



Dynamic Multi-period Freight Consolidation

Arturo Pérez Rivera, Martijn Mes

Beta Working Paper series 473

BETA publicatie	WP 473 (working paper)
ISBN	
ISSN	
NUR	804
Eindhoven	April 2015

Dynamic Multi-period Freight Consolidation

Arturo Pérez Rivera and Martijn Mes

University of Twente

Department of Industrial Engineering and Business Information Systems

P.O. Box 217, 7500 AE Enschede, The Netherlands

{a.e.perezrivera,m.r.k.mes}@utwente.nl

Abstract. Logistic Service Providers (LSPs) offering hinterland transportation face the trade-off between efficiently using the capacity of long-haul vehicles and minimizing the first and last-mile costs. To achieve the optimal trade-off, freights have to be consolidated considering the variation in the arrival of freight and their characteristics, the applicable transportation restrictions, and the interdependence of decisions over time. We propose the use of a Markov model and an Approximate Dynamic Programming (ADP) algorithm to consolidate the right freights in such transportation settings. Our model incorporates probabilistic knowledge of the arrival of freights and their characteristics, as well as generic definitions of transportation restrictions and costs. Using small test instances, we show that our ADP solution provides accurate approximations to the optimal solution of the Markov model. Using a larger problem instance, we show that our modeling approach has significant benefits when compared to common-practice heuristic approaches.

Keywords: Intermodal transportation, transportation planning, consolidation, time horizon, approximate dynamic programming

1 Introduction

Over the last decade, the hinterland transportation industry has experienced a change towards network oriented services. Many Logistic Service Providers (LSPs) now offer multiple services such as pick-up, storage, long-haul and final delivery of freight. With this change, new challenges arise for LSPs who organize their processes (and possibly carriers) in such a way that the efficiency of their entire transportation network is improved. We investigate one of such challenges encountered by a LSP in The Netherlands. On a daily basis, this Dutch LSP transports containers from the East of the country to different terminals in the port of Rotterdam. This LSP has reserved capacity on a barge to transport its containers. The costs of the long-haul are fixed, but the last-mile costs come from the time required for sailing, waiting, and handling of containers at 12 container terminals spread over a distance of 40km in the port of Rotterdam. The challenge is then to consolidate containers in such a way that each day, only a few close-by terminals are visited and the reserved barge capacity is used efficiently, over time.

In operations research terms, we study the planning problem that arises when a company wants to transport freights (e.g., containers) from a single origin to different destinations, periodically (e.g., daily). The destinations of these freights are always far away and closer among themselves than to the origin. For this reason, the long-haul is the same in every trip, independent of which freights were consolidated at the origin (e.g., a barge sailing over the same river). However, the last-mile route varies according to the destinations of the freights that were consolidated at the beginning of the long-haul. In addition, there is also an alternative mode (e.g., truck) that can be used to transport freights directly from their origin to their destination. The objective of the company is to reduce its total costs over time and to use the vehicle's capacity efficiently.

Companies with the aforementioned characteristics usually have fixed long-haul costs. Consequently, costs savings are only possible in the last-mile and in the use of the alternative transportation mode. The first source of costs is influenced by factors such as unloading time, waiting time, service reliability, etc. As a result, combinations of destinations might have different last-mile costs even when the transportation distance between them is the same. The second source of costs depends on the use of the alternative mode. This situation occurs when there are more urgent freights than the long-haul vehicle's capacity. Properly balancing the consolidation and postponement of freights is therefore a challenge for the company but also a necessity for its efficient operation.

For several reasons, consolidating freights in a way that minimize costs over time is not a straightforward task. First of all, the number of freights that arrive and their characteristics, vary from day to day. This uncertainty makes it difficult to know which freights to postpone for future consolidation. Second, each freight that arrives has a fixed time-window for transportation. Furthermore, not all freights which arrive on the same day have the same destination or time-window. Third, the objective of carrying as many freights as possible in the long-haul vehicle during each trip can be conflicting with the objective of reducing last-mile costs in the long run. To handle these planning challenges and to reduce costs over time, we propose the use of a Markov model and an Approximate Dynamic Programming (ADP) algorithm.

The remaining of this paper is organized as follows. In Section 2, we briefly introduce the relevant scientific literature on dynamic multi-period freight consolidation and outline our contribution to it. In Section 3, we present the mathematical notation of the problem characteristics and the formulation of the Markov model. Also in this section, we present our ADP approach. In Section 4, we carry out a series of numerical experiments. We close with conclusions and future research directions in Section 5.

2 Literature Review

In this section, we briefly analyze the scientific literature on freight consolidation in intermodal transportation networks. More specifically, we look at literature about problems where transportation modes are chosen dynamically for different

types of freights. We shortly examine the advantages, limitations, and extension opportunities of the models and the solution methods proposed in this type of papers. For a comprehensive literature review on strategical, tactical, and operational planning problems in intermodal transportation networks we refer to SteadieSeifi et al. [12] and Crainic and Kim [5].

The problem in this paper falls into the category of Dynamic Service Network Design (DSND) problems. Most DSND models assume deterministic demand [12]. In addition, most models consider the context of a single carrier and cyclically scheduled services where there are hardly any time-dependencies, even when there are multi-period horizons [5]. Although there are exceptions to these shortcomings, models seem to focus on one exception at a time and leave out the rest. For example, models which include *multiple* modes of transportation, such as [9], usually do not incorporate time issues. On the other hand, the few models which include *time dependencies*, such as [3], are developed for a single mode of transportation. The few models that include uncertainty in the demand, such as [6], are usually developed for the road transportation mode.

Most DSND solution approaches are based on graph theory, mathematical programming techniques, and heuristics [12, 14]. Solutions based on graph theory can not deal with time-dependencies for large instances and assume deterministic demand most of the time. To avoid these shortcomings and handle the complexities of large size problems, mathematical programming techniques such as cycle-based variables [2], branch-and-price [1], or column generation [9] have been proposed. Also to avoid such complexity issues, metaheuristic extensions such as Tabu Search [4, 13] have been vastly proposed [12]. A disadvantage of most of these heuristics and mathematical programming techniques is that they are less suitable for stochastic settings. Further design such as stochastic scenarios [6] or probabilistic constraints is required to incorporate stochastic elements on these solution approaches. Nevertheless, the need and the benefits of introducing stochastic elements into DSND formulations have been widely recognized in practice [8].

As mentioned by Wieberneit [14], realistic instances of DSND problems are difficult to solve with the exact approaches presented in the literature. Although these exact approaches have been studied for some years now [7], research about the use of approximations and decompositions, especially for stochastic multi-period problems, has been scarce [8, 12, 14]. Considering these challenges and opportunities, we believe our contribution to the scientific literature of DSND problems and intermodal transportation planning is two-fold. First, we develop a Markov model that handles complex time dependencies, incorporates stochastic demand (and its characteristics), for a multi-period horizon, and has a generic definition of costs depending on the destinations of freights. Second, we develop an Approximate Dynamic Programming (ADP) solution algorithm that makes the aforementioned model computationally applicable to realistic-size problems.

3 Problem Description and Formulation

We consider a dynamic multi-period long-haul freight consolidation problem in which decisions are made on consecutive periods t over a finite horizon $\mathcal{T} = \{0, 1, 2, \dots, T^{max} - 1\}$. For simplicity, in the remaining of the paper we refer to a period as a day. The main decision at each day is which of the known, and released-for-transport, freights to transport using the long-haul vehicle. Each freight must be delivered to a given destination d from a group of destinations \mathcal{D} within a given time-window. The time-window of a freight begins at a release-day $r \in \mathcal{R} = \{0, 1, 2, \dots, R^{max}\}$ and ends at a due-day $r + k$, where $k \in \mathcal{K} = \{0, 1, 2, \dots, K^{max}\}$ defines the length of the time-window. The arrival-day t of a freight is the moment when all its information is known to the planner. Note that r influences how long the freights are known before they can be transported, and thus influences the degree of uncertainty in the decisions. We consider that the destinations \mathcal{D} , release-days \mathcal{R} , and time-window lengths \mathcal{K} are known finite sets for the entire planning horizon \mathcal{T} .

New freights become available as time progresses. These freights and their characteristics are unknown before they arrive, but the planner has some probabilistic knowledge about them. First, we define F as the discrete and finite random variable describing the variation in the total number of freights arriving per day. Second, we define D as the random variable describing the variation in the destination of each freight, with possible values of $D \in \mathcal{D}$. Finally, R and K are two random variables with possible values of $R \in \mathcal{R}$ and $K \in \mathcal{K}$, which describe the variation of the time-windows of each freight. We consider that between two consecutive days, a number of freights f arrive with probability p_f^F , independent of the arrival day. Each freight has destination d with probability p_d^D , release-day r with probability p_r^R , and time-window length k with probability p_k^K , independent of the day and of other freights.

The last-mile costs depend on the subset of destinations visited. We denote a subset of destinations with $\mathcal{D}' \subseteq \mathcal{D}$, and denote its associated cost with $C_{\mathcal{D}'}$. At each day, there is only one long-haul vehicle with a maximum transport capacity of Q freights. There is also an alternative transport option for each destination d at a cost of B_d per freight. This alternative option has unlimited transport capacity, but can only be used for freights whose due-day is immediate (i.e., $r = k = 0$).

3.1 Markov Model (Dynamic Programming)

In this section, we transform the problem characteristics described before into stages, states, decision variables, transitions, and the optimality equations, or Dynamic Programming (DP) recursion for the Markov model. The *stages* of the model correspond to the days of the planning horizon. Thus, we denote discrete and consecutive stages by t . At each stage t , there is a known group of freights with different characteristics. We define $F_{t,d,r,k}$ as the number of known freights at stage t whose destination is d , whose release-day is r stages after t , and whose time-window length is k (i.e., its due-day is $r + k$ stages after t). The *state* of

the system at stage t is denoted by \mathbf{S}_t and is defined as the vector of all freight variables $F_{t,d,r,k}$, as seen in (1).

$$\mathbf{S}_t = [F_{t,d,r,k}]_{\forall d \in \mathcal{D}, r \in \mathcal{R}, k \in \mathcal{K}}, \forall t \in \mathcal{T} \quad (1)$$

The main decision made at a stage is which freights to consolidate in the long-haul vehicle of that stage. At each stage t , only freights which have been released (i.e., freights with $r = 0$) can be transported. Moreover, note that only one trip is carried out per stage and that its maximum transport capacity is Q freights. We use the integer variable $x_{t,d,k}$ as the number of freights that are transported in the long-haul vehicle at stage t , which have destination d and are due k stages after t . We denote the vector of decision variables at stage t as \mathbf{x}_t . Since only freights that have been released at the current stage can be transported, the possible values of these decision variables are state dependent. We define the feasible space of the vector of decision variables \mathbf{x}_t , given a state \mathbf{S}_t , as follows:

$$\mathbf{x}_t = [x_{t,d,k}]_{\forall d \in \mathcal{D}, k \in \mathcal{K}}, \forall t \in \mathcal{T} \quad (2a)$$

s.t.

$$\sum_{d \in \mathcal{D}} \sum_{k \in \mathcal{K}} x_{t,d,k} \leq Q, \quad (2b)$$

$$0 \leq x_{t,d,k} \leq F_{t,d,0,k}, \quad |F_{t,d,0,k} \in \mathbf{S}_t \\ \forall d \in \mathcal{D}, k \in \mathcal{K} \quad (2c)$$

As mentioned earlier, four discrete and independent random variables describe the arrival of freights, and their characteristics, over time: $\{F, D, R, K\}$. We combine all these random variables into a single arrival information variable $\tilde{F}_{t,d,r,k}$ which represents the freights that arrived from outside the system between stages $t-1$ and t , whose destination is d , whose release-day is r , and whose time-window length is k . We denote the vector of arrival information variables at stage t as \mathbf{W}_t , and define it in (3).

$$\mathbf{W}_t = [\tilde{F}_{t,d,r,k}]_{\forall d \in \mathcal{D}, r \in \mathcal{R}, k \in \mathcal{K}}, \forall t \in \mathcal{T} \quad (3)$$

The consolidation decision \mathbf{x}_t and arrival information \mathbf{W}_t have an influence on the transition of the state at stage $t-1$ to the state at stage t . Besides these two factors, we note that release-day r and due-day $r+k$ are indexed relative to stage t and therefore also have an influence on the transition of the freight variables $F_{t,d,r,k}$. To represent all of these transition factors and relations, we introduce the transition function S^M , as seen in (4a). In this function, we define freight variables $F_{t,d,r,k}$ at stage t with destination d according to their release-days and time-window length in three ways. First, freights which have been released at stage t (i.e., $r = 0$) and have a time-window length of k are the result of: (i) freights from the previous stage $t-1$ which were already released, had time-window length $k+1$, and were not transported (i.e., $F_{t-1,d,0,k+1} -$

$x_{t-1,d,k+1}$), (ii) freights from the previous stage $t - 1$ with next-stage release-day (i.e., $r = 1$) and time-window length k (i.e., $F_{t-1,d,1,k}$), and (iii) the new (random) arriving freights with the same characteristics (i.e., $\tilde{F}_{t,d,0,k}$) as seen in (4b). Second, freights which have not been released at stage t (i.e., $r \geq 1$) are the result of: (i) freights from the previous stage $t - 1$ with a release-day $r + 1$ and that have the same time-window length k , and (ii) the new freights with the same characteristics (i.e., $\tilde{F}_{t,d,r,k}$), as seen in (4c). Third, freights which have the maximum due-day (i.e., $k = K^{max}$) are the result only of the new freights with the same characteristics (i.e., $\tilde{F}_{t,d,r,|K|}$), as seen in (4d).

$$\mathbf{S}_t = S^M(\mathbf{S}_{t-1}, \mathbf{x}_{t-1}, \mathbf{W}_t), \forall t \in \mathcal{T} | t > 0 \quad (4a)$$

s.t.

$$F_{t,d,0,k} = F_{t-1,d,0,k+1} - x_{t-1,d,k+1} + F_{t-1,d,1,k} + \tilde{F}_{t,d,0,k}, |k < K^{max} \quad (4b)$$

$$F_{t,d,r,k} = F_{t-1,d,r+1,k} + \tilde{F}_{t,d,r,k}, |r \geq 1 \quad (4c)$$

$$F_{t,d,r,K^{max}} = \tilde{F}_{t,d,r,K^{max}}, \quad (4d)$$

$$\forall d \in \mathcal{D}, r \in \mathcal{R}, r + 1 \in \mathcal{R}, k \in \mathcal{K}, k + 1 \in \mathcal{K}$$

Now that stages, states, decision variables, and transitions are defined, we are left only with the optimality equations or DP recursion. Before defining this recursion, note that last-mile costs $C_{\mathcal{D}'}$ depend on the subset of destinations $\mathcal{D}' \subseteq \mathcal{D}$ from the freights consolidated in the long-haul vehicle. Note also that there is an alternative transportation cost B_d per urgent freight (i.e., $r = k = 0$) to destination d that is not consolidated in the long-haul vehicle. For determining the total costs, we introduce two auxiliary variables: (i) $y_{t,d} \in \{0, 1\}$, which gets a value of 1 if any freight with destination d is consolidated in the long-haul vehicle at stage t and 0 otherwise, as seen in (5b), and (ii) $z_{t,d} \in \mathbb{Z}$ which counts how many urgent freights to destination d were not transported in the long-haul vehicle, as seen in (5c). Thus, the costs at stage t are defined as a function of the vector of decision variables \mathbf{x}_t and the state \mathbf{S}_t as seen in (5a).

$$C(\mathbf{S}_t, \mathbf{x}_t) = \sum_{\mathcal{D}' \subseteq \mathcal{D}} \left(C_{\mathcal{D}'} \cdot \prod_{d' \in \mathcal{D}'} y_{t,d'} \cdot \prod_{d'' \in \mathcal{D} \setminus \mathcal{D}'} (1 - y_{t,d''}) \right) + \sum_{d \in \mathcal{D}} (B_d \cdot z_{t,d}) \quad (5a)$$

s.t.

$$y_{t,d} = \begin{cases} 1, & \text{if } \sum_{k \in \mathcal{K}} x_{t,d,k} > 0 \\ 0, & \text{otherwise} \end{cases}, \forall d \in \mathcal{D} \quad (5b)$$

$$z_{t,d} = F_{t,d,0,0} - x_{t,d,0}, \forall d \in \mathcal{D} \quad (5c)$$

The objective of the model is to minimize the costs in (5a), under the uncertainty in the arrival of freights and their characteristics, over a finite horizon. Therefore, we need an optimal decision for each of the possible states for each stage in the horizon, or in other words a policy. We define a policy π as a function that maps each possible state \mathbf{S}_t to a decision vector \mathbf{x}_t^π . Thus, the formal

objective of the Markov model is to find the best policy $\pi \in \Pi$ which minimizes the expected costs over the planning horizon, given an initial state \mathbf{S}_0 , as seen in (6):

$$\min_{\pi \in \Pi} \mathbb{E} \left\{ \sum_{t \in \mathcal{T}} C(\mathbf{S}_t, \mathbf{x}_t^\pi) | \mathbf{S}_0 \right\} \quad (6)$$

Following Bellman's principle of optimality, the best policy π for the entire planning horizon can be found solving a set of stochastic recursive equations which consider current-stage and expected next-stage costs. The recursion between stages t and $t+1$ can be written using the arrival information vector \mathbf{W}_{t+1} and the transition function S^M . Remind that \mathbf{W}_{t+1} is the result of the discrete and finite random variables describing the arrival process of freights, and thus has also a discrete and finite number of realizations. We denote the set of all possible realizations of the arrival information vector with Ω , i.e., $\mathbf{W}_t \in \Omega, \forall t \in \mathcal{T}$. For each realization $\omega \in \Omega$, there is an associated probability p_ω^Ω . With all of this in mind, the Bellman's optimality equations are defined as seen in (7). We consider that at the end of the horizon \mathcal{T} there are no next-stage costs, i.e., $V_{T^{max}}(\mathbf{S}_{T^{max}}) = 0$.

$$\begin{aligned} V_t(\mathbf{S}_t) &= \min_{\mathbf{x}_t} (C(\mathbf{S}_t, \mathbf{x}_t) + \mathbb{E}\{V_{t+1}(\mathbf{S}_{t+1})\}), \forall t \in \mathcal{T} \\ &= \min_{\mathbf{x}_t} (C(\mathbf{S}_t, \mathbf{x}_t) + \mathbb{E}\{V_{t+1}(S^M(\mathbf{S}_t, \mathbf{x}_t, \mathbf{W}_{t+1}))\}) \\ &= \min_{\mathbf{x}_t} \left(C(\mathbf{S}_t, \mathbf{x}_t) + \sum_{\omega \in \Omega} (p_\omega^\Omega \cdot V_{t+1}(S^M(\mathbf{S}_t, \mathbf{x}_t, \omega))) \right) \end{aligned} \quad (7)$$

The probability p_ω^Ω depends on the probabilities of the four discrete and independent random variables describing the arrival process: $\{F, D, R, K\}$. Remind that a number of freights f arrive with probability p_f^F and that each freight arriving has destination d with probability p_d^D , release-day r with probability p_r^R , and a time-window length k with probability p_k^K . Since the arrival process is independent of the stage, a realization ω is the vector of all freight variables $\tilde{F}_{d,r,k}^\omega$ without an index t , as seen in (8).

$$\omega = \left[\tilde{F}_{d,r,k}^\omega \right]_{\forall d \in \mathcal{D}, r \in \mathcal{R}, k \in \mathcal{K}} \quad (8)$$

The total number of freights arriving in realization ω is f , as seen in (9b), with probability p_f^F . Since the characteristics of freights are independent of each other, the probability that $\tilde{F}_{d,r,k}^\omega$ freights will have destination d , release-day r and time-window length k is the product of the probability of each characteristic raised to the power of that number of freights, as seen in the last part of (9a). However, the probability of a realization ω is not just the product of the probability of the number of freights and the probability of the characteristics of each freight variable. In our model, the order in which freights arrive at a given stage

t does not matter, but “repetition” in freight characteristics is allowed. From a combinatorial perspective [11], there are β ways of assigning the total number of arriving freights f to each freight variable $\tilde{F}_{d,r,k}^\omega$ (i.e., each combination of characteristics), as seen in (9c). Thus, we need to multiply the aforementioned probabilities with the multinomial coefficient β . Using this information, the probability p_ω^Ω can be computed as follows:

$$p_\omega^\Omega = \beta \cdot p_f^F \cdot \prod_{d \in \mathcal{D}, r \in \mathcal{R}, k \in \mathcal{K}} \left([p_d^D \cdot p_r^R \cdot p_k^K]^{\tilde{F}_{d,r,k}^\omega} \right) \quad (9a)$$

$$\text{s.t.} \\ f = \sum_{d \in \mathcal{D}, r \in \mathcal{R}, k \in \mathcal{K}} \tilde{F}_{d,r,k}^\omega \quad (9b)$$

$$\beta = \frac{f!}{\prod_{d \in \mathcal{D}, r \in \mathcal{R}, k \in \mathcal{K}} (\tilde{F}_{d,r,k}^\omega!)} \quad (9c)$$

Finally, with all the aforementioned definitions of stages, states, decision variables, transitions, and Bellman’s equations (i.e., DP recursion), the dynamic multi-period freight consolidation problem can be solved to optimality for all possible initial states. The way to do so is by solving the DP recursion in (7), starting at $t = T^{max} - 1$ where there are no next-stage costs, and then stepping backward in the horizon, considering all states at each stage, until $t = 0$. However, as with most Markov models, our model suffers from the three curses of dimensionality mentioned by Powell [10] and possibly a fourth one. First, the set of all possible realizations ω of the arrival information contains all possible permutations of the maximum number of freights that can arrive, for each possible combination of characteristics. Second, the state space of all possible states \mathcal{S}_t contains, for each possible realization of the arrival information, all possible permutation of accumulated freights (e.g., if at most two freights arrive with due-day of today or tomorrow, it is possible to have a state where there are four freights today). Third, the decision space of all possible decisions \mathbf{x}_t contains all permutations of each freight variable $F_{t,d,0,k}$. The fourth possible curse of dimensionality may arise due to the necessity of defining costs $C_{\mathcal{D}'}$ for each subset of destinations $\mathcal{D}' \subseteq \mathcal{D}$. For these reasons, our Markov model is directly applicable only in small, toy-sized, problems. Nevertheless, it provides the foundation for solving larger problems. In the following section we explain how to overcome these impediments for realistic-size problems using this Markov model as a basis.

3.2 Approximate Dynamic Programming Solution Algorithm

Approximate Dynamic Programming (ADP) is a modeling framework, based on a Markov model, that offers several strategies for tackling the curses of dimensionality in large, multi-period, stochastic optimization problems [10]. The output of ADP is the same as in the Markov model, i.e., a policy or function

π that maps each possible state \mathbf{S}_t to a decision vector \mathbf{x}_t^π , for each stage t in the planning horizon. This policy is derived from an approximation of the optimal values of the Bellman's equations. To do this approximation, a series of constructs and algorithmic manipulations of the base Markov model are needed. In this section we present the constructs and algorithmic manipulations used in our ADP algorithm, as shown in Algorithm 1.

Algorithm 1 Approximate Dynamic Programming Solution Algorithm

Require: $F, D, R, K, \mathcal{D}, T^{max}, R^{max}, K^{max}, [C_{\mathcal{D}'}]_{\forall \mathcal{D}' \subseteq \mathcal{D}}, B_d, Q, \mathbf{S}_0, N$

Ensure: Sets $\mathcal{T}, \mathcal{R}, \mathcal{K}, \Omega$ are defined

- 1: Initialize $\bar{V}_t^0, \forall t \in \mathcal{T}$
 - 2: $n \leftarrow 1$
 - 3: **while** $n \leq N$ **do**
 - 4: $\mathbf{S}_0^n \leftarrow \mathbf{S}_0$
 - 5: **for** $t = 0$ **to** $T^{max} - 1$ **do**
 - 6: $\hat{v}_t^n \leftarrow \min_{\mathbf{x}_t^n} (C(\mathbf{S}_t^n, \mathbf{x}_t^n) + \bar{V}_t^{n-1}(S^{M,x}(\mathbf{S}_t^n, \mathbf{x}_t^n)))$
 - 7: **if** $t > 0$ **then**
 - 8: $\bar{V}_{t-1}^{n,x^*} \leftarrow U^V(\bar{V}_{t-1}^{n-1}(\mathbf{S}_{t-1}^{n,x^*}), \mathbf{S}_{t-1}^{n,x^*}, \hat{v}_t^n)$
 - 9: **end if**
 - 10: $\mathbf{x}_t^{n*} \leftarrow \arg \min_{\mathbf{x}_t^n} (C(\mathbf{S}_t^n, \mathbf{x}_t^n) + \bar{V}_t^{n-1}(S^{M,x}(\mathbf{S}_t^n, \mathbf{x}_t^n)))$
 - 11: $\mathbf{S}_t^{n,x^*} \leftarrow S^{M,x}(\mathbf{S}_t^n, \mathbf{x}_t^{n*})$
 - 12: $\mathbf{W}_t^n \leftarrow \text{RandomFrom}(\Omega)$
 - 13: $\mathbf{S}_{t+1}^n \leftarrow S^M(\mathbf{S}_t^n, \mathbf{x}_t^{n*}, \mathbf{W}_t^n)$
 - 14: **end for**
 - 15: **end while**
 - 16: **return** $[\bar{V}_t^N]_{\forall t \in \mathcal{T}}$
-

The ADP solution can be applied to realistic-size instances because two of the dimensionality issues mentioned in the previous section are completely avoided. The first dimensionality issue corresponds to the set Ω containing all possible realizations ω of the arrival information. This issue is avoided through the construct of a post-decision state $\mathbf{S}_t^{n,x}$ and an approximated next-stage cost $\bar{V}_t^n(\mathbf{S}_t^{n,x})$, which we explain in the next paragraphs. The second dimensionality issue corresponds to the state space which contains all possible permutations of accumulated freights for each possible realization of the arrival information. This issue is avoided through the so-called ‘‘forward dynamic programming’’ algorithmic strategy, which solves the Bellman's equations by stepping forward in time, and repeats this process for N iterations. We elaborate on this strategy later in this section.

A post-decision state $\mathbf{S}_t^{n,x}$ is the state directly after decision \mathbf{x}_t^n given state \mathbf{S}_t^n but before the arrival information \mathbf{W}_t^n is known. In our model, the post-decision state contains all post-decision freight variables $F_{t,d,r,k}^{n,x}$, as seen in (10). To define the values of the post-decision state vector, we use the transition function $S^{M,x}$, as seen in (11). This function works in the same way of the

DP transition function defined in (4a), with the difference that the new arrival information $\mathbf{W}_t^n = \left[\tilde{F}_{t,d,r,k}^n \right]_{\forall d \in \mathcal{D}, r \in \mathcal{R}, k \in \mathcal{K}}$ is not included.

$$\mathbf{S}_t^{n,x} = \left[F_{t,d,r,k}^{n,x} \right]_{\forall d \in \mathcal{D}, r \in \mathcal{R}, k \in \mathcal{K}}, \quad \forall t \in \mathcal{T} \quad (10)$$

$$\mathbf{S}_t^{n,x} = S^{M,x}(\mathbf{S}_t^n, \mathbf{x}_t^n), \quad \forall t \in \mathcal{T} \quad (11)$$

is defined as

$$\begin{aligned} F_{t+1,d,0,k}^{n,x} &= F_{t,d,0,k+1}^n - x_{t,d,k+1}^n + F_{t,d,1,k}^n, \\ F_{t+1,d,r,k}^{n,x} &= F_{t,d,r+1,k}^n \mid r \geq 1, \\ \forall d \in \mathcal{D}, r \in \mathcal{R}, r+1 \in \mathcal{R}, k \in \mathcal{K}, k+1 \in \mathcal{K} \end{aligned}$$

In the forward dynamic programming algorithmic strategy, the Bellman's equations are solved only for one state at each stage. Just as in the Markov model, the feasible decisions \mathbf{x}_t^n for state \mathbf{S}_t^n in these equations are defined by (2a). However, some modifications are necessary to apply this algorithmic strategy. Besides the construct of the post-decision state, the construct of an approximated next-stage cost $\bar{V}_t^n(\mathbf{S}_t^{n,x})$ is necessary. This construct replaces the standard expectation in Bellman's equations, as seen in (12).

$$\bar{V}_t^n(\mathbf{S}_t^{n,x}) = \mathbb{E} \{ V_{t+1}(\mathbf{S}_{t+1}) \mid \mathbf{S}_t^x \} \quad (12)$$

Using the post-decision state and the approximated next-stage cost, the original Bellman's equations from (7) are converted to the ADP forward optimality equations, as seen in (13). Note that for each feasible decision \mathbf{x}_t^n , there is an associated post-decision state $\mathbf{S}_t^{n,x}$ obtained using (11). The ADP forward optimality equations are solved first at stage $t = 0$ and \mathbf{S}_0 , and then for subsequent stages and states until the end of the horizon. To advance "forward" in time, from stage t to $t + 1$, a Monte Carlo simulation of the random information Ω , defined in (8), is done. In this simulation, a sample \mathbf{W}_t^n from Ω is obtained. With this information, transition in the algorithm is done using the same DP transition function defined in (4a), as seen in Algorithm 1 lines 12 and 13.

$$\begin{aligned} \hat{v}_t^n &= \min_{\mathbf{x}_t^n} (C(\mathbf{S}_t^n, \mathbf{x}_t^n) + \bar{V}_t^{n-1}(\mathbf{S}_t^{n,x})) \\ &= \min_{\mathbf{x}_t^n} (C(\mathbf{S}_t^n, \mathbf{x}_t^n) + \bar{V}_t^{n-1}(S^{M,x}(\mathbf{S}_t^n, \mathbf{x}_t^n))) \end{aligned} \quad (13)$$

Immediately after the forward optimality equations are solved, the approximated next-stage cost $\bar{V}_t^n(\mathbf{S}_t^{n,x})$ is updated retrospectively, as seen in (14). The rationale behind this update is that, at stage t , the algorithm has seen new arrival information (via the Monte Carlo simulation) and has taken a decision in the new state \mathbf{S}_t^n which incurs a cost. This means that the approximated next-stage cost that was calculated at the previous stage $t - 1$, i.e., $\bar{V}_{t-1}^{n-1}(\mathbf{S}_{t-1}^{n,x})$, has now been observed at stage t . To take advantage of this observation and improve the

approximation, the algorithm updates this approximated next-stage cost using the old approximation, i.e., $\bar{V}_{t-1}^{n-1}(\mathbf{S}_{t-1}^{n,x})$, the new approximation, i.e., the value \hat{v}_t^n corresponding to the optimal decision that solves (13), and the decision \mathbf{x}_t^n that resulted in the value \hat{v}_t^n . We use U^V to denote the process that takes all of the aforementioned parameters and “tunes” the approximating function, as seen in (14). Note that in Algorithm 1 line 8, the parameters used for the update have the superscript * indicating the optimal decision made at stage $t - 1$ and its corresponding post-decision state.

$$\bar{V}_{t-1}^n(\mathbf{S}_{t-1}^{n,x}) \leftarrow U^V(\bar{V}_{t-1}^{n-1}(\mathbf{S}_{t-1}^{n,x}), \mathbf{S}_{t-1}^{n,x}, \hat{v}_t^n), \forall t \in \mathcal{T} \quad (14)$$

Two of the largest challenges of ADP are: (i) to find an accurate approximation function $\bar{V}_t^n(\mathbf{S}_t^{n,x})$ of the value of a post-decision state $\mathbf{S}_t^{n,x}$, and (ii) to define an appropriate updating process U^V for this function. For our problem, we use the concept of post-decision state “features”. A feature of a post-decision state is a quantitative characteristic that explains, to some extent, what the value of that post-decision state is. In our problem, features such as the number of urgent freights, the number of released freights that are not urgent, and the number of freights which have not been released for transport, can explain part of the value of a post-decision state. We define a set of features \mathcal{A} for which the value of each feature $a \in \mathcal{A}$ is obtained using a function $\phi_a(\mathbf{S}_t^{n,x})$. We assume the approximated next-stage value of a post-decision state can be expressed by a weighted linear combination of the features, using the weights θ_a for each feature $a \in \mathcal{A}$, as seen in (15).

$$\bar{V}_t^n(\mathbf{S}_t^{n,x}) = \sum_{a \in \mathcal{A}} (\phi_a(\mathbf{S}_t^{n,x}) \cdot \theta_a) \quad (15)$$

The use of features and weights for the approximating the value function $\bar{V}_t^n(\mathbf{S}_t^{n,x})$ is comparable to the use of regression models for fitting data to a (linear) function. In that sense, the independent variables of the regression would be the post-decision features and the dependent variable would be the post-decision value. However, in contrast to regression models, the data in our ADP is generated iteratively inside an algorithm and not all at once. Therefore, the updating process U^V for the approximating function in (15) cannot be based on solving systems of equations as in traditional regression models. Instead, we use a recursive least squares method for nonstationary data to “fine-tune” the weight θ_a for each feature $a \in \mathcal{A}$. This method is based in regression models, and explained in detail in [10].

4 Numerical Experiments

Through numerical experiments we want first, to show the accuracy of the approximation method, and second, to compare the benefits of our proposed Markov model with a heuristic commonly used in practice. To do these experiments, we use two test instances: (i) a small instance with three destinations,

and (ii) a large instance with seven destinations. The probability distributions of the four random variables describing the arrival process can be seen in Table 1. Note that for the Large Instance, the number of possible states is approximated by the combinatorial expression $\sum_{i=1}^{F^{max} \cdot (R^{max} + K^{max} + 1)} \frac{(|\mathcal{D}| \cdot |\mathcal{R}| \cdot |\mathcal{K}| + i - 1)!}{i! \cdot (|\mathcal{D}| \cdot |\mathcal{R}| \cdot |\mathcal{K}|)!}$.

Table 1. Random variables in the numerical experiments

Input Parameter	Small Instance	Large Instance
Freights arriving per day (F)	{1, 2}	{1, 2, 3, 4}
→Probability (p_f^F)	{0.8, 0.2}	{0.25, 0.25, 0.25, 0.25}
Destinations (D)	{1, 2, 3}	{1, 2, 3, 4, 5, 6, 7}
→Probability (p_d^D)	{0.1, 0.8, 0.1}	{0.1, 0.2, 0.1, 0.1, 0.3, 0.1, 0.1}
Release-days (R)	{0}	{0, 1, 2}
→Probability (p_r^R)	{1}	{0.3, 0.3, 0.4}
Time-window lengths (K)	{0, 1, 2}	{0, 1, 2}
→Probability (p_k^K)	{0.2, 0.3, 0.5}	{0.2, 0.3, 0.5}
# Possible arrival realizations $ \Omega $	54	766479
# Possible states	2884	$\approx 8.18 \cdot 10^{18}$

In the first experiment, we test the accuracy of the approximation method using the Small Instance. To do this, we compare the value of several initial states \mathbf{S}_0 of the Markov model in Section 3.1 against the value for the same states obtained with the ADP algorithm in Section 3.2. The remaining input parameters of the Small Instance are defined as follows. The planning horizon is $T^{max} = 5$ and the long-haul vehicle capacity is $Q = 3$. The long-haul vehicle cost $C_{\mathcal{D}'}$ for a subset of destinations $\mathcal{D}' \subseteq \mathcal{D}$ is defined between 250 and 1000, such that larger subsets have higher costs than smaller ones. The alternative cost B_d is defined between 500 and 1000 per destination $d \in \mathcal{D}$. The parameters of the ADP algorithm are set as follows. The number of iterations is $N = 2000$. The features \mathcal{A} are related to three characteristics of a post-decision state: (i) the number of freights with each combination of freight characteristics, (ii) the total number of urgent freights (i.e., $r = k = 0$), and (iii) the total number of released and non-urgent freights (i.e., $r = 0$ and $k > 0$). The features related to these characteristics include also counting the number of destinations that fulfill such characteristics (e.g., destinations having urgent freights). We also include a constant feature a' such that $\phi_{a'}(\mathbf{S}_t^{n,x}) = 1$ for all post-decision states and stages. Feature weights are initialized with 1, i.e. $\theta_a = 1$ for all features and all stages at iteration $n = 1$. The updating process U^V (i.e., the recursive least squares method for nonstationary data) requires a discount factor λ which is defined as $\lambda = 1 - \frac{0.5}{n}$. The results of this comparison can be seen in Figure 1. The state number counts the number of urgent freights contained in each state, i.e. State 3 has 3 urgent freights.

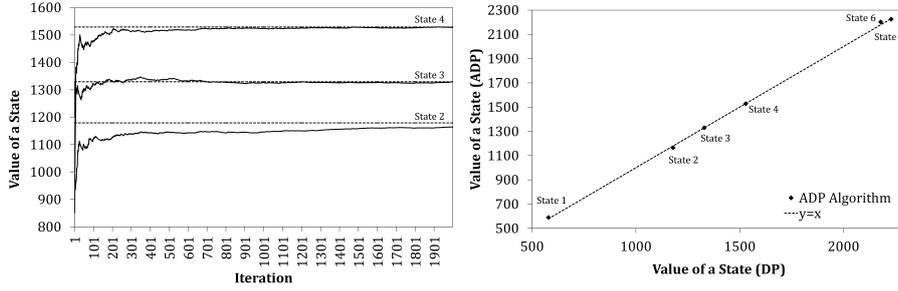


Fig. 1. Convergence and Approximation Accuracy of the ADP Algorithm

In the sample of states above, we can already observe several characteristics of the output from the ADP algorithm with the aforementioned settings. In the left part of Figure 1 we see that the values fluctuate during the first iterations for some states more than for others, but eventually converge for all states. In the right part of Figure 1, we see that the difference between the optimal values and the estimates produced by ADP are small. These experiments show that the ADP algorithm accurately predicts the value of initial states in the Markov model. However, these ADP values are an intermediate result since they are used to define a policy that should dictate decisions for all possible states during the horizon. Thus, to properly compare performance, the entire policy of the ADP should be compared to that of the Markov model.

In the second experiment, we compare the policy resulting from our ADP algorithm with (i) the policy of the Markov model (for the Small Instance only), and (ii) a heuristic commonly used in practice (for the Small and Large Instance). This experiment shows the benefits of incorporating stochastic information in the dynamic multi-period freight consolidation problem compared to using a heuristic. This benchmark heuristic consolidates the freights that yield the lowest direct costs (i.e., no future costs considered), and then if there is capacity left, fills the long-haul vehicle with released freights that go to the same destinations of the freights already consolidated (i.e., no extra costs). In the Large Instance, the planning horizon is $T^{max} = 5$ and the long-haul vehicle capacity is $Q = 10$. The long-haul vehicle cost $C_{\mathcal{D}'}$ for a subset of destinations $\mathcal{D}' \subseteq \mathcal{D}$ is defined between 250 and 2050, such that larger subsets have higher costs than smaller ones. The alternative cost B_d is defined between 300 and 800 per destination $d \in \mathcal{D}$. The ADP settings are the same as for the Small Instance. The experiment is done in a simulation of 2000 runs, using common random numbers for the comparison of the arrival information in the planning horizon. The results about the performance of the different policies are seen in Figure 2. Once again, the state number counts the number of urgent freights contained in each initial state.

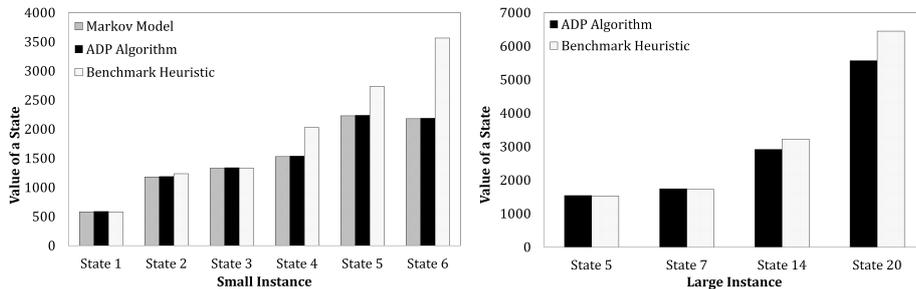


Fig. 2. Performance of the ADP Algorithm in the Small and Large Instance

From Figure 2 we conclude that the ADP policy performs at least as good as the benchmark heuristic, and always better than the heuristic in initial states with a large number of urgent freights. Furthermore, in the left part of Figure 2 we see that the performance of the ADP policy is approximately the same as the one from the Markov model, meaning that the ADP estimates result in a near-optimal policy.

5 Conclusions

We developed a Markov model and an Approximate Dynamic Programming (ADP) solution for the dynamic multi-period freight consolidation problem. The approach is designed to achieve the optimal balance between freights that are consolidated in a long-haul vehicle and freights that are postponed for further trips or alternative transportation modes. The optimal balance is achieved taking into account the probabilistic knowledge in the arrival of freights and their characteristics, the applicable transportation restrictions, and the interdependence of decisions over time.

Through a limited number of numerical experiments, proofs-of-concept of the accuracy of the ADP method and of the benefits of the Markov model were shown. These experiments showed that, even in small instances, there are some states where it pays off to have a look-ahead policy, and some others where a common heuristic is sufficient to achieve the optimal balance between direct shipment and postponement for future consolidation. This leads to the idea that further research is needed to identify in which problem settings looking ahead in the future (i.e., using the ADP approach) yield the largest benefits. Specifically, more experiments on large problem instances and different benchmark policies are crucial to analyze the value of look-ahead policies.

Bibliography

- [1] Andersen, J., Christiansen, M., Crainic, T.G., Gronhaug, R.: Branch and price for service network design with asset management constraints. *Transportation Science* 45(1), 33–49 (2011)
- [2] Andersen, J., Crainic, T.G., Christiansen, M.: Service network design with asset management: Formulations and comparative analyses. *Transportation Research Part C: Emerging Technologies* 17(2), 197 – 207 (2009), selected papers from the Sixth Triennial Symposium on Transportation Analysis (TRISTAN VI)
- [3] Andersen, J., Crainic, T.G., Christiansen, M.: Service network design with management and coordination of multiple fleets. *European Journal of Operational Research* 193(2), 377 – 389 (2009)
- [4] Crainic, T.G., Gendreau, M., Farvolden, J.M.: A simplex-based tabu search method for capacitated network design. *INFORMS Journal on Computing* 12(3), 223–236 (2000)
- [5] Crainic, T.G., Kim, K.H.: Chapter 8 Intermodal Transportation. In: Barnhart, C., Laporte, G. (eds.) *Transportation, Handbooks in Operations Research and Management Science*, vol. 14, pp. 467 – 537. Elsevier (2007)
- [6] Hoff, A., Lium, A.G., Lokketangen, A., Crainic, T.: A metaheuristic for stochastic service network design. *Journal of Heuristics* 16(5), 653–679 (2010)
- [7] Kim, D., Barnhart, C.: *Transportation service network design: Models and algorithms*. Springer (1999)
- [8] Lium, A.G., Crainic, T.G., Wallace, S.W.: A study of demand stochasticity in service network design. *Transportation Science* 43(2), 144–157 (2009)
- [9] Moccia, L., Cordeau, J.F., Laporte, G., Ropke, S., Valentini, M.P.: Modeling and solving a multimodal transportation problem with flexible-time and scheduled services. *Networks* 57(1), 53–68 (2011)
- [10] Powell, W.B.: *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, vol. 1. John Wiley & Sons (2007)
- [11] Riordan, J.: *Introduction to combinatorial analysis*. Courier Dover Publications (2002)
- [12] SteadieSeifi, M., Dellaert, N., Nuijten, W., Woensel, T.V., Raoufi, R.: Multimodal freight transportation planning: A literature review. *European Journal of Operational Research* 233(1), 1 – 15 (2014)
- [13] Verma, M., Verter, V., Zufferey, N.: A bi-objective model for planning and managing rail-truck intermodal transportation of hazardous materials. *Transportation Research Part E: Logistics and Transportation Review* 48(1), 132 – 149 (2012), select Papers from the 19th International Symposium on Transportation and Traffic Theory
- [14] Wieberneit, N.: Service network design for freight transportation: a review. *OR Spectrum* 30(1), 77–112 (2008)