CIRRELT

Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport

Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation

An Optimized Target Level Inventory Replenishment Policy for Vendor-Managed Inventory Systems

Leandro C. Coelho
Gilbert Laporte

January 2013

CIRRELT-2013-05

Bureaux de Montréal :

Université de Montréal
C.P. 6128, succ. Centre-ville
Montréal (Québec)
Canada H3C 3J7
Téléphone : 514 343-7575
Télécopie : 514 343-7121

Bureaux de Québec :

Université Laval
2325, de la Terrasse, bureau 2642
Québec (Québec)
Canada G1V 0A6
Téléphone : 418 656-2073
Télécopie : 418 656-2624

www.cirrelt.ca

An Optimized Target Level Inventory Replenishment Policy for Vendor-Managed Inventory Systems

Leandro C. Coelho[1,2,*], **Gilbert Laporte**[1,3]

[1] Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT)

[2] Department of Logistics and Operations Management, HEC Montréal, 3000 Côte-Sainte-Catherine, Montréal, Canada H3T 2A7

[3] Department of Management Sciences, HEC Montréal, 3000 Côte-Sainte-Catherine, Montréal, Canada H3T 2A7

Abstract. In Vendor-Managed Inventory systems the supplier is responsible for replenishing customers and for deciding when and how much to deliver. One of two inventory policies is typically employed by the supplier. The first one, called the maximum level (ML) policy, gives full freedom to the supplier to deliver any quantity as long as it respects customer inventory capacities. The alternative, much more constrained, is called the order-up-to (OU) policy. It states that the supplier has to bring the customer inventory up to its maximum capacity level upon delivery. We propose a new tactical policy, called optimized target level, under which when the supplier visits a customer, the quantity delivered is such that the final inventory will always be at the same customer-dependent optimized target level. We perform a computational evaluation of this new policy against both traditional strategies on benchmark instances. We show that it yields lower costs and inventory levels than the OU policy, and is only marginally more expensive than the ML policy, while being easier to implement.

Keywords. Vendor-managed inventory, inventory-routing, replenishment policy, inventory management.

Acknowledgements. We thank Jean-François Cordeau and Raf Jans for suggesting this research topic. This work was partly supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) under grant 39682-10. This support is gratefully acknowledged. We also thank the Réseau québécois de calcul de haute performance (RQCHP) for providing parallel computing facilities.

* Corresponding author: Leandro.Coelho@cirrelt.ca

Dépôt légal – Bibliothèque et Archives nationales du Québec
Bibliothèque et Archives Canada, 2013

1 Introduction

In Vendor-Managed Inventory (VMI) systems, the supplier is responsible not only for delivering the products and routing its vehicles to serve its customers, but also for determining when and how much to deliver to them. The combined optimization of inventory control and vehicle routing gives rise to the Inventory-Routing Problem (IRP). There exist several real-life applications of this problem (see, e.g., the recent oil delivery case described by Uzar and Çatay [11]). The survey paper of Andersson et al. [1] concentrates on the applications of the IRP whereas Coelho et al. [8] provide an overview more focused on the methodological aspects of the problem.

In the IRP literature, two inventory replenishment policies are traditionally used. The first one, called maximum level (ML), gives full flexibility to the supplier. The quantity delivered to a customer only has to comply with its inventory capacity. The other policy, which is much more constrained, is called order-up-to (OU). Under an OU policy, whenever a customer is visited, its inventory level is brought up to its maximum capacity.

The OU policy simplifies the solution process since it effectively removes one decision dimension from the problem. Indeed, whenever a customer is visited, the quantity delivered is automatically computed as the difference between the inventory capacity and the current inventory level. In a recent study [7], the application of an OU policy was presented as a way to increase consistency and predictability within an IRP context. It therefore improves customer satisfaction since the delivery amounts are then more predictable. An interesting application of consistency in the context of a parcel distribution network can be found in Bard and Jarrah [4].

While the implementation of an OU policy simplifies the supplier's decision process and can increase customer satisfaction, it is not without disadvantages since it leads to cost increases which can be as high as 20% over an ML policy [2, 7]. An interesting question is therefore whether one can determine an optimal consistent replenishment level for each customer without incurring the full cost of an OU policy. This question has already been answered in the context of traditional (i.e., non-VMI) systems where each customer can individually optimize its replenishment level through the application of an (s, S) policy [3, 5], for example. However, as far as the authors are aware, the determination of a system-optimal and customer-dependent stable replenishment level has never been investigated in a VMI context.

Our objective is therefore to propose a new tactical replenishment policy that would combine the customer-related advantages of OU and the supplier-related benefit of ML which affords more flexibility and lower system costs. This new policy, which we call *optimized target level* (OTL), determines an optimal replenishment target level for each customer. It can be viewed as an optimized OU policy, except that instead replenishing up to the customer's capacity, the supplier fills the customer's inventory up to an OTL. In order to take advantage of this new idea, the OTL of each customer is computed simultaneously with the remaining routing and inventory decisions, in order to jointly optimize the inventory holding costs and the routing costs.

Figure 1 depicts the three inventory replenishment policies just described. Figure 1a illustrates the ML policy in which replenishment quantities only have to respect the maximum inventory capacity; in Figure 1b, all replenishments must bring the inventory level up to the maximum customer capacity C; finally, in Figure 1c, a target level L is optimized and all deliveries must ensure that the inventory level L is attained.

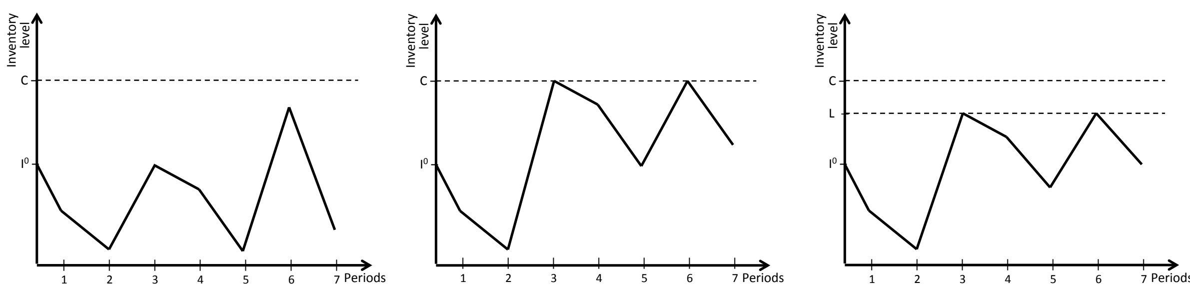

(a) The ML policy **(b)** The OU policy **(c)** The OTL policy

Figure 1: Different inventory replenishment policies

The remainder of this paper is organized as follows. Section 2 describes the problem at hand. In Section 3 we provide a formal mathematical model of each of the three inventory policies under study. In Section 4 we briefly describe the exact algorithm developed for the problem. Section 5 reports on computational experiments carried out to evaluate the performance of the new policy. This is followed by our conclusions in Section 6.

2 Description of the OTL Problem

We now describe the IRP framework that will be used to introduce the OTL policy. Since the routing costs are assumed to be symmetric, we define the problem on an undirected graph $\mathcal{G} = (\mathcal{V}, E)$, where $\mathcal{V} = \{0, ..., n\}$ is the vertex set and $E = \{(i,j) : i,j \in \mathcal{V}, i < j\}$ is the edge set. Vertex 0 represents the supplier and the remaining vertices $\mathcal{V}' = \mathcal{V} \setminus \{0\}$ represent n customers. A routing cost c_{ij} is associated with edge $(i, j) \in E$. Both the supplier and customers incur unit inventory holding costs h_i per period ($i \in \mathcal{V}$), and each customer has a maximum inventory holding capacity C_i. The length of the planning horizon is p and, at each time period $t \in \mathcal{T} = \{1, ..., p\}$, the supplier holds a quantity r^t of a single product. We assume the supplier has sufficient inventory to meet the full customer demand during the planning horizon, and all demand also has to be satisfied, i.e., backlogging is not allowed. At the beginning of the planning horizon the decision maker knows the current inventory level of the supplier and of the customers ($I_i^0, i \in \mathcal{V}$), and receives information on the demand d_i^t of each customer i for each time period t. The quantity r^t made available at the supplier in period t can be used for deliveries to customers in the same period, and the delivery amount received by customer i in period t can be used to meet the demand in that period. A set $\mathcal{K} = \{1, ..., K\}$ of vehicles are available. We denote by Q_k the capacity of vehicle k. Each vehicle can perform one route per time period, from the supplier to a subset of customers.

The aim is to set an OTL for each customer, to determine vehicle routes for each period and to compute delivery quantities for each period and each customer, in line with an OTL policy and so that all constraints are satisfied and the total cost is minimized.

3 Mathematical Formulations

In this section we provide a formal mathematical formulation for the OTL, starting in Section 3.1 with the basic IRP under the ML policy. We then present the OU policy in Section 3.2 and, in Section 3.3, the adaptations and transformations needed to model the IRP under an OTL policy as a mixed-integer linear program.

The models work with routing variables x_{ij}^{kt} equal to the number of times edge (i, j) is used on the route of vehicle k in period t. We also use binary variables y_i^{kt} equal to one if and only if node i is visited by vehicle k in period t. Let I_i^t denote the inventory level at vertex $i \in \mathcal{V}$ at the end of period $t \in \mathcal{T}$. We denote by q_i^{kt} the product quantity delivered by vehicle k to customer i in period t.

3.1 The IRP under an ML policy

The problem can be formulated under an ML inventory policy as follows:

$$\text{(ML) minimize} \sum_{i \in \mathcal{V}} \sum_{t \in \mathcal{T}} h_i I_i^t + \sum_{(i,j) \in E} \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} c_{ij} x_{ij}^{kt}, \tag{1}$$

subject to

$$I_0^t = I_0^{t-1} + r^t - \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{V}'} q_i^{kt} \quad t \in \mathcal{T} \tag{2}$$

$$I_i^t = I_i^{t-1} + \sum_{k \in \mathcal{K}} q_i^{kt} - d_i^t \quad i \in \mathcal{V}' \quad t \in \mathcal{T} \tag{3}$$

$$I_i^t \leq C_i \quad i \in \mathcal{V} \quad t \in \mathcal{T} \tag{4}$$

$$\sum_{k \in \mathcal{K}} q_i^{kt} \leq C_i - I_i^{t-1} \quad i \in \mathcal{V}' \quad t \in \mathcal{T} \tag{5}$$

$$q_i^{kt} \leq C_i y_i^{kt} \quad i \in \mathcal{V}' \quad k \in \mathcal{K} \quad t \in \mathcal{T} \tag{6}$$

$$\sum_{i \in \mathcal{V}'} q_i^{kt} \leq Q_k y_0^{kt} \quad k \in \mathcal{K} \quad t \in \mathcal{T} \tag{7}$$

$$\sum_{j \in \mathcal{V}, i<j} x_{ij}^{kt} + \sum_{j \in \mathcal{V}, j<i} x_{ji}^{kt} = 2 y_i^{kt} \quad i \in \mathcal{V} \quad k \in \mathcal{K} \quad t \in \mathcal{T} \tag{8}$$

$$\sum_{i \in \mathcal{S}} \sum_{j \in \mathcal{S}, i<j} x_{ij}^{kt} \leq \sum_{i \in \mathcal{S}} y_i^{kt} - y_m^{kt} \quad \mathcal{S} \subseteq \mathcal{V}' \quad k \in \mathcal{K} \quad t \in \mathcal{T} \quad m \in \mathcal{S} \tag{9}$$

$$I_i^t \geq 0 \quad i \in \mathcal{V} \quad t \in \mathcal{T} \tag{10}$$

$$q_i^{kt} \geq 0 \quad i \in \mathcal{V}' \quad k \in \mathcal{K} \quad t \in \mathcal{T} \tag{11}$$

$$x_{i0}^{kt} \in \{0,1,2\} \quad i \in \mathcal{V}' \quad k \in \mathcal{K} \quad t \in \mathcal{T} \tag{12}$$

$$x_{ij}^{kt} \in \{0,1\} \quad i,j \in \mathcal{V}' \quad k \in \mathcal{K} \quad t \in \mathcal{T} \tag{13}$$

$$y_i^{kt} \in \{0,1\} \quad i \in \mathcal{V} \quad k \in \mathcal{K} \quad t \in \mathcal{T}. \tag{14}$$

The objective function (1) minimizes the total inventory and routing costs. Constraints (2) and (3) define the inventory conservation at the supplier and at the customers. Constraints (4) impose maximal inventory levels at the customers. Constraints (5) and (6) link the quantities delivered to the routing variables. In particular, they only allow a vehicle to deliver products to a customer if the customer is visited by this vehicle. Constraints (7) ensure the vehicle capacities are respected, while constraints (8) and (9) are degree constraints and subtour elimination constraints, respectively. Constraints (10)−(14) enforce integrality and non-negativity conditions on the variables.

This model can be strengthened through the inclusion of the following families of valid inequalities [2, 6]:

$$x_{i0}^{kt} \leq 2y_i^{kt} \quad i \in \mathcal{V} \quad k \in \mathcal{K} \quad t \in \mathcal{T} \tag{15}$$

$$x_{ij}^{kt} \leq y_i^{kt} \quad i,j \in \mathcal{V} \quad k \in \mathcal{K} \quad t \in \mathcal{T} \tag{16}$$

$$y_i^{kt} \leq y_0^{kt} \quad i \in \mathcal{V}' \quad k \in \mathcal{K} \quad t \in \mathcal{T} \tag{17}$$

$$y_0^{kt} \leq y_0^{k-1,t} \quad k \in \mathcal{K}\backslash\{1\} \quad t \in \mathcal{T} \tag{18}$$

$$y_i^{kt} \leq \sum_{j<i} y_j^{k-1,t} \quad i \in \mathcal{V} \quad k \in \mathcal{K}\backslash\{1\} \quad t \in \mathcal{T} \tag{19}$$

$$\sum_{k \in \mathcal{K}} \sum_{l=1}^{t} y_i^{kl} \geq \left\lceil \left(\sum_{l=1}^{t} d_i^l - I_i^0 \right) / C_i \right\rceil \quad i \in \mathcal{V} \quad t \in \mathcal{T}. \tag{20}$$

Constraints (15) and (16) enforce the condition that if the supplier is the immediate successor of a customer in the route of vehicle k in period t, then i must be visited by the same vehicle. A similar reasoning is applied to customer j in inequalities (16). Constraints (17) ensure that the supplier is visited if any customer i is visited by vehicle k in period t.

When the fleet is homogeneous, one can break some vehicle symmetry by constraints (18), thus ensuring that vehicle k cannot leave the depot if vehicle $k-1$ is not used. This symmetry breaking rule is then extended to the customer vertices by constraints (19) which state that if customer i is assigned to vehicle k in period t, then vehicle $k-1$ must serve a customer with an index smaller than i in the same period.

Finally, constraints (20) ensure that customer i is visited at least the number of times corresponding to the right-hand side of the inequality. For each customer i, the supplier has to deliver until time t at least the total demand of customer i up to period t, minus its initial inventory. Since the maximum delivery quantity to customer i is its inventory capacity C_i, the minimum number of visits to i is determined by the right-hand side of (20). Again, this inequality is only valid if the fleet is homogeneous.

3.2 The IRP under an OU policy

In order to model the problem under an OU policy, it suffices the add the following constraints to the ML model. These ensure that a delivery should fill the customer's inventory capacity:

$$q_i^{kt} \geq C_i y_i^{kt} - I_i^{t-1} \quad i \in \mathcal{V}' \quad k \in \mathcal{K} \quad t \in \mathcal{T}. \tag{21}$$

The valid inequalities (15)−(20) described for the ML formulation are also valid for the OU policy.

3.3 The IRP under an OTL policy

The most straightforward way of adapting the ML and OU formulations to OTL is to introduce a variable target level to be attained whenever a customer is visited. To this end, the right-hand sides of constraints (5), (6) and (21) are defined as functions of a target level for customer i, denoted by an integer continuous variable L_i, which replaces C_i. One also needs to set the bounds on the new variables. The following constraints are therefore introduced:

$$\sum_{k \in \mathcal{K}} q_i^{kt} \leq L_i - I_i^{t-1} \quad i \in \mathcal{V}' \quad t \in \mathcal{T} \tag{22}$$

$$q_i^{kt} \leq L_i y_i^{kt} \quad i \in \mathcal{V}' \quad k \in \mathcal{K} \quad t \in \mathcal{T} \tag{23}$$

$$q_i^{kt} \geq L_i y_i^{kt} - I_i^{t-1} \quad i \in \mathcal{V}' \quad k \in \mathcal{K} \quad t \in \mathcal{T} \tag{24}$$

$$0 \leq L_i \leq C_i \quad i \in \mathcal{V}'. \tag{25}$$

The difficulty of this problem lies of course in constraints (23) and (24) which are non-linear since they contain products of a continuous and binary variables. These constraints state that the quantity delivered is that needed to bring the customer's inventory up to level L_i if there is a visit, and zero otherwise. Constraints (25) ensure that the optimized target level L_i is at most the customer's maximum inventory level C_i. Setting $L_i = C_i$ means that the OTL and OU policies coincide.

In order to linearize this model and ease its resolution, we replace constraints (23) and (24) with

$$0 \leq q_i^{kt} \leq L_i - I_i^t \quad i \in \mathcal{V}' \quad k \in \mathcal{K} \quad t \in \mathcal{T} \tag{26}$$

$$q_i^{kt} \leq C_i y_i^{kt} \quad i \in \mathcal{V}' \quad k \in \mathcal{K} \quad t \in \mathcal{T} \tag{27}$$

$$q_i^{kt} \geq L_i - I_i^t - (1 - y_i^{kt}) C_i \quad i \in \mathcal{V}' \quad k \in \mathcal{K} \quad t \in \mathcal{T}. \tag{28}$$

The role of these constraints is to ensure that if a delivery is made to customer i by vehicle k in period t, i.e., $y_i^{kt} = 1$, then the quantity delivered should be that to bring the inventory of i up to level L_i, i.e., $q_i^{kt} = L_i - I_i^t$; if, on the other hand, no delivery takes place, i.e., $y_i^{kt} = 0$, then the associated quantity should be zero as well, i.e., $q_i^{kt} = 0$. This is achieved by the combination of the proposed linear constraints. If $y_i^{kt} = 0$, then constraints (28) state that the associated quantity has to be larger than a negative value, and constraints (27) state that it has to be less than or equal to zero. Coupled with the left-hand side of constraints (26), they will force the quantity delivered to be equal to zero. Note that the left-hand side of (26) is presented here for the sake of completeness, since it was already stated by constraints (11). If $y_i^{kt} = 1$, then (27) is superseded by the right-hand side of (26), stating that the quantity delivered has to be smaller than or equal to $L_i - I_i^t$ which, coupled with constraints (28), ensures that exactly this quantity will be delivered to the customer. Note that rewriting these constraints in this way eliminates all non-linearities.

Finally, the model can be further strengthened by changing (4) to

$$I_i^t \leq L_i \quad i \in \mathcal{V} \quad t \in \mathcal{T}. \tag{29}$$

Thus, the OTL model is finally defined as (1)−(3), (7)−(14), (22), (25)−(29), and can be further improved through the inclusion of the valid inequalities (15)−(20).

4 Exact branch-and-cut algorithm

We now sketch the branch-and-cut scheme we have implemented to solve the model. At any given node of the search tree, the linear program defining the problem, without the subtour elimination constraints, is solved. At every node we perform a search for violated subtour elimination constraints (9), and some of these constraints are added to the current node which is then reoptimized. This process is repeated until a solution is found, or until there are no more cuts to be added. At this point branching on a fractional variable occurs in a typical branch-and-bound fashion.

We provide a sketch of the branch-and-bound-and-cut scheme in Algorithm 1.

Algorithm 1 Branch-and-cut algorithm

At the root node of the search tree, generate and insert all valid inequalities into the program.
Subproblem solution. Solve the LP relaxation of the node.
Termination check:
if there are no more nodes to evaluate **then**
  Stop.
else
  Select one node from the branch-and-cut tree.
end if
while the solution of the current LP relaxation contains subtours **do**
  Identify connected components as in Padberg and Rinaldi [10].
  Determine whether the component containing the supplier is weakly connected as in Gendreau et al. [9].
  Add all violated subtour elimination constraints (9).
  Subproblem solution. Solve the LP relaxation of the node.
end while
if the solution of the current LP relaxation is integer **then**
  Go to the termination check.
else
  Branching: branch on one of the fractional variables.
  Go to the termination check.
end if

5 Computational Experiments

In order to evaluate the impact of the new OTL inventory policy in the IRP, we have assessed its performance on benchmark instances. In particular, we have measured the impact of the OTL policy on cost, inventory levels, and computing time with respect to the classical ML and OU policies. To this end, we have solved the models of Section 3 with the algorithm described in Section 4, coded in C++, and using IBM Concert Technology and CPLEX 12.5 running in parallel with six threads. All computations were executed on a grid of Intel Xeon™ processors running at 2.66 GHz with up to 96 GB of RAM installed per node, with the Scientific Linux 6.1 operating system.

We have used the instance set designed by Archetti et al. [2] for the IRP to evaluate the OTL policy. It is composed of 160 instances with up to three time periods and 50 customers, and six time periods and 30 customers, under two levels of inventory holding cost. These instances will be referred to as *low-p* and *high-p*, corresponding to low and high inventory costs, respectively, p being the number of periods. Optimal solutions for the the ML and OU policies with single and multiple vehicles are known for these instances. All instances and detailed computational results are available from `http://www.leandro-coelho.com/instances`.

We provide in Table 1 average computational results for the IRP instances under the OTL policy. Optimal solutions were computed within 12 hours of CPU time for all instances with one vehicle, and for 129 out of the 160 instances for two vehicles. For the instances without proven optimality, the remaining gap is small. From a computational perspective, solving the problem under an OTL policy is much more difficult than under either of the other two policies. To provide a comparison, under the ML and the OU policies and one vehicle, no instance took more than one hour of computing time, and most instances were solved within a few seconds. Under the OTL policy, the average running time becomes much higher and one particular instance required more than eight hours of computing time before optimality was proved. However, this can be expected given the complexity of the problem, but once the OTL levels have been computed, they are fixed and the remaining problem can then be solved relatively easily at an operational level.

Table 1: Summary of the computational results for the IRP under an OTL policy

Vehicles	Instance	Cost	Gap (%)	Time (s)
	low-3	2968.92	0.00	23.08
	low-6	5937.99	0.00	2782.46
$K = 1$	high-3	9234.07	0.00	16.98
	high-6	13060.76	0.00	2077.90
	Average	7375.70	0.00	923.83
	low-3	3309.30	0.01	1626.22
	low-6	7012.81	5.68	25066.33
$K = 2$	high-3	9576.18	0.00	1439.66
	high-6	14115.41	2.76	23726.46
	Average	7988.25	1.58	10106.73

In addition to comparing CPU times, we have assessed the effect of the OTL policy on cost and inventory levels. The results of these experiments are summarized in Table 2. Specifically, we provide average solution cost increases of the OU policy with respect to the ML policy, and of the OTL over the ML policy. We also report the average cost decrease of the OTL policy with respect to OU, and the average ratio of inventory levels under the OTL and OU policies.

Table 2: Comparison of the computational results for the IRP under the OTL, the OU and the ML policies

Vehicles	Instance	% Cost increase OU over ML	% Cost increase OTL over ML	% Cost decrease OTL over OU	Average ratio of inventory levels L_i/C_i
$K = 1$	low-3	12.30	0.00	12.30	0.82
	low-6	3.07	2.48	0.57	0.93
	high-3	4.29	0.04	4.27	0.82
	high-6	1.80	1.21	0.58	0.92
	Average	6.09	0.70	5.39	0.85
$K = 2$	low-3	11.65	0.00	11.64	0.81
	low-6	5.35	4.57	0.75	0.92
	high-3	4.31	0.01	4.30	0.81
	high-6	2.88	2.09	0.77	0.92
	Average	6.53	1.25	5.26	0.85

From Table 2 one can see that the OTL policy costs slightly more than the ML policy but tends to provide larger cost reductions with respect to the OU policy. Over all instances, the average increase of the OTL with respect to the ML is less than 1% for the instances with one vehicle, and 1.25% for the instances with two vehicles. The average replenishment level under an OTL policy goes down by approximately 15% compared with the OU policy. Note that when $K = 2$, the cost increases presented in Table 2 are overestimated since not all instances were solved optimally and the costs associated with the OTL policy correspond to heuristic solutions when the optimum is not known. Similarly, the cost decreases and the inventory level reductions are underestimated.

A deeper analysis of the solutions shows that on instances with a shorter planning horizon, the OTL and ML policies yield about the same results. This is due to the fact that under an ML policy, the inventory at the end of horizon is zero because of the minimization of the inventory costs, and thus the supplier delivers only the amount needed to satisfy the demand. When the planning horizon is short, customers typically receive only one delivery each. It is then possible to set the OTL to the exact quantity needed to satisfy future demand without surplus. The benefits of the OTL really start being felt on longer planning horizons over which more deliveries are made.

6 Conclusions

We have introduced a new inventory replenishment policy, called OTL, within the framework of the IRP. We have developed a mathematical model corresponding to this policy, and we have solved it through an exact branch-and-cut algorithm. Our results have shown that the OTL policy yields less costly solutions than the classical OU policy, at the expense of an increased computational time. Computational results on benchmark instances confirm the managerial interest of the proposed policy both in terms of costs and inventory levels. We hope this study will stimulate further research into the proposed OTL policy.

References

[1] H. Andersson, A. Hoff, M. Christiansen, G. Hasle, and A. Løkketangen. Industrial aspects and literature survey: Combined inventory management and routing. *Computers & Operations Research*, 37(9):1515–1536, 2010.

[2] C. Archetti, L. Bertazzi, G. Laporte, and M. G. Speranza. A branch-and-cut algorithm for a vendor-managed inventory-routing problem. *Transportation Science*, 41(3):382–391, 2007.

[3] K. J. Arrow, T. Harris, and J. Marschak. Optimal inventory policy. *Econometrica*, 19(3):250–272, 1951.

[4] J. F. Bard and A. I. Jarrah. Integrating commercial and residential pickup and delivery networks: A case study. *Omega*, 41(4):706–720, 2013.

[5] A. S. Caplin. The variability of aggregate demand with (S, s) inventory policies. *Econometrica*, 53 (6):1395–1409, 1985.

[6] L. C. Coelho and G. Laporte. The exact solution of several classes of inventory-routing problems. *Computers & Operations Research*, 40(2):558–565, 2013.

[7] L. C. Coelho, J.-F. Cordeau, and G. Laporte. Consistency in multi-vehicle inventory-routing. *Transportation Research Part C: Emerging Technologies*, 24(1):270–287, 2012.

[8] L. C. Coelho, J.-F. Cordeau, and G. Laporte. Thirty years of inventory-routing. Technical Report CIRRELT-2012-52, Montreal, Canada, 2012.

[9] M. Gendreau, G. Laporte, and F. Semet. The covering tour problem. *Operations Research*, 45(4): 568–576, 1997.

[10] M. W. Padberg and G. Rinaldi. A branch-and-cut algorithm for the resolution of large-scale symmetric traveling salesman problems. *SIAM Review*, 33(1):60–100, 1991.

[11] M. F. Uzar and B. Çatay. Distribution planning of bulk lubricants at BP Turkey. *Omega*, 40(6): 870–881, 2012.