALTERNATIVE FORMULATIONS

THE DEVELOPMENT of a stochastic programming heuristic raises the question of whether explicit treatment of randomness provides a measurable improvement than a deterministic model, especially since it generally results in a substantially more complicated model. To address that issue we contrast the SLAP model with two deterministic approximations, which were chosen based on an extensive series of experiments reported in [15]. The deterministic models are divided into single stage (or static) and dynamic models. Section 3.1 presents the single-stage formulation, which not only ignores the uncertainty of the problem, but its dynamic nature as well. Noteworthy is that, despite their obvious theoretical limitations, these models are the most commonly used in practice. Section 3.2 presents a dynamic formulation with a planning horizon of length P. Certain modeling issues are tackled with different approaches, giving rise to several alternative dynamic formulations.

3.1. Single-Stage Deterministic Formulation

A successful single-stage deterministic formulation is an assignment model which assigns available drivers to available loads with the objective of maximizing total profits in the first time period plus an estimate of the future profits, where the second term is included to induce empty repositioning moves to handle the anticipated future demand.

The goal is to find desirable vehicle allocations for the beginning of the 2nd period based on estimates about the future load demands, without losing the attractive simplicity of the model. One way is to introduce "regional salvage values" on each vehicle left in a region i at the end of the 1st period. This value, denoted by $p_i(2)$ should represent the average return of a vehicle in i for the beginning of the 2nd period until the end of an appropriate planning horizon H. Those salvage values can be calculated using a backward recursion scheme, where starting from period $H + 1$ (where they are uniformly equal to zero) the procedure goes backward assuming that in every period dispatches are done in an average way. To do this, it uses the historical frequencies of loaded and empty movements between regions. The complete procedure is described in Appendix B of [15]. In this paper the salvage values, whenever necessary, were calculated using a planning horizon of $H = 21$ days.

In this representative model:

$$\bar{\Psi}^2(\mathbf{S}(2), \mathbf{\Phi}(2)) \triangleq \textstyle\sum_{j\in R} p_j(2) S_j(2) \qquad (17)$$

and since:

$$S_j(2) = \textstyle\sum_{i\in R} (x_{ij}(1) + y_{ij}(1)) \qquad \forall j = 1, \dots, R \qquad (18)$$

the objective function reduces to:

$$\Psi^1(\mathbf{S}(1)) = \max_{\{x,y\}} \textstyle\sum_{i\in R} \sum_{j\in R} x_{ij}(1)[r_{ij} + p_j(2)] + y_{ij}(1)[-c_{ij} + p_j(2)] \qquad (19)$$

The approximation of the quite complex function $\bar{\Psi}^2(s, \mathbf{\Phi}(2))$ by a separable and linear function leads to an easily handled network formulation. A weakness of this approach is that the resulting model may prefer empty repositioning moves, which only represent anticipation of random future demands, over already existing loads $L_{ij}(1)$. This problem can be tackled by adding a large profit M to the objective coefficients of the known loads appropriately, to give them preference in the optimal solution. Additionally, as all linear models seek their optimal solutions at the extreme points of the feasible region, such a model may force too many trucks to a region j with an attractive salvage value $p_j(2)$. This tendency can be reduced by imposing upper bounds on the number of vehicles into each region. Incorporating the big M profit bonus, the objective function becomes:

$$\Psi^1(\mathbf{S}(1)) = \max_{\{x(1),y(1)\}} \sum_{i\in R}\sum_{j\in R} x_{ij}(1)[r_{ij} + M + p_j(2)] + y_{ij}(1)[-c_{ij} + p_j(2)] \qquad (20)$$

The network representation of this model, which we will refer to as DETAS, is depicted in Figure 5.

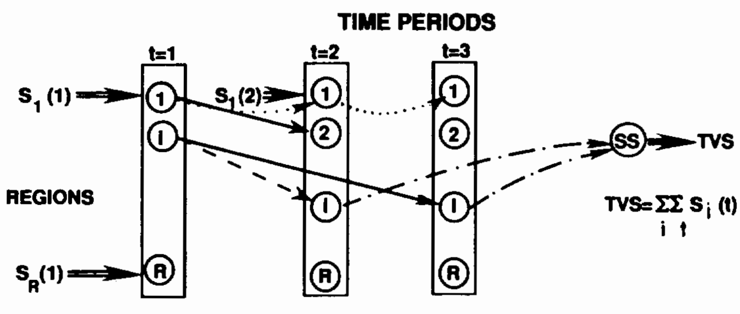

MODEL	NETWORK LINKS (OUT OF REGION I)							
	LOAD (DESTINATION J)		EMPTY MOVE (DESTINATION J)		HOLDING		ENDING LINKS	
	————		— — —		....		—·—·—·—	
	COST	UPPER BOUNDS	COST	UPPER BOUNDS	COST	UPPER BOUNDS	COST	UPPER BOUNDS
DETAS	$r_{ij}+M$	1	c_{ij}	M	0	M	$p_i(t)$	M

Fig. 5. Network representation of the single-stage deterministic formulation.

Notice the use of a time-space network to accommodate the empty movements and of a super sink node to accommodate the inclusion of upper bounds on the truck distribution after the empty repositioning. It should also be emphasized that the network representation of this model (as well as the ones to follow) does not consider simplifications A4 and A5, but depicts the model in its actually applied form. That is why we have initial vehicle supplies in regions on periods other than the first, and why a complete time-space network is needed to accommodate travel times of more than 1 day.

3.2. Dynamic Deterministic Formulations

The major weakness of the static single-stage deterministic formulations is the use of the salvage values $p_i(2)$ to approximate the forecasted activities after the assignment of loads and the empty repositioning moves. Dynamic formulations model forecasted activities deterministically over a finite planning horizon. The future activities beyond that planning horizon may be ignored, or may be approximated by the salvage values $p_i(P+1)$. Additionally, the random variables representing the demands $\Phi_{ij}(t)$ for the periods $t = 1, \ldots, P$ are substituted by their expected values, $\bar{\Phi}_{ij}(t)$. In our models we used a planning horizon $P = 7$. Use of a larger planning horizon is expected to improve the performance of these models, but since it produces a much larger network it requires prohibitive computational effort.

Then, the deterministic dynamic model is given by:

$$\Psi^1(\mathbf{S}(1)) = \max_{\{x(1),y(1),\ldots,x(P),y(P)\}} \sum_{i=1}^{P} [\mathbf{r}^{\mathrm{T}}\mathbf{x}(t) - \mathbf{c}^{\mathrm{T}}\mathbf{y}(t)] + \sum_{j\in R} p_j(P+1)S_j(P+1) \qquad (21)$$

subject to:

$$\mathbf{x}(1) \leqslant \mathbf{L}(1) \qquad (21a)$$

$$\mathbf{x}(t) \leqslant \bar{\Phi}(t) \qquad \forall t = 2, \ldots, P \qquad (21b)$$

$$[\mathbf{x}(t) + \mathbf{y}(t)] * 1 = \mathbf{S}(t) \qquad \forall t = 1, \ldots, P \qquad (21c)$$

$$[\mathbf{x}(t) + \mathbf{y}(t)]^{\mathrm{T}} * 1 = \mathbf{S}(t+1) \qquad \forall t = 1, \ldots, P \qquad (21d)$$

This model can be represented as a time-space network, where a node represents a region at a time period. There are three types of links in this network:

—Links representing known loads: they have a cost coefficient equal to the direct contribution of the load, and an upper bound of 1. Because of this assumption that only loads for the first day are known, these links emanate only from nodes representing regions at the first time period.

—Links representing empty movements: they have a cost coefficient equal to minus the cost of moving empty between the origin and destination regions, and are unbounded.

—Links representing forecasted loads: they have a cost coefficient equal to the historical average direct contribution of loaded movements between the origin and destination regions involved. Because of the assumption that all loads for the first day are known, these links emanate only from nodes that represent regions at time periods greater than or equal to 2. A natural upper bound for the links representing forecasted loads would be the historical frequencies of loaded movements, $\bar{\Phi}_{ij}(w)$, between the origin and destination regions on a weekday w corresponding to that time period. These frequencies are typically fractional and 90% of them are between 0 and 1 and to use them results in a network with fractional upper bounds. To solve it, the units of the network are changed to 1/10 or 1/100 of a truck and after the optimal solution is obtained, the optimal flows are rounded to multiples of 10 or 100 accordingly to correspond to "whole" trucks. The network representation of that model, which we will refer to as DETD, is depicted in Figure 6.

There are many variations of these deterministic approximations. Frantzeskakis and Powell[15] compare a number of these variations. The models used here produced the best results among the static and dynamic models that were tested.

4. EXPERIMENTAL DESIGN

THIS SECTION provides details about our experimental design. The following sections include comments made on the implementation of the rolling horizon procedure and a discussion of the data requirements. The performance indicators that we are interested in are explained in the third section. Finally, a useful deterministic model which can be used as an upper bound on the performance of the models is presented in the last section.

4.1. Comments on the RHP Procedure

In order to contrast the performance of the alternative models, we use a rolling horizon procedure (RHP). The implementation of an RHP for a model is quite simple. At any time t, the RHP models time from t up to $t + P - 1$, where P is the planning horizon of that model. The model is solved for that time t and its recommended actions are identified. However, we only implement the instructions for time t, and then we advance the clock to handle time $t + 1$. This is done for $t = 1, \ldots, T$, where T is the length of the simulation. It is important to choose T sufficiently

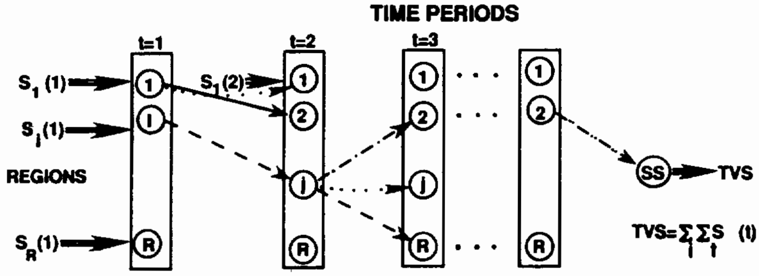

		NETWORK LINKS (OUT OF REGION i)									
	VEHICLE SUPPLIES $VS_i(t)$	LOAD (DESTINATION j)		EMPTY MOVE (DESTINATION j)		FORECASTED LOADED MOVE (WEEKDAY W)		HOLDING		ENDING LINKS	
MODEL		COST	UPPER BOUNDS	COST	UPPER BOUNDS	COST	UPPER BOUNDS	COST	UPPER BOUNDS	COST	UPPER BOUNDS
DETD	$s_i(t)$ 100	r_{ij}	100	c_{ij}	M	r_{ij}	$f_{ij}(w)$ 100	0	M	0	M

Fig. 6. Network representation of the dynamic deterministic function.

large to capture the dynamic effects and to mitigate statistical sampling errors.

As the rolling horizon procedure (RHP) proceeds, the profits of the implemented recommended actions of the model are accumulated. When the RHP terminates, the total profits obtained by each model serve as the evaluating criterion for its performance. Notice that the sum does not involve discounting, since the length of a time period is relatively small (1 day). If financial models were to be evaluated with an RHP approach, discounting should probably be present, since they usually involve monthly or even longer time periods.

The RHP spans a period of $T = 12$ days and actually simulates the operation of a carrier that uses that model as his dispatching tool during that period. All the models are tested under uniform conditions, that is under the same vehicle allocation in the beginning of the RHP ($t = 1$) and with the same load opportunities during the period that the RHP spans. Since no loads are carried from one day to another, the results are easier to evaluate and the performance of each model on a particular day of the RHP is directly and entirely dependent on whether the model had successfully predicted the future in the previous days of the RHP and arranged for an appropriate vehicle allocation.

4.2. Data Requirements

This section describes the data used in the numerical testing of the models; all the experiments are based on data obtained from a major motor carrier.

All models were tested in the RHP environment with the same initial vehicle allocation. For each vehicle, the region and the time of first availability have to be known. The peformance of the models is expected to depend on the fleet size. In the extreme case of too many vehicles, the models that are myopic, which do not recommend any empty movements, may do as well as the more advanced models that forecast the future, or even better. When the fleet is too small, all models basically assign the vehicles to the best available loads and the different approaches basically converge, in terms of the recommended solution. Thus, the models are tested for fleet sizes of 100, 250, 500, 750, 1000 and 1500 vehicles. All vehicles are assumed of the same type and all loads can be carried from that type of vehicle.

Necessary data for the dispatch of the vehicles include average travel times and costs for empty and loaded movements between each pair of regions. Also (to evaluate forecasted loaded movements) the average revenue obtained by carrying a load is needed for each pair of regions. Note that the part of the loaded move between the location of the vehicle (assumed to be the centroid of the origin region) and the location from which the load is to be picked up (within the same region) is ignored. Also ignored is the part of the move between the unloading location of a load at its destination region and the centroid of that region. Thus, as soon as a vehicle unloads, it is assumed to be available for the next dispatch immediately. The cost of holding a vehicle overnight in a region is assumed equal to zero and the vehicles can load or unload anytime during the day or any day during the week.

Historical data on the average frequencies of loaded and empty moves between each pair of regions are needed for the dynamic models. The loaded moves are assumed to follow a weekly pattern (so the frequencies are a function of the day of the week) and, since the span of the RHP is small, no adjustments were made for seasonal effects.

The load opportunities that the models are faced with during the time span of the RHP are represented by a list of loads that will materialize in that period of T days. This realization of the random demands $\Phi_{ij}(t)$ for $t = 2, \ldots, T$, (with the known loads $L_{ij}(1)$ attached to it) was created randomly, using the historical average frequencies $\bar{\Phi}_{ij}$ described above. The origin and the destination of each load is needed, as well as the date of departure of the load and the revenue associated with that move. Remember that, because of our assumptions, the date that the shipper calls in to request a vehicle coincides with the date that the vehicle is to pick-up the load from the carrier. That list of loads contains on the average 316 loads per day for the regular weekdays and 32 loads per day for the weekend days.

4.3. Performance Measures

The major criterion in assessing the performance of the models was the total profits that the carrier would obtain by actually implementing their dispatching recommendations. These profit figures obviously refer to the period span of the RHP and to the particular realization of the loads for that period.

Since the motor carrier environment is highly competitive, it is reasonable for the carriers to be concerned with the level of service that they provide; this may be represented, in a way, by the percentage of loads accepted by the carrier. Additional measures are the number of empty movements recommended by the model and the percentage of the empty miles (over the total miles) traveled, as well as the average revenue per accepted load and the average distance of an empty movement recommended by the model. Additional insight may be obtained by examining the daily values of the above mentioned performance measures.

Finally, the required CPU time on a MICROVAX-2 is reported to indicate the usefulness of each model as a real-time dispatching tool. These times include only the processing time needed to solve the models and did not include all the necessary updating of the vehicle allocation and the list of known loads for the rolling horizon procedure.

4.4. Upper Bound

The numerical experiments focus on the relative performance of the SLAP algorithm versus the deterministic approximations. A separate measure of performance is obtained by comparing these results to a theoretical upper bound, obtained as follows. We take the realizations of the forecasted loads that were used in the RHP procedure (that is, the list of loads described above) and build a single, deterministic network spanning the 12 days of the simulation. The result of this single optimization provides an upper bound, which we will refer to as UB, since decisions made in time t are based on what will actually happen in the future. In addition, the optimization can take into account the actual length of the simulation, allowing it to "cheat" toward the end of the simulation period. For example, its solution does not include any empty movements on the last day of the simulation.

5. EVALUATION OF EXPERIMENTAL RESULTS

THE COMPUTATIONAL results of the performance of the models for the 12-day period are presented in Table I for the different fleet sizes. For each model the following performance measures are reported: total profits, percentage of loads accepted, revenue per accepted load, number of empty movements recommended, average distance per empty movement recommended, percentage of empty miles traveled (over total miles), total revenues and costs. Additional results involving the daily patterns of these performance measures, as well as of the pattern of accesssible loads, were provided, but are not presented in this paper.

As far as the three SLAP models are concerned, the inclusion of the Poisson distribution in SLAP2 generally resulted in slightly higher total profits than SLAP1 (up to 4%), while SLAP3 had equivalent performance for small fleet sizes but performed worse as the fleet size increased. Hence, SLAP2 will be the representative of the stochastic models used in the following comparisons to the deterministic modeling approaches.

It should be emphasized that the strongest point of the stochastic models is not necessarily the percentage of accepted loads. Actually, for fleet sizes up to 750 vehicles they were outperformed by the dynamic deterministic models, only to reverse that when greater fleet sizes were involved. Nevertheless for fleet sizes of up to 250 vehicles, model DETD was accepting more loads than the UB model, which indicates that DETD was forcing into its optimal solution loads that ultimately were not as profitable. Notice also that for the biggest fleet size (1500 vehicles), the optimal solution of the UB model accepted virtually all (99.5%) of the materialized loads.

As far as the profitability of the accepted loads is concerned, the DETD model consistently yielded the lowest average revenue per accepted load, even when it was accepting fewer loads. When the models SLAP and DETD are compared, the difference in profitabil-

TABLE I

Results Obtained by the 12-Day RHP for the Alternative Models

Fleet Size	Model	Profits (th. $)	% Loads Accept.	Reven./ Load Accept. ($)	# Empty Moves	% Empty Miles	Miles/ Empty Move (mi.)s	Total Revenue (th. $)	Total Cost (th. $)
100	DETAS	209.7	10.3	1875	13	0.5	214	624	415
	DETD	201.6	14.6	1236	18	0.9	233	580	378
	SLAP-1	214.2	12.9	1476	4	0.1	169	614	400
	SLAP-2	214.5	12.8	1494	5	0.2	236	617	403
	SLAP-3	215.7	12.8	1510	2	0.0	107	620	405
	UB	249.7	13.3	1743	25	0.2	64	746	496
250	DETAS	486.4	28.0	1618	75	1.3	224	1461	975
	DETD	465.1	32.6	1299	75	1.7	261	1367	902
	SLAP-1	504.2	30.6	1497	32	0.5	185	1476	972
	SLAP-2	505.0	30.3	1517	36	0.5	189	1482	977
	SLAP-3	505.2	30.0	1537	34	0.5	194	1484	979
	UB	545.0	31.7	1607	35	0.6	244	1644	1099
500	DETAS	813.9	52.6	1478	349	3.0	198	2505	1691
	DETD	805.2	57.0	1329	255	3.5	289	2442	1603
	SLAP-1	893.3	54.9	1522	169	1.6	228	2691	1798
	SLAP-2	893.0	54.9	1522	170	1.6	225	2692	1789
	SLAP-3	891.0	54.1	1546	183	1.8	245	2699	1808
	UB	923.9	57.7	1526	157	1.6	253	2841	1917
750	DETAS	984.0	67.6	1421	667	4.6	198	3096	2112
	DETD	1026.4	76.6	1277	424	4.1	263	3152	2126
	SLAP-1	1140.6	74.7	1462	413	3.3	253	3522	2382
	SLAP-2	1154.2	76.2	1453	429	3.3	251	3568	2414
	SLAP-3	1133.9	74.7	1468	421	3.9	302	3537	2403
	UB	1188.0	79.2	1455	355	2.6	244	3712	2524
1000	DETAS	1045.5	74.4	1405	1012	6.2	193	3367	2322
	DETD	1128.7	85.9	1242	451	3.3	215	3439	2310
	SLAP-1	1255.8	88.4	1398	669	5.1	283	3983	2727
	SLAP-2	1268.2	89.4	1392	638	4.9	283	4012	2743
	SLAP-3	1217.3	86.4	1408	741	6.2	305	3923	2706
	UB	1332.6	95.7	1377	536	3.7	267	4246	2913
1500	DETAS	1087.9	82.0	1391	1681	9.2	194	3677	2590
	DETD	1181.2	89.3	1229	282	1.7	187	3537	2355
	SLAP-1	1203.8	90.4	1401	1104	9.3	334	4082	2878
	SLAP-2	1254.6	92.6	1395	1008	8.0	317	4166	2911
	SLAP-3	1132.0	87.6	1413	1294	11.5	352	3989	2857
	UB	1394.8	99.5	1362	395	2.3	229	4367	2972

ity was very clear, ranging from 13% to 22%, and generally dropping as the fleet size increases. When models SLAP and DETAS are contrasted, for small fleet sizes (up to 250 vehicles) the stochastic model fell short of the deterministic model in terms of the quality of the accepted loads. This is, nevertheless, misleading since the DETAS model was accepting much fewer loads and was in that sense able to be more selective about their profitability. For medium fleet sizes (500 and 750 vehicles), the stochastic model maintained higher (up to 5%) revenue per accepted load, although it was accepting slightly more loads (2–8%). For bigger fleet sizes, the two models had virtually equal revenue per accepted load, while the stochastic model was accepting substantially more loads (11–14%).

In terms of the empty repositioning moves, the DETAS model generally created more than any other model, while the number of empty moves recommended by SLAP increased with the fleet size. Thus, while the stochastic model created fewer empty moves than DETD for fleet sizes of up to 500 vehicles, the situation reversed dramatically after that fleet size. In terms of the average distance of the empty moves, model DETAS maintained an approximately steady figure of about 200 miles, while model DETD increased the distance from 230 miles (100 vehicles) to 290 miles (500 vehicles) and then decreased it to 190 miles (1500 vehicles). Meanwhile, the stochastic model recommended empty moves with average distance that increased from 170 to 320 miles as the fleet size increased. Nevertheless, for fleet sizes of

TABLE II

Comparative Results of Profits Obtained by RHP and of Upper Bound

Fleet Size	Best Static Determ. (a)	Best Dynam. Determ. (b)	(b) − (a)/(a) % (c)	Best Stoch. Model (d)	(d) − (a)/(a) % (e)	(d) − (b)/(b) % (f)	Upper Bound (g)	(g) − (d)/(d) % (h)
100	209.7	201.6	−3.8	215.7	2.8	7.0	249.7	15.8
250	486.4	465.1	−4.4	505.2	3.9	8.6	545.0	7.9
500	813.9	805.2	−1.1	893.3	9.8	10.9	923.9	3.4
750	984.0	1026.4	4.3	1154.2	17.3	12.5	1188.0	2.9
1000	1045.5	1128.7	8.0	1268.2	21.3	12.4	1332.6	5.1
1500	1087.9	1181.2	8.6	1254.6	15.3	6.2	1394.8	11.2

500–1000 vehicles, the UB model recommended empty movements of average distance that was closer to that of the SLAP model's empty movements, rather than the other models. The combination of these two factors caused the substantial increase in percentage of empty mileage recommended by the stochastic model as the fleet size increased, with obvious effects in the total costs of implementing those decisions.

The total profits obtained by the RHP procedure for the alternative models are reproduced in Table II in columns (a), (b) and (d). Columns (c) and (e) indicate the percent improvement in profits from using the DETD and SLAP model over using the DETAS model. They range from −4.4% to 8.6% for the DETD model and from 2.8% to 21.3% for the SLAP model and they generally increase with the fleet size. Column (f) presents the percent improvement from using the SLAP model over using the DETD model, ranging from 7.0% to 12.5% and (with the exception of the 1500 vehicles) increases with the fleet size.

The fact that the gap in profits between the best representatives of the three major approaches increases generally with the fleet size, indicates that the diminishing returns effect is stronger when the vehicles are dispatched with deterministic models. Thus, it is clear that a poor strategy of managing a fleet may lead to false strategic planning decisions concerning the optimum fleet size of a company.

In column (g) of the same table, the upper bound in profits is presented. This bound is obtained by solving optimally a 12-day network with the same initial vehicle allocation, where all the loads that will materialize all through that period are known in advance. In column (h) the gap between the SLAP model and the upper bound is presented, as a percent of the upper bound value. This gap ranged from 15.8% to only 2.9%, and reaffirmed our belief that there would be a mid-range "optimum" fleet size for the SLAP model. This is inherent in the construction of the SLAP algorithm, since big fleet sizes make the assumption of a constant "average value" of a truck in a region at a time period less valid and forces into the solution a higher number of empty moves (some of which will turn out to be not so unprofitable in the future).

Finally, average CPU times for solving the models for one day of the rolling horizon were of the order of 31 sec for DETAS, 72 secs for DETD and 767 secs for SLAP.

6. GENERAL CONCLUSIONS AND FUTURE RESEARCH

WE BELIEVE that the experiments conducted showed a clear superiority of the SLAP algorithm, as far as the carrier's total profits are concerned. Future research should probably concentrate into this area and especially in algorithms that would produce an accurate representation of the stochastic vehicle supplies of the regions at any time period.

Additionally, more strict upper bounds on the total profits could be obtained that would not totally disregard the presence of uncertainty, as the one used in this paper was. This would indicate more clearly where these models stand and how much room for their improvement exists.

APPENDIX: CALCULATION OF THE DISPATCH PROBABILITIES

THE DERIVATION of the dispatch probabilities d^t_{ikn} is described in this Appendix. Under the demand realizations $\phi(t)$, the event that the kth vehicle is dispatched on the 1st option is equivalent to the event that the demand for that option is more than or equal to k (if that option is a loaded movement) or it is an event with probability 1 (if that option is an empty one). So:

$$d^t_{ik1} = \text{Prob}[\phi_{ij}(t) \geq k]$$

if the first option is a loaded movement to region j, or $d_{ik1} = 1$ if it is an empty movement.

Under the demand realizations $\phi(t)$, the event that the kth vehicle is dispatched on the nth option out of region i is equivalent to the joint event that cumulatively the first (best) $n - 1$ options take less than k vehicles and the first n options take more than or equal to k vehicles. To state it mathematically, let:

$q_{in} = j$, if the nth dispatch option out of region i at time t is to move to region j.

Then:

$$\phi_i^n(t) = \sum_{m=1}^{n} \phi_{iq_{in}}(t)$$

and

$$\begin{aligned} d_{ikn}^t &= \text{Prob}[\phi_i^{n-1}(t) < k \text{ and } \phi_i^n(t) \geq k] \\ &= \text{Prob}[\phi_i^{n-1}(t) < k] + \text{Prob}[\phi_i^n(t) \geq k] \\ &\quad - \text{Prob}[\phi_i^{n-1}(t) < k \text{ or } \phi_i^n(t) \geq k] \\ &= \text{Prob}[\phi_i^{n-1}(t) < k] + \text{Prob}[\phi_i^n(t) \geq k] - 1 \\ &= \text{Prob}[\phi_i^{n-1}(t) < k] - \text{Prob}[\phi_i^n(t) < k] \end{aligned}$$

We may assume that the actual number of loads $\phi_{ij}(t)$ follows a Poisson distribution with mean equal to the historical frequency $\overline{\Phi}_{ij}(t)$. Since the cumulative variable $\boldsymbol{\phi}_i^n$ is the sum of variables following Poisson distributions, it follows Poisson distribution as well with mean equal to the sum of the individual variables' means. We should note that the assumption of the Poisson distribution is not necessary but greatly simplifies the calculation of the dispatch probabilities and can be justified theoretically.

The empty options, on the other hand, can accommodate an unbounded number of vehicles. Thus, if the highest ranking empty option out of region i has rank n_e and if the kth vehicle is dispatched on that option, then all the other vehicles with indices l, $\forall l \geq k$, are dispatched on n_e as well. So, for $k < S_i(t)$,

$$\begin{aligned} d_{ikn_e}^t &= \text{Prob}[\phi_i^{n_e-1}(t) < k \text{ and } \phi_i^{n_e}(t) \geq k] \\ &= \text{Prob}[\phi_i^{n_e-1}(t) < k] + \text{Prob}[\phi_i^{n_e}(t) \geq k] \\ &\quad - \text{Prob}[\phi_i^{n_e-1}(t) < k \text{ or } \phi_i^{n_e}(t) \geq k] \\ &= \text{Prob}[\phi_i^{n_e-1}(t) < k] + 1 - 1 \\ &= \text{Prob}[\phi_i^{n_e-1}(t) < k] \end{aligned}$$

and $d_{ikn_1}^t = 0$, $\forall n_1 > n_e$. Thus, the summation in (12) actually runs up to $n = n_e$.

We would like to emphasize that no assumption has been made to limit the number of vehicles dispatched on the highest ranked empty option. Although such an assumption seems very reasonable for practical applications, it is inconsistent with our linear approximation and was thus avoided.

REFERENCES

1. L. Baker, "Overview of Computer Based Models Applicable to Freight Car Utilization," Report prepared for the U.S. Department of Transportation, NTIS, Springfield, Va., 1977.
2. E. M. Beale, "On Minimizing a Convex Function Subject to Linear Inequalities," *J. R. Stat. Soc.* **B17,** 173–184 (1955).
3. E. M. Beale, J. J. Forest and C. J. Taylor, "Multi-time Period Stochastic Programming," in *Stochastic Programming*, M. A. Dempster (ed.), Academic Press, New York, 1980.
4. E. M. Beale, G. B. Dantzig and R. D. Watson, "A First Order Approach to a Class of Multi-time Period Stochastic Programming Problems," *Math. Program. Study* **27,** 103–177 (1986).
5. J. Birge, "Solution Methods for Stochastic Dynamic Linear Programs," Technical Report SOL 80-29, Systems Optimization Laboratory, Stanford University, 1980.
6. J. Birge, "Decomposition and Partitioning Techniques for Multistage Stochastic Linear Programs," *Opns. Res.* **33,** 989–1007 (1985).
7. J. Birge and R. Wets, "Designing Approximation Schemes for Stochastic Optimization Problems in Particular for Stochastic Programs with Recourse," *Math. Program Study* **27,** 54–102 (1986).
8. J. H. Bookbinder and S. P. Sethi, "The Dynamic Transportation Problem: A Survey," *Naval Res. Logist. Quart.* **27,** 447–452 (1980).
9. L. Cooper and L. J. LeBlanc, "Stochastic Transportation Problems and Other Network Related Convex Problems," *Naval Res. Logist. Quart.* **24,** 327–337 (1977).
10. G. Dantzig, "Linear Programming under Uncertainty," *Mgmt. Sci.* **1,** 197–206 (1955).
11. G. Dantzig and P. Wolfe, "Decomposition Principle for Linear Programs," *Opns. Res.* **8,** 101–111 (1960).
12. G. Dantzig and A. Madansky, "On the Solution of 2-Stage Linear Programs with Uncertainty," in *Proceedings, 4th Berkeley Symposium on Mathematics Statistics and Probability,* University of California Press, Berkeley, 1961.
13. P. J. Dejax and T. G. Crainic, "Models for Empty Freight Vehicle Transportation Logistics," *Trans. Sci.* **21,** 227–247 (1987).
14. M. Frank and P. Wolfe, "An Algorithm for Quadratic Programming," *Naval Res. Logist. Quart.* **3,** 95–110 (1956).
15. L. F. Frantzeskakis and W. B. Powell, "Development and Evaluation of a Successive Linear Approximation Procedure for Stochastic, Dynamic Networks," Technical Report SOR-88-13, Operations Research and Statistics Department, Princeton University, 1988.
16. W. C. Jordan, "The Impact of Uncertain Demand and Supply on Empty Rail Car Distribution," Ph.D. dissertation, Cornell University, Ithaca, N.Y., 1982.
17. W. C. Jordan and M. A. Turnquist, "A Stochastic, Dynamic Model for Railroad Car Distribution," *Trans. Sci.* **17,** 123–145 (1983).
18. P. Kall, "Computational Methods for Solving 2-Stage Stochastic Linear Programming Problems," *J. Appl. Math. Physics* **30,** 261–271 (1979).

19. F. LOUVEAUX, "A Solution Method for Multi-Stage Stochastic Programs with Recourse, with Application to an Energy Investment Problem," *Opns. Res.* **28,** 889-902 (1980).
20. S. C. MISRA, "Linear Programming of Empty Wagon Disposition," *Rail Int.* **3,** 151–158 (1972).
21. W. B. POWELL, Y. SHEFFI AND S. THIRIEZ, "The Dynamic Vehicle Allocation Problem with Uncertain Demands," in *Ninth International Symposium on Transportation and Traffic Theory*, VNU Science Press, Utrecht, 1984.
22. W. B. POWELL, "A Stochastic Model of the Dynamic Vehicle Allocation Problem," *Trans. Sci.* **20,** 117–129 (1986).
23. W. B. POWELL, "A Comparative Review of Alternative Algorithms for the Dynamic Vehicle Allocation Problem," in *Vehicle Routing: Methods and Studies*, B. I. Golden and A. A. Assad (eds.), North-Holland, Amsterdam, 1988.
24. L. QI, "Forest Iteration Method for Stochastic Transportation Problem," *Math. Program Study* **25,** 142–163 (1985).
25. R. ROCKAFELLAR AND R. WETS, "Scenarios and Policy Aggregation in Optimization Under Uncertainty," Working Paper WP-87-119, IIASA, Laxenburg, Austria (December), 1987.
26. B. STRAZICKY, "Some Results Concerning an Algorithm for the Discrete Recourse Problem," in *Stochastic Programming*, M. A. Dempster (ed.), Academic Press, New York, 1980.
27. R. VAN SLYKE AND R. WETS, "L-Shaped Linear Programs with Applications to Optimal Control and Stochastic Programming," *SIAM J. Appl. Math.* **17,** 638-663 (1969).
28. D. WALKUP AND R. WETS, "Stochastic Programs with Recourse," *SIAM J. Appl. Math.* **15,** 1299-1314 (1967).
29. D. WALKUP AND R. WETS, "Stochastic Programs with Recourse II: On the Continuity of the Objective," *SIAM J. Appl. Math.* **17,** 98-103 (1969).
30. S. WALLACE, "Decomposing the Requirement Space of a Transportation Problem into Polyhedral Cones," *Math. Program Study* **28,** 29–47 (1986).
31. S. WALLACE, "Solving Stochastic Programs with Network Recourse," *Networks* **16,** 295–317 (1986).
32. S. WALLACE, "A Two-Stage Stochastic Facility Location Problem with Time-Dependent Supply," in *Numerical Methods in Stochastic Programming*, Y. Ermoliev and R. Wets (eds.), Springer-Verlag, Berlin, 1988.
33. S. WALLACE, "Minimum Cost Problems in Stochastic Networks," Working Paper, Chr. Michelsen Institute, Norway.
34. S. WALLACE AND K. BREKKE, "Optimal Fleet Size when National Quotas can be Traded," *Mar. Resource Econ.* **2,** 315–329 (1986).
35. R. WETS, "Programming under Uncertainty: The Equivalent Convex Program," *SIAM J. Appl. Math.* **14,** 89–105 (1966).
36. R. WETS, "Programming under Uncertainty: The Solution Set," *SIAM J. Appl. Math.* **14,** 1143–1151 (1966).
37. R. WETS, "Stochastic Programs with Fixed Resources: The Equivalent Deterministic Problem," *SIAM Rev.* **16,** 309–339 (1974).
38. R. WETS, "Stochastic Programming: Solution Techniques and Approximation Schemes," in *Mathematical Programming: The State of the Art*, pp. 566–603, A. Bachen, M. Grotchel and B. Korte (eds.), Springer-Verlag, Berlin, 1982.
39. R. WETS, "Solving Stochastic Programs with Simple Recourse," *Stochastics* **10,** 219–242 (1983).
40. W. W. WHITE AND A. M. BOMBERAULT, "A Network Algorithm for Empty Freight Car Allocation," *IBM Syst. J.* **8,** 147–169 (1969).
41. R. WOLLMER, "Two State Linear Programming Under Uncertainty with 0–1 Integer First Stage Variables," *Math. Program.* **19,** 279–288 (1980).
42. W. ZIEMBA, "Computational Algorithms for Convex Stochastic Programs with Simple Recourse," *Opns. Res.* **18,** 414–431 (1970).

(Received, November 1988; revision received June 1989; accepted August 1989)