

An approximate dynamic programming approach to urban freight distribution with batch arrivals

Wouter van Heeswijk, Martijn Mes, Marco Schutten

Beta Working Paper series 474

BETA publicatie	WP 474 (working paper)
ISBN	
ISSN	
NUR	804
Eindhoven	April 2015

An approximate dynamic programming approach to urban freight distribution with batch arrivals

Wouter van Heeswijk, Martijn Mes and Marco Schutten

University of Twente

Department of Industrial Engineering and Business Information Systems

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

{w.j.a.vanheeswijk,m.r.k.mes,m.schutten}@utwente.nl

Abstract. We study a dispatch problem with uncontrolled batch arrivals of LTL orders at an urban consolidation center. These arrivals reflect the delivery of goods by independent carriers. The specific order properties (e.g., destination, size, delivery window) may be highly varying in city logistics, and directly distributing an incoming batch may yield high costs. Instead, the hub operator may decide to wait for incoming batches that allow for more efficient distribution. A waiting policy is required to decide which orders to ship and which orders to hold. We model the dispatching problem as a Markov decision problem. Dynamic Programming (DP) is applied to solve toy-sized instances. To solve realistic instances, we propose an Approximate Dynamic Programming (ADP) approach. Through numerical experiments, we show that the ADP approach closely approximates the optimal values of DP for small instances, and outperforms two benchmark policies for larger instances.

Keywords: urban distribution, transportation planning, consolidation, approximate dynamic programming

1 Introduction

In the field of urban freight logistics, the need for consolidation centers at the edge of urban areas is becoming increasingly important [10]. Due to the external costs of freight transport – such as congestion, air pollution, and noise hindrance – there is a desire for higher efficiency of goods transport within the city center. Governments actively seek to reduce the negative influence of large trucks in urban areas. Imminent regulations are, for example, restricted access areas and road pricing for heavy vehicles. Such developments spur the need for transshipments at the edge of urban areas. Transshipments allow both for bundling goods – such that more efficient routes can be generated – and dispatching environment-friendly vehicles such as electric vehicles on the last mile. However, the introduction of an additional transshipment in the supply chain also poses new challenges. We adopt the perspective of the party in charge of the consolidation center, to which we refer as the ‘hub operator’. We focus on

the timing of dispatching urban freighters in an environment where the batch arrivals of goods at the hub are not controlled by the hub operator.

In our problem settings, orders arrive at the consolidation center in batches (e.g., a trailer load), and subsequently need to be distributed to the customers in the city. A batch may contain orders with dispersed destinations, various delivery windows, distinct load sizes, etc. Directly distributing an arriving batch may therefore render poor solutions. Instead, waiting for other batches to arrive could yield order clusters for which better solutions are available. However, various uncertainties may play a role, such as when a new batch arrives and what the exact contents of such a batch are. Though faced with uncertainty, the operator may have a certain degree of knowledge regarding orders that may arrive in the future, such as forecasts or probability distributions. Based on the available knowledge regarding current orders and expectations on orders yet to arrive, the operator is able to make informed waiting decisions. This can be accomplished by deploying a waiting policy that makes dispatch decisions given the information and beliefs of the operator. Three basic types of waiting policies can be distinguished [2]. *Time-based* policies consider only the time until delivery as basis for dispatching decisions. *Quantity-based* policies base the decision whether or not to depart on the achieved fill rate of the vehicle (e.g., when it can make a tour utilizing at least 70% of its capacity). *Time-and-quantity* based approaches take both factors into account. Time-based policies generally result in better service to the customer in terms of lateness, while quantity-based policies perform better in terms of capacity utilization [3]. In practice, policies are generally time-and-quantity based [2]. Furthermore, a distinction can be made between myopic policies and policies that explicitly take into account information regarding future demand, the latter vastly increasing complexity [5, 8]. The use of information regarding future events in transportation is recognized as an important aspect of optimization, yet its incorporation in solution methods is still an ongoing development [9, 5, 12].

Our dispatch problem can be classified as a Delivery Dispatch Problem (DDP). The concept of queuing with bulk service – where customers arrive independently and are served in batches – was introduced by Bailey [1]. Neuts [7] expanded on the subject by explicitly studying control policies for the timing of vehicle dispatch. Traditional solutions to the DDP rely on analytical approaches such as queuing models. Route properties are considered to be a given input [6], such that routing decisions do not play a role. A more generic approach based on a batch Markovian arrival process is presented in [2], allowing to define arrival properties that do not necessarily follow well-defined distributions. Wang and Odoni [14] study a problem with batch arrivals of passengers at a train station, who are subsequently transported to their final destinations. They present a queuing approach aimed at the design of such transportation systems. To estimate the routing costs, they apply the approximation of Daganzo [4]. Using this approach, they find an average performance gap of 2% with the solutions of well-known routing heuristics. A key distinction between our DDP and the problem studied by Wang and Odoni is that they do not consider delivery windows;

their objective is to minimize service time for the passengers. Transportation requests are planned directly; the opportunity to defer dispatch decisions to future decision moments is not included in their approach. When considering more complex variants of the DDP, the stochastic and dynamic nature of the problem gives rise to Markov decision problems [6]. Although a Markov decision model is a useful framework to describe finite-horizon decision models with stochastic elements, practical implementations generally suffer from intractably large state spaces and expected values that cannot be calculated exactly [8]. The DDP is therefore too hard to solve exactly in most instances [6].

To model the stochastic and dynamic nature of our order arrival process, we formulate our DDP as a Markov decision process. Dynamic programming (DP) can be used to solve such models to optimality, but as stated in the previous paragraph, the sheer size of such problems prevents us from applying DP on larger instances. Mathematical programming has traditionally been applied to handle high-dimensional problems in transportation. However, this method generally does not cope well with stochastic information revealed over time [9]. Topaloglu and Powell [13] present a generic framework for solving dynamic resource-allocation problems. In this type of problems, the arrival of tasks is probabilistic, and coupling tasks to resources yields a reward. The methodology they apply is known as approximate dynamic programming (ADP). This simulation-based approach uses intelligent sampling to approximate the optimal solutions of DP [8]. ADP allows to retain the complexities embedded in the dispatch problem, while circumventing the computational effort required to solve the Markov model with stochastic dynamic programming. Various successful ADP applications can be found in transportation literature [12]. Provided that the value functions of the defined problem can be accurately approximated, ADP is a powerful approach to solve realistic transportation problems. Following frameworks such as [9, 13, 8], we develop an ADP approach to solve our dispatch problem. We aim to contribute to existing literature with (i) the formulation of a Markov model for DDPs with time windows, and (ii) an ADP approach to provide high-quality solutions for realistic-sized DDPs.

2 Problem Formulation

This section introduces the planning problem. We describe the problem in a generic way, making it applicable to a variety of instances. We assume that the characteristics of arriving orders are stochastic and have a known associated probability distribution. Certain variables that are treated as stochastic in our description may be known in practice. In that case, we can replace the stochastic variable by the actual information regarding future orders, creating a restricted instance of our problem. We consider a finite planning horizon, during which batches of orders can arrive at the consolidation center. Dispatching decisions are made at fixed decision moments within the planning horizon, with constant time intervals separating the decision moments. We model the arrival rates of order batches, the number of orders in a batch, order sizes, order destinations,

and dispatch times as stochastic variables. When a batch of orders arrives at the center, the exact properties of the orders are revealed.

The decision problem that we address is the choice which orders to dispatch at the current decision moment. To make an informed decision, we require insight in the effects of postponing the dispatch of orders, e.g., distributing them from the center to their destination in the city. Postponed orders may be combined with future orders against lower costs than when dispatched at the current decision moment. We therefore consider optimization over a planning horizon; all orders arriving after the current decision moment are probabilistic. We assess all possible realizations of future order arrivals as scenarios, and plan the future arrivals as if they were actual orders. For both the actual and future orders belonging to a given scenario, we compute the costs of dispatching. The costs of dispatching probabilistic orders are required to quantify the expected costs. By assessing the realizations of probabilistic orders, we can compute the expected costs of the various dispatch decisions for the orders currently in inventory.

Consider an urban area with a fixed set of recipients (i.e., order destinations), which are delivered via a consolidation center at the edge of the area. Our representation of the urban distribution network is as follows. Let $\mathcal{G} = \{\mathcal{V}, \mathcal{A}\}$ be a directed and complete graph with \mathcal{V} being the set of vertices and \mathcal{A} being the set of arcs. The element $\{0\} \in \mathcal{V}$ represents the consolidation center in the network. The remaining vertices signify the subset of order destinations $\mathcal{V}' = \{1, 2, \dots, |\mathcal{V}'|\}$, with $\mathcal{V}' = \mathcal{V} \setminus \{0\}$. A distance matrix gives the distance between any pair of vertices in the graph.

We consider a planning horizon that contains decision moments with fixed intermediate time intervals. For instance, we may have a one-day horizon with hourly decision moments. Let $\mathcal{T} = \{0, 1, \dots, T\}$ be the set containing all decision moments, and $t \in \mathcal{T}$ describe any decision moment within the planning horizon. Furthermore, we consider a homogeneous fleet (i.e., a set of identical vehicles), though our method is able to cope with heterogeneous fleets as well. We distinguish between a primary fleet Q^{pr} and a secondary fleet Q^{se} . We assume that the secondary fleet has an infinite size, and is either an actual transport alternative (e.g., renting an additional vehicle in case of shortage) or a dummy fleet with infinite costs. A dummy fleet serves as bound on capacity, without having to explicitly calculate the capacity constraints for each decision. We only assign vehicles in Q^{se} if no more vehicles in Q^{pr} are available. We assume that dispatching a secondary vehicle is always more expensive than dispatching a primary vehicle. To ease the presentation, we assume that every dispatched vehicle has a fixed route duration of $\tau_{route} \geq 1$ (this assumption can be easily relaxed). When dispatching at t , the vehicle will be available again at $t + \tau_{route}$. For decision-making purposes, we keep track of the availability of primary vehicles now and in the future. Let $r \in [0, \tau_{route} - 1]$ be the number of time intervals before a dispatched vehicle returns. Because all vehicle are available again at $t + \tau_{route}$, we only keep track of availability up to $t + \tau_{route} - 1$. Let $q_{t,r}$ denote the number of primary vehicles available for dispatch at $t + r$. It follows that

$q_{t,0}$ vehicles are available for dispatch at t . We record primary fleet availability in the vector $Q_t = (q_{t,0}, q_{t,1}, \dots, q_{t,t+\tau_{route}-1})$.

Every order is characterized by four properties: destination, load size, earliest dispatch time, and latest dispatch time. The order destination (i.e., the customer location) is represented by a vertex $v \in \mathcal{V}'$. Let $\mathcal{L} = \{\frac{1}{k}, \frac{2}{k}, \dots, 1\}$ be the discretized set of viable load sizes, with integer $k \geq 1$, and 1 representing a full load for an urban vehicle. The size of an order is given by $l \in \mathcal{L}$. The hard dispatch window of an order is given by earliest dispatch time $e \in \mathcal{E}$ and latest dispatch time $d \in \mathcal{D}$. Both indices are relative to the decision moment t ; at the decision moment $t + 1$ all indices of orders in inventory are reduced by 1. Note that order types with $e > 0$ can be used to describe future orders. We define a maximum length of the dispatch window τ_{window} , such that $d \in [e, e + \tau_{window}]$. Every unique combination of the four properties represents an order type. Let $I_{t,v,l,e,d} \in \mathbb{Z}_+$ be the number of a given order type in inventory at t . We denote the information regarding all orders in inventory at t as $I_t = (I_{t,v,l,e,d})_{v \in \mathcal{V}', l \in \mathcal{L}, e \in \mathcal{E}, d \in \mathcal{D}}$. The state of the system at t , $S_t \in \mathcal{S}$, combines primary fleet availability with available orders, and is represented by

$$S_t = (Q_t, I_t)_{\forall v \in \mathcal{V}', l \in \mathcal{L}, e \in \mathcal{E}, d \in \mathcal{D}}, \forall t \in \mathcal{T}. \quad (1)$$

For $t \geq 1$, let $\mathcal{O}_t = \{0, 1, \dots, o_t^{max}\}$ be the set containing the number of possible order arrivals between decision moments $t - 1$ and t . Let $o_t \in \mathcal{O}_t$ be a realization of the number of orders arriving between $t - 1$ and t . Furthermore, we set $l^{max} \in \mathbb{Z}_+$ as the maximum number of orders that can be held in inventory, i.e., the maximum inventory remaining after a decision.

For every decision moment t in the planning horizon, we decide which orders in inventory to dispatch. Orders that are not dispatched remain in inventory, and are available at the next decision moment. Let the integer variable $x_{t,v,l,e,d}$ describe the number of a specific order type to be dispatched at t . A feasible action at decision moment t is given by

$$x_t(S_t) = (x_{t,v,l,e,d})_{\forall v \in \mathcal{V}', l \in \mathcal{L}, e \in \mathcal{E}, d \in \mathcal{D}}, \quad (2)$$

where

$$\sum_{v \in \mathcal{V}', l \in \mathcal{L}, e \in \mathcal{E}, d \in \mathcal{D}} (I_{t,v,l,e,d} - x_{t,v,l,e,d}) \leq l^{max}, \quad (3)$$

$$x_{t,v,l,e,d} \leq I_{t,v,l,e,d}, \forall v \in \mathcal{V}', \forall l \in \mathcal{L}, \forall e \in \mathcal{E}, \forall d \in \mathcal{D}, \quad (4)$$

$$x_{t,v,l,e,0} = I_{t,v,l,e,0}, \forall v \in \mathcal{V}', \forall l \in \mathcal{L}, \forall e \in \mathcal{E}, \quad (5)$$

$$x_{t,v,l,e,d} = 0, e > 0, \forall v \in \mathcal{V}', \forall l \in \mathcal{L}, \forall d \in \mathcal{D}, \quad (6)$$

$$x_{t,v,l,e,d} \in \mathbb{Z}_+, \forall v \in \mathcal{V}', \forall l \in \mathcal{L}, \forall e \in \mathcal{E}, \forall d \in \mathcal{D}. \quad (7)$$

Constraint (3) ensures that after dispatching, no more than the maximum inventory remains at the consolidation center. According to (4), it is not possible

to dispatch more orders of a certain type than available at the decision moment t . Constraint (5) states that all orders that have a latest dispatch time equal to the maximum delay must be dispatched. (6) prevents orders with an earliest dispatch time in the future from being dispatched. (7) states that only nonnegative integer amounts of orders can be dispatched. The set of feasible actions in a given state is described by $\mathcal{X}_t(S_t)$.

3 Markov Model

We model the operator's decision problem as a Markov model. This model considers all possible realizations of orders arrivals during the planning horizon. With this knowledge, we can make the optimal dispatch decision for the current decision moment. In realistic instances, the state space, action space, and outcome space for such a model will be intractably large. Solving the Markov model in an exact manner is therefore not possible within a reasonable time. In Section 5, we solve some toy-sized instances of the Markov model using DP. The ADP approach as outlined in Section 4 is applied to larger instances.

Every action $x_t(S_t)$ has associated direct costs $C(S_t, x_t)$. The direct costs are the sum of fixed dispatching costs per vehicle, variable transportation costs, and handling costs. As the focus of this paper is on the dispatching decisions, we do not explicitly consider routing. Instead, we use the classic procedure of [4] to estimate the transportation costs for a dispatched set of orders. This procedure is known to provide good estimates of total route distances [11], given constraints on vehicle capacity, number of destinations, and shape of the service area. These constraints are likely to be fulfilled in an urban distribution setting. The approximation is based on the average distances between the depot and the customers, the number of customer locations visited, the size of the service area, and the capacity of the vehicles. We consider fixed handling costs per visited customer; note that this provides an incentive to simultaneously deliver multiple orders to a customer.

To model the uncertainties with respect to the properties of arriving orders, we introduce six stochastic variables. These are (i) the number of orders arriving O_t , (ii) the destination V , (iii) the order size L , (iv) the earliest dispatch time E , (v) the length of the dispatch window D_{window} , and (vi) the latest dispatch time $D = E + D_{window}$. The corresponding probability distributions are discrete and finite. To capture all probability distributions into a single variable, we define the exogenous information variable $\tilde{I}_{t,v,l,e,d} \in \mathbb{Z}_+$, $t \geq 1$, which indicates the number of arrivals of a specific order type. Furthermore, we introduce a generic variable W_t that describes all exogenous information, i.e., all orders arriving between $t-1$ and t :

$$W_t = [\tilde{I}_{t,v,l,e,d}]_{\forall v \in \mathcal{V}, \forall l \in \mathcal{L}, \forall e \in \mathcal{E}, \forall d \in \mathcal{D}}, t \geq 1. \quad (8)$$

There exists a finite number of realizations of W_t . Let Ω_t be the set of possible batch arrivals between $t-1$ and t , and $\omega_t \in \Omega_t$ be a realization of the random variables occurring with $P(W_t = \omega_t)$.

We proceed to describe the transition from a state S_t to the next state S_{t+1} . The transition is affected by the action x_t and the new arrivals W_{t+1} . We first describe the effects of x_t . Orders not dispatched at t remain in inventory, hence must be included in S_{t+1} . As indices e and d are adjusted over time, we introduce two new variables to properly process the conversion. Let $e' = \max\{0, e - 1\}$ and $d' = d - 1$. Since $e < 0$ does not affect our decision making, capping e' at 0 reduces the number of possible order types. Let $\bar{q}_t \in \{0, \dots, |\mathcal{Q}^{pr}|\}$ be the number of primary vehicles dispatched at t ; combined with Q_t this information suffices to compute Q_{t+1} . We represent new arrivals with the information variable W_{t+1} . This gives us the transition function

$$S_{t+1} = S^M(S_t, x_t, W_{t+1}) , \quad (9)$$

where

$$I_{t+1,v,l,e',d'} = I_{t,v,l,e,d} - x_{t,v,l,e,d} + \tilde{I}_{t+1,v,l,e',d'} , \quad (10)$$

$$\forall t \in \mathcal{T}, \forall v \in \mathcal{V}', \forall l \in \mathcal{L}, \forall e \in \mathcal{E}, \forall d \in \mathcal{D} ,$$

$$q_{t+1,r} = \begin{cases} q_{t,r+1} - \bar{q}_t & \text{if } r < \tau_{route} - 1 , \\ |\mathcal{Q}^{pr}| & \text{if } r = \tau_{route} - 1 , \end{cases} \forall r \in [0, \tau_{route} - 1] . \quad (11)$$

Constraint (10) states that for every order type, we have the amount of the order type in state S_t , minus the amount of the order type that was dispatched, plus the amount of the order type that arrived between t and $t + 1$. Constraint (11) ensures that Q_{t+1} is consistently updated. Having described the transition function, we now introduce the optimality equation that must be solved:

$$V_t(S_t) = \min_{x_t \in \mathcal{X}_t(S_t)} \left(C(S_t, x_t) + \sum_{\omega_{t+1} \in \Omega_{t+1}} P(W_{t+1} = \omega_{t+1}) V_{t+1}(S_{t+1} | S_t, x_t, \omega_{t+1}) \right) . \quad (12)$$

We proceed to describe the state space. Between every two consecutive decision moments $t - 1$ and t , we can have $o_t \in \{0, \dots, |\mathcal{O}_t| - 1\}$ new orders arriving. Every arriving order can have any of the unique order types, given the constraints on the dispatch window. Before the new arrivals occur, we can have up to l^{max} orders in inventory. Hence, we can have at most $l^{max} + |\mathcal{O}_t| - 1$ orders at a given decision moment. A state can be any feasible combination of order types available at t , combined with any vector Q_t .

Next, we describe the action space. At every decision moment, we decide which orders to dispatch. Every combination of orders to dispatch represents a unique action. Orders that are not dispatched remain in inventory, and may be dispatched at the next decision moment. As we do not consider routing options, a unique selection of orders to dispatch equals exactly one action.

The transition from one state to another is determined by the current state, the used action, and the realization of the random variables. The remaining

inventory before new orders arrive is deterministic. The probability of o_t orders arriving is given by $P(O_t = o_t)$. The probability of an arriving order being of a certain order type is given by the multivariate distribution $P(V, L, E, D)$. V , L and E are independent random variables, while D is the sum of the realizations of E and D_{window} .

The outcome space is dependent on the state S_t , the action x_t , and the realization of new arrivals ω_{t+1} . Orders not shipped at decision moment t are with certainty included in S_{t+1} . As route duration is deterministic, so is the change in fleet availability. Therefore, only the order arrivals account for stochasticity. To account for the multiple permutations corresponding to o_t order arrivals, we must multiply the probability of o_t orders arriving with a multinomial coefficient [15]. We obtain the following probability function for new arrivals:

$$P(W_t = \omega_t) = P(O_t = o_t) \frac{o_t!}{\prod_{\tilde{I}_{t,v,l,e,d} \in \omega_t} \tilde{I}_{t,v,l,e,d}!} \prod_{v \in \mathcal{V}, l \in \mathcal{L}, e \in \mathcal{E}, d \in \mathcal{D}} P(V = v, L = l, E = e, D = d)^{\tilde{I}_{t,v,l,e,d}}, t \geq 1 . \quad (13)$$

4 Solution approach

Realistic problems become extremely large in terms of state space, action space, and outcome space, making them intractable for the DP method. We address this problem with an approximate dynamic programming approach. With ADP, we retain the full level of detail in the state description, without having to enumerate the full state space. By means of Monte Carlo simulation, we are able to approximate the exact values of the DP method [8].

In our ADP implementation, we use the concept of the post-decision state [8]. The post-decision state S_t^x is the state immediately after action x_t , but before the arrival of new information ω_{t+1} . Given our action x_t , we have a deterministic transition from S_t to the so-called post-decision state S_t^x . We express this transition in the function

$$S_t^x = S^{M,x}(S_t) , \quad (14)$$

where

$$I_{t,v,l,e,d} = I_{t,v,l,e,d} - x_{t,v,l,e,d} , \quad (15)$$

$$\forall t \in \mathcal{T}, \forall v \in \mathcal{V}, \forall l \in \mathcal{L}, \forall e \in \mathcal{E}, \forall d \in \mathcal{D} ,$$

$$q_{t,r} = q_{t,r} - \bar{q}_t , \forall r \in [0, \tau_{route} - 1] . \quad (16)$$

The post-decision state has a corresponding value function

$$V_t(S_t^x) = \mathbb{E}\{V_{t+1}(S_{t+1})|S_t^x\} . \quad (17)$$

Adopting the concept of the post-decision state allows us to represent our problem as a deterministic minimization problem. Although this reduces the computational effort, Equation (17) still requires to evaluate all states in the outcome space. In ADP, we therefore replace this value function with a single value function approximation $\bar{V}_t^{n-1}(S_t^x)$. n is an iteration counter, representing that we use an estimate from iteration $n - 1$ at iteration n . At every decision moment, we take the best decision given our value function approximation. Incoming arrivals are generated according to Equation (13). Utilizing the post-decision state and value function approximation for future costs, we find the best action by solving the following deterministic minimization problem for the n -th iteration:

$$\tilde{x}_t^n = \arg \min_{x_t \in \mathcal{X}_t(S_t)} (C_t(S_t, x_t) + \bar{V}_t^{n-1}(S_t^x)) . \quad (18)$$

Equation (18) provides the action that minimizes the value \hat{v}_t^n , with the function to compute \hat{v}_t^n being

$$\hat{v}_t^n = \min_{x_t \in \mathcal{X}_t(S_t)} (C_t(S_t, x_t) + \bar{V}_t^{n-1}(S_t^x)) . \quad (19)$$

Once we obtain our estimate \hat{v}_t^n , we can update $\bar{V}_{t-1}^{n-1}(S_{t-1}^x)$. For this, we use the following function:

$$\bar{V}_{t-1}^n(S_{t-1}^x) \leftarrow U^V(\bar{V}_{t-1}^{n-1}(S_{t-1}^x), S_{t-1}^x, \hat{v}_t^n) . \quad (20)$$

Table 1 provides an outline of our ADP algorithm.

We briefly discuss two options for the function U^V to update $\bar{V}_t^n(S_{t-1}^x)$: lookup and value function approximation (VFA). With the lookup approach, we store an estimate $\bar{V}_t^n(S_t^x)$ for every post-decision state, which is updated based on our observation at the next decision moment. We can speed up this procedure by first completing a full iteration, and then update all post-decision values at once (a procedure known as double pass, see [8]). Though the lookup ADP resolves several computational challenges of dynamic programming, we still need to visit a state to learn about its value. Hence, in order to properly learn the value of being in a specific state, we need to visit a large number of states. Instead, we want to learn about the value of many states with a single observation. To achieve this, we make use of VFA with the so-called basis function approach, see [8]. Let \mathcal{F} be a set of features, with $f \in \mathcal{F}$ being some variable that partially explains the costs of being in a state. Relevant features for our dispatch problem are, e.g., the total volume of orders in inventory, the number of orders with

Table 1. ADP algorithm with post-decision states

Step 0 Initialize	
Step 0a:	Initialize $\bar{V}_t^0(S_t), \forall t \in \mathcal{T}, \forall S_t \in \mathcal{S}$
Step 0b:	Set iteration counter to $n = 1$, and set the maximum number of iterations to N .
Step 0c:	Select an initial state S_0 .
Step 1 For $t = 0$ to T do:	
Step 1a:	Find the best action \tilde{x}_t^n by solving Equation (18).
Step 1b:	If $t > 0$, then update $\bar{V}_t^{n-1}(S_t)$ using Equation (20).
Step 1c:	Obtain the post-decision state S_t^x via Equation (14).
Step 1d:	Obtain a sample realization W_{t+1} , calculate S_{t+1} with Equation (9)
Step 2 Set $n := n + 1$.	
If $n \leq N$, then go to Step 1.	
Step 3 Return $\bar{V}_t^N(S_t^x), \forall t \in \mathcal{T}$.	

$d = 0$, and the number of distinct destinations. Let $\phi_f(S_t^x)$ be a basis function of feature f – for example, a cross-product or a polynomial of f – that returns a certain value given S_t^x . Let θ_f^n be a weight corresponding to feature f . Our value function approximation becomes

$$\bar{V}_t^n(S_t^x) = \sum_{f \in \mathcal{F}} \theta_f^n \phi_f(S_t^x), \forall t \in \mathcal{T} . \quad (21)$$

Following [8], the weights θ_f^n are updated using recursive least squares for nonstationary data. Using this procedure, we are able to learn about the value of many states by sampling just a single state. Using VFA, it is therefore not necessary to visit all states in the state space to learn about their value, allowing to handle large state spaces. The key difficulty with VFA is to define basis functions that closely approximate the exact values of states. Good insight in the structure of the problem is required to select features that allow to accurately approximate the true values.

After learning the appropriate weights by completing the procedure in Table 1, VFA can be applied for practical decision making. By calculating the values for the post-decision states corresponding to our initial state, we are able to obtain the best action given the estimate. Only the features of the states, the basis functions, and the corresponding weights are necessary for decision making.

5 Numerical experiments

First, we solve a toy-sized instance with dynamic programming. We show how both the lookup approach and the VFA approach approximate the exact DP-values. Next, we consider larger problems. As an exact benchmark is missing for larger instances, we cannot show convergence results for these. Instead, we

assess the *performance* of ADP. We take the results from the ADP procedure – i.e., the feature weights obtained from the VFA approach – to make decisions. We then compare ADP-based simulation results to the results of two benchmark policies. The first benchmark policy (‘Postpone’) we deploy in this paper is given in Table 2, and essentially aims to postpone as many orders as possible, until a suitable consolidation opportunity arises. The second benchmark policy (‘DirectShipment’) always ships orders, as long as primary vehicle capacity is available. ‘DirectShipment’ sorts and assigns orders just as ‘Postpone’ describes, and dispatches in the same way when secondary vehicles are required.

Table 2. Benchmark policy – Postpone

Step 0	Sort orders.	
	Step 0a:	Sort available orders based on lowest d .
	Step 0b:	Sort available orders with same d based on smallest size.
Step 1	While orders with $d = 0$ are unassigned do:	Assign order with $d = 0$ to vehicle.
Step 2	While remaining inventory exceeds l^{max} do:	Assign first order on list to vehicle.
Step 3	While capacity from already dispatched vehicles remains do:	Assign first order on list to vehicle.

We first describe the properties of our toy problem. We consider a fleet of two primary vehicles; secondary vehicles are twice as expensive as primary vehicles. We consider three distinct customer locations, a random order size from $\{0.2, 0.4, 0.6, 0.8, 1\}$, a maximum inventory of two orders, and a maximum of two arrivals per decision moment. We fix the tour length at $\tau_{route} = 1$. We set $e = 0$ for all orders, and randomly select d from $\{0, 1\}$. All probability distributions are uniform. We define a planning horizon with five decision moments. To provide some insight in the dimensionality of our dispatch problem: this toy-sized instance already has a state space of about 140,000 states.

The features we use for our VFA are (i) a constant, (ii) the number of vehicles available at the decision moment, (iii) the number of distinct order destinations, (iv) the total volume of orders in inventory, and (v) the square of the volume of orders in inventory. In Figure 1 and Figure 2, we show for two initial states (one without initial inventory, the other with four orders at the decision moment) how both the lookup approach and the VFA approach converge to the optimal values found with DP. In the first number of iterations, the estimates fluctuate due to the inability to accurately compute expected costs. However, by learning the values of visited states, ADP starts recognizing good actions.

From here on, we focus only on ADP with VFA using basis functions. The values we learn with VFA – by completing the algorithm in Table 1 – result in a set of weights for every decision moment. When in a given state, these weights allow us to estimate the values of all post-decision states that can be




Fig. 1. Approximation of exact value with Lookup




Fig. 2. Approximation of exact value with VFA

reached. Hence, ADP results in a policy, which we can use to solve a deterministic decision problem. We apply the learned policy in a Monte Carlo simulation on a variety of initial states, comparing the performance of the learned policy to both DP and the benchmark policies. For all simulations, we use the same realizations of order arrivals, and perform 10,000 simulation replications over the planning horizon. We do this for ten initial states, selected to represent a variety of properties. Table 3 shows the comparison between DP, ADP, and the two benchmark policies. The percentages indicate the average difference in costs between the optimal solution and the simulation results. By applying our ADP policy, we incur costs that are on average 0.60% higher than the optimal solution. ADP thereby outperforms the policy ‘Postpone’ by 3.14%, and the policy ‘DirectShipment’ by 11.56%. With ADP, we postpone 24% less orders than ‘Postpone’ does. For the initial states where ‘DirectShipment’ actually postpones orders – for some initial states it never will – ADP postpones 203% more.

Table 3. Comparison between ADP and benchmark policies for small instances

	Average costs	Average deviation from optimal	Lowest deviation from optimal	Highest deviation from optimal
DP	876	0.00%	0.00%	0.00%
ADP	881	0.60%	0.45%	0.99%
Postpone	908	3.76%	1.98%	9.41%
DirectShipment	1033	12.23%	8.52%	18.66%

Finally, we perform tests on five larger instances, with 10 customers, 2-5 vehicles, 10 order sizes, 5-10 arrivals per time unit, a maximum inventory of 5-10, a maximum dispatch window of 2, and a time horizon of 10. To illustrate how quickly the size of the state space grows, we give a small numerical example. With the mentioned settings, we have $10 \cdot 10 \cdot 3 = 300$ order types. Suppose we have 10 orders – all of different types – at a given decision moment. We then already have $\frac{300!}{290!10!}$ states only for this particular combination of order arrivals. This multinomial coefficient should be computed for all possible combinations of order arrivals; it follows that the size of the state space is $\gg 10^{18}$. Clearly, an exact benchmark for such instances cannot be provided.

Table 4 shows the results of our experiments on our larger instances. On average, ADP outperforms the policy ‘Postpone’ with 10.11% and the policy ‘DirectShipment’ with 19.71%. Particularly in cases with scarce vehicle capacity, ADP performs notably better. Under both benchmark policies, secondary vehicles are dispatched more frequently. Also, the benefits of ADP are greater when considering higher numbers of arrivals; the benchmark heuristics have no rules to cope with the larger complexity in dispatching decisions.

Table 4. Comparison between ADP and benchmark policies for larger instances

	Average costs	Average deviation from ADP	Lowest deviation from ADP	Highest deviation from ADP
ADP	3116	0.00%	0.00%	0.00%
Postpone	3443	10.11%	5.01%	12.93%
DirectShipment	3658	19.71%	10.33%	27.41%

6 Conclusions

In this paper, we proposed an ADP approach to make dispatch decisions at urban consolidation centers. We optimized decisions for a certain planning horizon, taking into account stochastic order arrivals during this horizon. We have shown that ADP is able to closely approximate the optimal values obtained by DP for toy-sized instances of our problem. For larger instances, ADP clearly and consistently outperforms both benchmark policies.

The ADP approach as described in this paper resolves the intractability of the state space and outcome space. However, we have not yet addressed the size of the action space, which can become problematically large as well. A possible approach to tackle this problem – without affecting the quality of decision-making – is to express the decision problem as an integer linear program that can be solved with limited computational effort. This requires the basis functions in VFA to be defined in such a way that they are linear with the decision problem.

Our numerical experiments have shown that even for small instances, simple waiting policies are inadequate to capture the complexity of waiting decisions. We expect that the ability of ADP to capture such complexities will be even more pertinent when increasing the size and complexity of instances. Further research will focus on the evaluation of realistically-sized instances and comparison with more sophisticated benchmark policies. The basis functions as proposed in this paper may not work well on every instance. Insights in appropriate VFAs for a variety of problem structures will yield a valuable contribution to existing literature. Both the ADP approach and its benchmarks need to be refined in order to provide an in-depth analysis of the applicability of ADP in realistic-sized dispatch problems.

Bibliography

- [1] Bailey, N.T.: On queueing processes with bulk service. *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 80–87 (1954)
- [2] Cai, Q.: *Shipment Consolidation in Discrete Time and Discrete Quantity: Matrix-Analytic Methods*. Master’s thesis, University of Waterloo (2011)
- [3] Cetinkaya, S., Lee, C.Y.: Stock replenishment and shipment scheduling for vendor-managed inventory systems. *Management Science* 46(2), 217–232 (2000)
- [4] Daganzo, C.F.: The distance traveled to visit n points with a maximum of c stops per vehicle: An analytic model and an application. *Transportation Science* 18(4), 331–350 (1984)
- [5] Ichoua, S., Gendreau, M., Potvin, J.Y.: Exploiting knowledge about future demands for real-time vehicle dispatching. *Transportation Science* 40(2), 211–225 (2006)
- [6] Minkoff, A.S.: A markov decision model and decomposition heuristic for dynamic vehicle dispatching. *Operations Research* 41(1), 77–90 (1993)
- [7] Neuts, M.F.: A general class of bulk queues with poisson input. *The Annals of Mathematical Statistics* pp. 759–770 (1967)
- [8] Powell, W.B.: *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, vol. 842. John Wiley & Sons (2011)
- [9] Powell, W.B., Topaloglu, H.: Stochastic programming in transportation and logistics. *Handbooks in operations research and management science* 10, 555–635 (2003)
- [10] Quak, H.: Sustainability of urban freight transport: Retail distribution and local regulations in cities. No. EPS-2008-124-LIS, Erasmus Research Institute of Management (ERIM) (2008)
- [11] Robusté, F., Estrada, M., López-Pita, A.: Formulas for estimating average distance traveled in vehicle routing problems in elliptic zones. *Transportation Research Record: Journal of the Transportation Research Board* 1873(1), 64–69 (2004)
- [12] SteadieSeifi, M., Dellaert, N., Van Woensel, T.: A multimodal network flow problem with product quality preservation, transshipment, and asset management (2014)
- [13] Topaloglu, H., Powell, W.B.: Dynamic-programming approximations for stochastic time-staged integer multicommodity-flow problems. *INFORMS Journal on Computing* 18(1), 31–42 (2006)
- [14] Wang, H., Odoni, A.: Approximating the performance of a last mile transportation system. *Transportation Science* (2014)
- [15] Wells, M.B.: *Elements of combinatorial computing*. Elsevier (2014)

Nr.	Year	Title	Author(s)
474	2015	An approximate dynamic programming approach to urban freight distribution with batch arrivals	Wouter van Heeswijk, Martijn Mes, Marco Schutten
473	2015	Dynamic Multi-period Freight Consolidation	Arturo Pérez Rivera, Martijn Mes
472	2015	Maintenance policy selection for ships: finding the most important criteria and considerations	A.J.M. Goossens, R.J.I. Basten
471	2015	Using Twitter to Predict Sales: A Case Study	Remco Dijkman, Panagiotis Ipeirotis, Freek Aertsen, Roy van Helden
470	2015	The Effect of Exceptions in Business Processes	Remco Dijkman, Geoffrey van IJzendoorn, Oktay Türetken, Meint de Vries
469	2015	Business Model Prototyping for Intelligent Transport Systems. A Service-Dominant Approach	Konstantinos Traganos, Paul Grefen, Aafke den Hollander, Oktay Türetken, Rik Eshuis
468	2015	How suitable is the RePro technique for rethinking care processes?	Rob J.B. Vanwersch, Luise Pufahl, Irene Vanderfeesten, Jan Mendling, Hajo A. Reijers
467	2014	Where to exert abatement effort for sustainable operations considering supply chain interactions?	Tarkan Tan, Astrid Koomen
466	2014	An Exact Algorithm for the Vehicle Routing Problem with Time Windows and Shifts	Said Dabia, Stefan Ropke, Tom Van Woensel
465	2014	The RePro technique: a new, systematic technique for rethinking care processes	Rob J.B. Vanwersch, Luise Pufahl, Irene Vanderfeesten, Hajo A. Reijers
464	2014	Exploring maintenance policy selection using the Analytic Hierarchy Process: an application for naval ships	A.J.M. Goossens, R.J.I. Basten
463	2014	Allocating service parts in two-echelon networks at a utility company	D. van den Berg, M.C. van der Heijden, P.C. Schuur
462	2014	Freight consolidation in networks with transshipments	W.J.A. van Heeswijk, M.R.K. Mes, J.M.J. Schutten, W.H.M. Zijm
461	2014	A Software Architecture for a Transportation Control Tower	Anne Baumgrass, Remco Dijkman, Paul Grefen, Shaya Pourmirza, Hagen Völzer, Mathias Weske
460	2014	Small traditional retailers in emerging markets	Youssef Boulaksil, Jan C. Fransoo, Edgar E. Blanco, Sallem Koubida
459	2014	Defining line replaceable units	J.E. Parada Puig, R.J.I. Basten
458	2014	Inventories and the Credit Crisis: A Chicken and Egg Situation	Maximiliano Udenio, Vishal Gaur, Jan C. Fransoo
457	2014	An Exact Approach for the Pollution-Routing Problem	Said Dabia, Emrah Demir, Tom Van Woensel
456	2014	Fleet readiness: stocking spare parts and high-tech assets	Rob J.I. Basten, Joachim J. Arts
455	2014	Competitive Solutions for Cooperating Logistics Providers	Behzad Hezarkhani, Marco Slikker, Tom Van Woensel
454	2014	Simulation Framework to Analyse Operating Room Release Mechanisms	Rimmert van der Kooij, Martijn Mes, Erwin Hans
453	2014	A Unified Race Algorithm for Offline Parameter Tuning	Tim van Dijk, Martijn Mes, Marco Schutten, Joaquim Gromicho
452	2014	Cost, carbon emissions and modal shift in intermodal network design decisions	Yann Bouchery, Jan Fransoo
451	2014	Transportation Cost and CO2 Emissions in Location Decision Models	Josue C. Vélazquez-Martínez, Jan C. Fransoo, Edgar E. Blanco, Jaime Mora-Vargas
450	2014	Tracebook: A Dynamic Checklist Support System	Shan Nan, Pieter Van Gorp, Hendrikus H.M. Korsten, Richard Vdovjak, Uzay Kaymak

Nr.	Year	Title	Author(s)
449	2014	Intermodal hinterland network design with multiple actors	Yann Bouchery, Jan Fransoo
448	2014	The Share-a-Ride Problem: People and Parcels Sharing Taxis	Baoxiang Li, Dmitry Krushinsky, Hajo A. Reijers, Tom Van Woensel
447	2014	Stochastic inventory models for a single item at a single location	K.H. van Donselaar, R.A.C.M. Broekmeulen
446	2014	Optimal and heuristic repairable stocking and expediting in a fluctuating demand environment	Joachim Arts, Rob Basten, Geert-Jan van Houtum
445	2014	Connecting inventory control and repair shop control: a differentiated control structure for repairable spare parts	M.A. Driessen, W.D. Rustenburg, G.J. van Houtum, V.C.S. Wiers
444	2014	A survey on design and usage of Software Reference Architectures	Samuil Angelov, Jos Trienekens, Rob Kusters
443	2014	Extending and Adapting the Architecture Tradeoff Analysis Method for the Evaluation of Software Reference Architectures	Samuil Angelov, Jos J.M. Trienekens, Paul Grefen
442	2014	A multimodal network flow problem with product quality preservation, transshipment, and asset management	Maryam SteadieSeifi, Nico Dellaert, Tom Van Woensel
441	2013	Integrating passenger and freight transportation: Model formulation and insights	Veaceslav Ghilas, Emrah Demir, Tom Van Woensel
440	2013	The Price of Payment Delay	K. van der Vliet, M.J. Reindorp, J.C. Fransoo
439	2013	On Characterization of the Core of Lane Covering Games via Dual Solutions	Behzad Hezarkhani, Marco Slikker, Tom van Woensel
438	2013	Destocking, the Bullwhip Effect, and the Credit Crisis: Empirical Modeling of Supply Chain Dynamics	Maximiliano Udenio, Jan C. Fransoo, Robert Peels
437	2013	Methodological support for business process redesign in healthcare: a systematic literature review	Rob J.B. Vanwersch, Khurram Shahzad, Irene Vanderfeesten, Kris Vanhaecht, Paul Grefen, Liliane Pintelon, Jan Mendling, Geofridus G. van Merode, Hajo A. Reijers
436	2013	Dynamics and equilibria under incremental horizontal differentiation on the Salop circle	B. Vermeulen, J.A. La Poutré, A.G. de Kok
435	2013	Analyzing Conformance to Clinical Protocols Involving Advanced Synchronizations	Hui Yan, Pieter Van Gorp, Uzay Kaymak, Xudong Lu, Richard Vdovjak, Hendriks H.M. Korsten, Huilong Duan
434	2013	Models for Ambulance Planning on the Strategic and the Tactical Level	J. Theresia van Essen, Johann L. Hurink, Stefan Nickel, Melanie Reuter
433	2013	Mode Allocation and Scheduling of Inland Container Transportation: A Case-Study in the Netherlands	Stefano Fazi, Tom Van Woensel, Jan C. Fransoo
432	2013	Socially responsible transportation and lot sizing: Insights from multiobjective optimization	Yann Bouchery, Asma Ghaffari, Zied Jemai, Jan Fransoo
431	2013	Inventory routing for dynamic waste collection	Martijn Mes, Marco Schutten, Arturo Pérez Rivera
430	2013	Simulation and Logistics Optimization of an Integrated Emergency Post	N.J. Borgman, M.R.K. Mes, I.M.H. Vliegen, E.W. Hans
429	2013	Last Time Buy and Repair Decisions for Spare Parts	S. Behfard, M.C. van der Heijden, A. Al Hanbali, W.H.M. Zijm
428	2013	A Review of Recent Research on Green Road Freight Transportation	Emrah Demir, Tolga Bektas, Gilbert Laporte
427	2013	Typology of Repair Shops for Maintenance Spare Parts	M.A. Driessen, V.C.S. Wiers, G.J. van Houtum, W.D. Rustenburg
426	2013	A value network development model and implications for innovation and production network management	B. Vermeulen, A.G. de Kok
425	2013	Single Vehicle Routing with Stochastic Demands: Approximate Dynamic Programming	C. Zhang, N.P. Dellaert, L. Zhao, T. Van Woensel, D. Sever

Nr.	Year	Title	Author(s)
424	2013	Influence of Spillback Effect on Dynamic Shortest Path Problems with Travel-Time-Dependent Network Disruptions	Derya Sever, Nico Dellaert, Tom Van Woensel, Ton de Kok
423	2013	Dynamic Shortest Path Problem with Travel-Time-Dependent Stochastic Disruptions: Hybrid Approximate Dynamic Programming Algorithms with a Clustering Approach	Derya Sever, Lei Zhao, Nico Dellaert, Tom Van Woensel, Ton de Kok
422	2013	System-oriented inventory models for spare parts	R.J.I. Basten, G.J. van Houtum
421	2013	Lost Sales Inventory Models with Batch Ordering and Handling Costs	T. Van Woensel, N. Erkip, A. Curseu, J.C. Fransoo
420	2013	Response speed and the bullwhip	Maximiliano Udenio, Jan C. Fransoo, Eleni Vatamidou, Nico Dellaert
419	2013	Anticipatory Routing of Police Helicopters	Rick van Urk, Martijn R.K. Mes, Erwin W. Hans
418	2013	Supply Chain Finance: research challenges ahead	Kasper van der Vliet, Matthew J. Reindorp, Jan C. Fransoo
417	2013	Improving the Performance of Sorter Systems by Scheduling Inbound Containers	S.W.A. Haneyah, J.M.J. Schutten, K. Fikse
416	2013	Regional logistics land allocation policies: Stimulating spatial concentration of logistics firms	Frank P. van den Heuvel, Peter W. de Langen, Karel H. van Donselaar, Jan C. Fransoo
415	2013	The development of measures of process harmonization	Heidi L. Romero, Remco M. Dijkman, Paul W.P.J. Grefen, Arjan van Weele
414	2013	BASE/X. Business Agility through Cross-Organizational Service Engineering. The Business and Service Design Approach developed in the CoProFind Project	Paul Grefen, Egon Lüftenegger, Eric van der Linden, Caren Weisleder
413	2013	The Time-Dependent Vehicle Routing Problem with Soft Time Windows and Stochastic Travel Times	Duygu Tas, Nico Dellaert, Tom van Woensel, Ton de Kok
412	2013	Clearing the Sky - Understanding SLA Elements in Cloud Computing	Marco Comuzzi, Guus Jacobs, Paul Grefen
411	2013	Approximations for the waiting time distribution in an M/G/c priority queue	A. Al Hanbali, E.M. Alvarez, M.C. van der Heijden
410	2013	To co-locate or not? Location decisions and logistics concentration areas	Frank P. van den Heuvel, Karel H. van Donselaar, Rob A.C.M. Broekmeulen, Jan C. Fransoo, Peter W. de Langen
409	2013	The Time-Dependent Pollution-Routing Problem	Anna Franceschetti, Dorothée Honhon, Tom van Woensel, Tolga Bektas, Gilbert Laporte
408	2013	Scheduling the scheduling task: A time management perspective on scheduling	J.A. Larco, V. Wiers, J. Fransoo
407	2013	Clustering Clinical Departments for Wards to Achieve a Prespecified Blocking Probability	J. Theresia van Essen, Mark van Houdenhoven, Johann L. Hurink
406	2013	MyPHRMachines: Personal Health Desktops in the Cloud	Pieter Van Gorp, Marco Comuzzi
405	2013	Maximising the Value of Supply Chain Finance	Kasper van der Vliet, Matthew J. Reindorp, Jan C. Fransoo
404	2013	Reaching 50 million nanostores: retail distribution in emerging megacities	Edgar E. Blanco, Jan C. Fransoo
403	2013	A Vehicle Routing Problem with Flexible Time Windows	Duygu Tas, Ola Jabali, Tom van Woensel
402	2013	The Service Dominant Business Model: A Service Focused Conceptualization	Egon Lüftenegger, Marco Comuzzi, Paul Grefen, Caren Weisleder

Nr.	Year	Title	Author(s)
401	2013	Relationship between freight accessibility and logistics employment in US counties	Frank P. van den Heuvel, Liliana Rivera, Karel H. van Donselaar, Ad de Jong, Yossi Sheffi, Peter W. de Langen, Jan C. Fransoo
400	2012	A Condition-Based Maintenance Policy for Multi-Component Systems with a High Maintenance Setup Cost	Qiyushi Zhu, Hao Peng, Geert-Jan van Houtum
399	2012	A flexible iterative improvement heuristic to support creation of feasible shift rosters in self-rostering	E. van der Veen, J.L. Hurink, J.M.J. Schutten, S.T. Uijland
398	2012	Scheduled Service Network Design with Synchronization and Transshipment Constraints for Intermodal Container Transportation Networks	K. Sharypova, T.G. Crainic, T. van Woensel, J.C. Fransoo
397	2012	Destocking, the bullwhip effect, and the credit crisis: empirical modeling of supply chain dynamics	Maximiliano Udenio, Jan C. Fransoo, Robert Peels
396	2012	Vehicle routing with restricted loading capacities	J. Gromicho, J.J. van Hoorn, A.L. Kok, J.M.J. Schutten
395	2012	Service differentiation through selective lateral transshipments	E.M. Alvarez, M.C. van der Heijden, I.M.H. Vliegen, W.H.M. Zijm
394	2012	A Generalized Simulation Model of an Integrated Emergency Post	Martijn Mes, Manon Bruens
393	2012	Business Process Technology and the Cloud: defining a Business Process Cloud Platform	Vassil Stoitsev, Paul Grefen
392	2012	Vehicle Routing with Soft Time Windows and Stochastic Travel Times: A Column Generation and Branch-and-Price Solution Approach	D. Tas, M. Gendreau, N. Dellaert, T. van Woensel, A.G. de Kok
391	2012	Improve OR-Schedule to Reduce Number of Required Beds	J. Theresia van Essen, Joël M. Bosch, Erwin W. Hans, Mark van Houdenhoven, Johann L. Hurink
390	2012	How does development lead time affect performance over the ramp-up lifecycle? Evidence from the consumer electronics industry	Andreas Pufall, Jan C. Fransoo, Ad de Jong, A.G. (Ton) de Kok
389	2012	The Impact of Product Complexity on Ramp-Up Performance	Andreas Pufall, Jan C. Fransoo, Ad de Jong, A.G. (Ton) de Kok
388	2012	Co-location synergies: specialized versus diverse logistics concentration areas	Frank P. van den Heuvel, Peter W. de Langen, Karel H. van Donselaar, Jan C. Fransoo
387	2012	Proximity matters: Synergies through co-location of logistics establishments	Frank P. van den Heuvel, Peter W. de Langen, Karel H. van Donselaar, Jan C. Fransoo
386	2012	Spatial concentration and location dynamics in logistics: the case of a Dutch province	Frank P. van den Heuvel, Peter W. de Langen, Karel H. van Donselaar, Jan C. Fransoo
385	2012	FNet: An Index for Advanced Business Process Querying	Zhiqiang Yan, Remco Dijkman, Paul Grefen
384	2012	Defining Various Pathway Terms	W.R. Dalinghaus, P.M.E. Van Gorp
383	2012	The Service Dominant Strategy Canvas: Defining and Visualizing a Service Dominant Strategy through the Traditional Strategic Lens	Egon Lüftenegger, Paul Grefen, Caren Weisleder
382	2012	A Stochastic Variable Size Bin Packing Problem with Time Constraints	Stefano Fazi, Tom van Woensel, Jan C. Fransoo
381	2012	Coordination and Analysis of Barge Container Hinterland Networks	K. Sharypova, T. van Woensel, J.C. Fransoo
380	2012	Proximity matters: Synergies through co-location of logistics establishments	Frank P. van den Heuvel, Peter W. de Langen, Karel H. van Donselaar, Jan C. Fransoo
379	2012	A literature review in process harmonization: a conceptual framework	Heidi Romero, Remco Dijkman, Paul Grefen, Arjan van Weele

Nr.	Year	Title	Author(s)
378	2012	A Generic Material Flow Control Model for Two Different Industries	S.W.A. Haneyah, J.M.J. Schutten, P.C. Schuur, W.H.M. Zijm
377	2012	Dynamic demand fulfillment in spare parts networks with multiple customer classes	H.G.H. Tiemessen, M. Fleischmann, G.J. van Houtum, J.A.E.E. van Nunen, E. Pratsini
376	2012	Paper has been replaced by wp 417	K. Fikse, S.W.A. Haneyah, J.M.J. Schutten
375	2012	Strategies for dynamic appointment making by container terminals	Albert Douma, Martijn Mes
374	2012	MyPHRMachines: Lifelong Personal Health Records in the Cloud	Pieter van Gorp, Marco Comuzzi
373	2012	Service differentiation in spare parts supply through dedicated stocks	E.M. Alvarez, M.C. van der Heijden, W.H.M. Zijm
372	2012	Spare parts inventory pooling: how to share the benefits?	Frank Karsten, Rob Basten
371	2012	Condition based spare parts supply	X. Lin, R.J.I. Basten, A.A. Kranenburg, G.J. van Houtum
370	2012	Using Simulation to Assess the Opportunities of Dynamic Waste Collection	Martijn Mes
369	2012	Aggregate overhaul and supply chain planning for rotables	J. Arts, S.D. Flapper, K. Vernooij
368	2012	Operating Room Rescheduling	J.T. van Essen, J.L. Hurink, W. Hartholt, B.J. van den Akker
367	2011	Switching Transport Modes to Meet Voluntary Carbon Emission Targets	Kristel M.R. Hoen, Tarkan Tan, Jan C. Fransoo, Geert-Jan van Houtum
366	2011	On two-echelon inventory systems with Poisson demand and lost sales	Elisa Alvarez, Matthieu van der Heijden
365	2011	Minimizing the Waiting Time for Emergency Surgery	J.T. van Essen, E.W. Hans, J.L. Hurink, A. Oversberg
364	2012	Vehicle Routing Problem with Stochastic Travel Times Including Soft Time Windows and Service Costs	Duygu Tas, Nico Dellaert, Tom van Woensel, Ton de Kok
363	2011	A New Approximate Evaluation Method for Two-Echelon Inventory Systems with Emergency Shipments	Erhun Özkan, Geert-Jan van Houtum, Yasemin Serin
362	2011	Approximating Multi-Objective Time-Dependent Optimization Problems	Said Dabia, El-Ghazali Talbi, Tom Van Woensel, Ton de Kok
361	2011	Branch and Cut and Price for the Time Dependent Vehicle Routing Problem with Time Windows	Said Dabia, Stefan Röpkke, Tom Van Woensel, Ton de Kok
360	2011	Analysis of an Assemble-to-Order System with Different Review Periods	A.G. Karaarslan, G.P. Kiesmüller, A.G. de Kok
359	2011	Interval Availability Analysis of a Two-Echelon, Multi-Item System	Ahmad Al Hanbali, Matthieu van der Heijden
358	2011	Carbon-Optimal and Carbon-Neutral Supply Chains	Felipe Caro, Charles J. Corbett, Tarkan Tan, Rob Zuidwijk
357	2011	Generic Planning and Control of Automated Material Handling Systems: Practical Requirements Versus Existing Theory	Sameh Haneyah, Henk Zijm, Marco Schutten, Peter Schuur
356	2011	Last time buy decisions for products sold under warranty	Matthieu van der Heijden, Bermawi Iskandar
355	2011	Spatial concentration and location dynamics in logistics: the case of a Dutch province	Frank P. van den Heuvel, Peter W. de Langen, Karel H. van Donselaar, Jan C. Fransoo
354	2011	Identification of Employment Concentration Areas	Frank P. van den Heuvel, Peter W. de Langen, Karel H. van Donselaar, Jan C. Fransoo
353	2011	BPMN 2.0 Execution Semantics Formalized as Graph Rewrite Rules: extended version	Pieter van Gorp, Remco Dijkman

Nr.	Year	Title	Author(s)
352	2011	Resource pooling and cost allocation among independent service providers	Frank Karsten, Marco Slikker, Geert-Jan van Houtum
351	2011	A Framework for Business Innovation Directions	E. Lüftenegger, S. Angelov, P. Grefen
350	2011	The Road to a Business Process Architecture: An Overview of Approaches and their Use	Remco Dijkman, Irene Vanderfeesten, Hajo A. Reijers
349	2011	Effect of carbon emission regulations on transport mode selection under stochastic demand	K.M.R. Hoen, T. Tan, J.C. Fransoo, G.J. van Houtum
348	2011	An improved MIP-based combinatorial approach for a multi-skill workforce scheduling problem	Murat Firat, Cor Hurkens
347	2011	An approximate approach for the joint problem of level of repair analysis and spare parts stocking	R.J.I. Basten, M.C. van der Heijden, J.M.J. Schutten
346	2011	Joint optimization of level of repair analysis and spare parts stocks	R.J.I. Basten, M.C. van der Heijden, J.M.J. Schutten
345	2011	Inventory control with manufacturing lead time flexibility	Ton G. de Kok
344	2011	Analysis of resource pooling games via a new extension of the Erlang loss function	Frank Karsten, Marco Slikker, Geert-Jan van Houtum
343	2011	Vehicle refueling with limited resources	Murat Firat, C.A.J. Hurkens, Gerhard J. Woeginger
342	2011	Optimal Inventory Policies with Non-stationary Supply Disruptions and Advance Supply Information	Bilge Atasoy, Refik Güllü, Tarkan Tan
341	2011	Redundancy Optimization for Critical Components in High-Availability Capital Goods	Kurtulus Baris Öner, Alan Scheller-Wolf, Geert-Jan van Houtum
340	2011	Making Decision Process Knowledge Explicit Using the Product Data Model	Razvan Petrusel, Irene Vanderfeesten, Cristina Claudia Dolean, Daniel Mican
339	2010	Analysis of a two-echelon inventory system with two supply modes	Joachim Arts, Gudrun Kiesmüller
338	2010	Analysis of the dial-a-ride problem of Hunsaker and Savelsbergh	Murat Firat, Gerhard J. Woeginger
335	2010	Attaining stability in multi-skill workforce scheduling	Murat Firat, Cor Hurkens
334	2010	Flexible Heuristics Miner (FHM)	A.J.M.M. Weijters, J.T.S. Ribeiro
333	2010	An exact approach for relating recovering surgical patient workload to the master surgical schedule	P.T. Vanberkel, R.J. Boucherie, E.W. Hans, J.L. Hurink, W.A.M. van Lent, W.H. van Harten
332	2010	Efficiency evaluation for pooling resources in health care	Peter T. Vanberkel, Richard J. Boucherie, Erwin W. Hans, Johann L. Hurink, Nelly Litvak
331	2010	The Effect of Workload Constraints in Mathematical Programming Models for Production Planning	M.M. Jansen, A.G. de Kok, I.J.B.F. Adan
330	2010	Using pipeline information in a multi-echelon spare parts inventory system	Christian Howard, Ingrid Reijnen, Johan Marklund, Tarkan Tan
329	2010	Reducing costs of repairable spare parts supply systems via dynamic scheduling	H.G.H. Tiemessen, G.J. van Houtum
328	2010	Identification of Employment Concentration and Specialization Areas: Theory and Application	Frank P. van den Heuvel, Peter W. de Langen, Karel H. van Donselaar, Jan C. Fransoo
327	2010	A combinatorial approach to multi-skill workforce scheduling	M. Firat, C. Hurkens
326	2010	Stability in multi-skill workforce scheduling	M. Firat, C. Hurkens, A. Laugier
325	2010	Maintenance spare parts planning and control: A framework for control and agenda for future research	M.A. Driessen, J.J. Arts, G.J. van Houtum, W.D. Rustenburg, B. Huisman
324	2010	Near-optimal heuristics to set base stock levels in a two-echelon distribution network	R.J.I. Basten, G.J. van Houtum

Nr.	Year	Title	Author(s)
323	2010	Inventory reduction in spare part networks by selective throughput time reduction	M.C. van der Heijden, E.M. Alvarez, J.M.J. Schutten
322	2010	The selective use of emergency shipments for service-contract differentiation	E.M. Alvarez, M.C. van der Heijden, W.H.M. Zijm
321	2010	Heuristics for Multi-Item Two-Echelon Spare Parts Inventory Control Problem with Batch Ordering in the Central Warehouse	Engin Topan, Z. Pelin Bayindir, Tarkan Tan
320	2010	Preventing or escaping the suppression mechanism: intervention conditions	Bob Walrave, Kim E. van Oorschot, A. Georges L. Romme
319	2010	Hospital admission planning to optimize major resources utilization under uncertainty	Nico Dellaert, Jully Jeunet
318	2010	Minimal Protocol Adaptors for Interacting Services	R. Seguel, R. Eshuis, P. Grefen
317	2010	Teaching Retail Operations in Business and Engineering Schools	Tom Van Woensel, Marshall L. Fisher, Jan C. Fransoo
316	2010	Design for Availability: Creating Value for Manufacturers and Customers	Lydie P.M. Smets, Geert-Jan van Houtum, Fred Langerak
315	2010	Transforming Process Models: executable rewrite rules versus a formalized Java program	Pieter van Gorp, Rik Eshuis
314	2010	Working paper 314 is no longer available	----
313	2010	A Dynamic Programming Approach to Multi-Objective Time-Dependent Capacitated Single Vehicle Routing Problems with Time Windows	S. Dabia, T. van Woensel, A.G. de Kok
312	2010	Tales of a So(u)rcerer: Optimal Sourcing Decisions Under Alternative Capacitated Suppliers and General Cost Structures	Osman Alp, Tarkan Tan
311	2010	In-store replenishment procedures for perishable inventory in a retail environment with handling costs and storage constraints	R.A.C.M. Broekmeulen, C.H.M. Bakx
310	2010	The state of the art of innovation-driven business models in the financial services industry	E. Lüftenegger, S. Angelov, E. van der Linden, P. Grefen
309	2010	Design of Complex Architectures Using a Three Dimension Approach: the CrossWork Case	R. Seguel, P. Grefen, R. Eshuis
308	2010	Effect of carbon emission regulations on transport mode selection in supply chains	K.M.R. Hoen, T. Tan, J.C. Fransoo, G.J. van Houtum
307	2010	Interaction between intelligent agent strategies for real-time transportation planning	Martijn Mes, Matthieu van der Heijden, Peter Schuur
306	2010	Internal Slackening Scoring Methods	Marco Slikker, Peter Borm, René van den Brink
305	2010	Vehicle Routing with Traffic Congestion and Drivers' Driving and Working Rules	A.L. Kok, E.W. Hans, J.M.J. Schutten, W.H.M. Zijm
304	2010	Practical extensions to the level of repair analysis	R.J.I. Basten, M.C. van der Heijden, J.M.J. Schutten
303	2010	Ocean Container Transport: An Underestimated and Critical Link in Global Supply Chain Performance	Jan C. Fransoo, Chung-Yee Lee
302	2010	Capacity reservation and utilization for a manufacturer with uncertain capacity and demand	Y. Boulaksil; J.C. Fransoo; T. Tan
300	2009	Spare parts inventory pooling games	F.J.P. Karsten; M. Slikker; G.J. van Houtum
299	2009	Capacity flexibility allocation in an outsourced supply chain with reservation	Y. Boulaksil, M. Grunow, J.C. Fransoo
298	2010	An optimal approach for the joint problem of level of repair analysis and spare parts stocking	R.J.I. Basten, M.C. van der Heijden, J.M.J. Schutten
297	2009	Responding to the Lehman Wave: Sales Forecasting and Supply Management during the Credit Crisis	Robert Peels, Maximiliano Udenio, Jan C. Fransoo, Marcel Wolfs, Tom Hendriks

Nr.	Year	Title	Author(s)
296	2009	An exact approach for relating recovering surgical patient workload to the master surgical schedule	Peter T. Vanberkel, Richard J. Boucherie, Erwin W. Hans, Johann L. Hurink, Wineke A.M. van Lent, Wim H. van Harten
295	2009	An iterative method for the simultaneous optimization of repair decisions and spare parts stocks	R.J.I. Basten, M.C. van der Heijden, J.M.J. Schutten
294	2009	Fujaba hits the Wall(-e)	Pieter van Gorp, Ruben Jubeh, Bernhard Grusie, Anne Keller
293	2009	Implementation of a Healthcare Process in Four Different Workflow Systems	R.S. Mans, W.M.P. van der Aalst, N.C. Russell, P.J.M. Bakker
292	2009	Business Process Model Repositories - Framework and Survey	Zhiqiang Yan, Remco Dijkman, Paul Grefen
291	2009	Efficient Optimization of the Dual-Index Policy Using Markov Chains	Joachim Arts, Marcel van Vuuren, Gudrun Kiesmuller
290	2009	Hierarchical Knowledge-Gradient for Sequential Sampling	Martijn R.K. Mes; Warren B. Powell; Peter I. Frazier
289	2009	Analyzing combined vehicle routing and break scheduling from a distributed decision making perspective	C.M. Meyer; A.L. Kok; H. Kopfer; J.M.J. Schutten
288	2010	Lead time anticipation in Supply Chain Operations Planning	Michiel Jansen; Ton G. de Kok; Jan C. Fransoo
287	2009	Inventory Models with Lateral Transshipments: A Review	Colin Paterson; Gudrun Kiesmuller; Ruud Teunter; Kevin Glazebrook
286	2009	Efficiency evaluation for pooling resources in health care	P.T. Vanberkel; R.J. Boucherie; E.W. Hans; J.L. Hurink; N. Litvak
285	2009	A Survey of Health Care Models that Encompass Multiple Departments	P.T. Vanberkel; R.J. Boucherie; E.W. Hans; J.L. Hurink; N. Litvak
284	2009	Supporting Process Control in Business Collaborations	S. Angelov; K. Vidyasankar; J. Vonk; P. Grefen
283	2009	Inventory Control with Partial Batch Ordering	O. Alp; W.T. Huh; T. Tan
282	2009	Translating Safe Petri Nets to Statecharts in a Structure-Preserving Way	R. Eshuis
281	2009	The link between product data model and process model	J.J.C.L. Vogelaar; H.A. Reijers
280	2009	Inventory planning for spare parts networks with delivery time requirements	I.C. Reijnen; T. Tan; G.J. van Houtum
279	2009	Co-Evolution of Demand and Supply under Competition	B. Vermeulen; A.G. de Kok
278	2010	Toward Meso-level Product-Market Network Indices for Strategic Product Selection and (Re)Design Guidelines over the Product Life-Cycle	B. Vermeulen, A.G. de Kok
277	2009	An Efficient Method to Construct Minimal Protocol Adaptors	R. Seguel, R. Eshuis, P. Grefen
276	2009	Coordinating Supply Chains: a Bilevel Programming Approach	Ton G. de Kok, Gabriella Muratore
275	2009	Inventory redistribution for fashion products under demand parameter update	G.P. Kiesmuller, S. Minner
274	2009	Comparing Markov chains: Combining aggregation and precedence relations applied to sets of states	A. Busic, I.M.H. Vliegen, A. Scheller-Wolf
273	2009	Separate tools or tool kits: an exploratory study of engineers' preferences	I.M.H. Vliegen, P.A.M. Kleingeld, G.J. van Houtum
272	2009	An Exact Solution Procedure for Multi-Item Two-Echelon Spare Parts Inventory Control Problem with Batch Ordering	
271	2009	Distributed Decision Making in Combined Vehicle Routing and Break Scheduling	C.M. Meyer, H. Kopfer, A.L. Kok, M. Schutten

Nr.	Year	Title	Author(s)
270	2009	Dynamic Programming Algorithm for the Vehicle Routing Problem with Time Windows and EC Social Legislation	A.L. Kok, C.M. Meyer, H. Kopfer, J.M.J. Schutten
269	2009	Similarity of Business Process Models: Metrics and Evaluation	Remco Dijkman, Marlon Dumas, Boudewijn van Dongen, Reina Kaarik, Jan Mendling
267	2009	Vehicle routing under time-dependent travel times: the impact of congestion avoidance	A.L. Kok, E.W. Hans, J.M.J. Schutten
266	2009	Restricted dynamic programming: a flexible framework for solving realistic VRPs	J. Gromicho; J.J. van Hoorn; A.L. Kok; J.M.J. Schutten;

Working Papers published before 2009 see: <http://beta.ieis.tue.nl>