DEPARTMENT OF ENGINEERING MANAGEMENT

From storage to shipment: The effect of ignoring inventory when planning routes

Florian Arnold & Kenneth Sörensen

UNIVERSITY OF ANTWERP
Faculty of Applied Economics

City Campus
Prinsstraat 13, B.226
B-2000 Antwerp
Tel. +32 (0)3 265 40 32
Fax +32 (0)3 265 47 99
www.uantwerpen.be

FACULTY OF APPLIED ECONOMICS

DEPARTMENT OF ENGINEERING MANAGEMENT

From storage to shipment:
The effect of ignoring inventory when planning routes

Florian Arnold & Kenneth Sörensen

RESEARCH PAPER 2017-002
JANUARY 2017

University of Antwerp, City Campus, Prinsstraat 13, B-2000 Antwerp, Belgium
Research Administration – room B.226
phone: (32) 3 265 40 32
fax: (32) 3 265 47 99
e-mail: joeri.nys@uantwerpen.be

D/2017/1169/002

From storage to shipment: The effect of ignoring inventory when planning routes

Florian Arnold [1a], Kenneth Sörensen[a]

[a] *University of Antwerp, Departement of Engineering Management, ANT/OR - Operations Research Group*

Abstract

Routing problems have been widely studied, yet, the interdependences between routing and inventory levels at the supplier are not well-explored. In reality, the optimal distribution of different goods from potentially multiple depots to customers depends on the inventory levels in the depots. Customers can only be served if sufficient inventory of the demanded product is available. In this paper, we present a model and a corresponding heuristic to capture inventory constraints in routing problems with multiple depots (MPMDVRPI). As the main contribution, we study the effect that different inventory levels have on the quality of the respective distribution routes. Depending on the number of depots and products, we observe cost increases of the routing between 4% and 30% if inventory levels are sparse. Furthermore, we find that a different allocation of inventory to depots can affect the routing costs by up to 9%.

Keywords: multi-depot vehicle routing problem, multi-product, inventory management, inventory allocation, metaheuristics

1. Introduction

The Vehicle Routing problem (VRP) is one of the most studied problems in the field of Operations Research. This can be attributed to its complexity, generality and practical relevance. Many logistic systems involve the task of planning deliveries from a source (also called depot) to destination points (also called customers), and thereby to solve a VRP to minimize a certain objective (e.g. travel

[1]corresponding author. Email: florian.arnold@uantwerpen.be

time, cost, or externalities). In practice, the deliveries are often made from different depots. The resulting problem has been studied as the multi-depot routing problem (MDVRP) (Montoya-Torres et al., 2015). In this problem, one has to solve the allocation task which customer should be supplied by which depot, as well as to plan the delivery routes. Naturally, a customer can only be supplied from a depot that has the demanded products on stock. Therefore, it might happen that a customer's demand has to be satisfied from a far-away depot, if closer depots do not have sufficient inventory. Inventory levels can thus affect delivery routes. These interdependencies become more complex if we consider different products, rather than one generic product like in most routing problems.

Take the example of a large mail order company with a complex distribution infrastructure, like Amazon, shipping many product, e.g. flashlights and batteries, from many depots. Most customers will buy batteries and flashlights in a bundle, so it makes sense to store both products in the same depots; otherwise the customer needs to be supplied by two different depots which results in less efficient delivery routes. However, in which depots should the products be stored and how much? Is it necessary to store all products in sufficient quantity (whatever this is) in all depots? This would certainly result in optimal delivery routing decisions, but on the downside, it might result in overwhelming inventory holding costs. This perspective cannot be ignored, considering that approximately 33 % of logistics costs can be attributed to holding inventory (Wilson, 2006). On the other hand, an optimal inventory management plan (with minimized holding costs) might lead to expensive delivery routes, because customers potentially have to be delivered from far-away depots.

These questions relate to different topics in the field of logistic optimization. The problem of inventory allocation tries to put the right amount of inventory at the right location (Federgruen and Zipkin, 1984). However, this perspective simplifies the distribution activities, which are optimised with MDVRPs (Montoya-Torres et al., 2015). Reversely, routing problems usually consider inventory as a black box, implicitly assuming that products are available in sufficient quantities. An exception is the class of inventory–routing problems (Moin and Salhi, 2007), which studies vendor-managed inventory, but does not consider inventory levels at the supplier's warehouses. Thus, there seems to be a lack of findings about the interdependencies between decisions on the inventory and on the distribution level. Considering the possible benefits of an integrated decision-making along the logistic chain (Reimann et al., 2014; Chandra and Fisher, 1994), quantitative studies on the interplay between decision levels are necessary.

In this paper, we will remove the inventory black-box assumption in general routing problems and account for the following:

- Multiple products are distinguished and stored in multiple depots.
- Customers can only be served from a depot, if the demanded products are stocked and available in sufficient quantity.

We incorporate inventory levels for multiple products in MDVRPs and formulate a novel mathematical optimization problem: the multi-product multi-depot vehicle routing problem with inventory restrictions. To solve the resulting distribution problem we develop an efficient heuristic. Hereby, we do no optimise inventory levels, or the allocation of inventory, but rather optimise the routing that results from a given inventory situation. Thus, we embed inventory levels as a restriction in the planning problem. As the main contribution of this paper, we conduct a simulation study with the goal to investigate the effects of different inventory levels on routing costs. In the lines of (Salhi and Rand, 1989), where the authors quantify the effect of ignoring distribution planning in strategic facility location decisions, we investigate

- How do routing costs change if inventory levels are considered?
- How does the number of depots and products affect this relation?

The paper is structured as follows. In Section 2 we introduce the planning problem, and in Section 3 we present an effective heuristic to solve it. Our experimental setup and motivation is outlined in Section 4. Finally, Section 5 discusses the most important findings and Section 6 summarizes our contributions.

2. Problem Description

We extend the basic MDVRP (Montoya-Torres et al., 2015) by introducing inventory levels at the depots and distinguishing between different products. The resulting problem is called the multi-product multi-depot vehicle routing problem with inventory restrictions (MPMDVRPI). The MPMDVRPI is defined as follows: given a set of customers with known demands (for several products), given several depots, each with their stock levels defined, find the lowest-cost routes (starting and ending at one of the depots) for a fleet of vehicles such that all demands are satisfied, the capacities of the vehicles are not exceeded, and the total demand for each product served from each depot does not exceed the inventory of that product at the depot. The integration of inventory levels and the consideration of different products in the MDVRP is modelled by additional constraints.

More formally, let us consider a routing problem with N customers. Each customer $i \in \{1, \ldots, N\}$ has a demand d_{ip} for each product $p \in \{1, \ldots, P\}$, which might potentially be 0. The demand can be supplied from any depot $j \in \{1, \ldots, M\}$, but only if depot j has sufficient inventory level I_{jp} to cover the customer's demand for product p. We denote the allocation of customers to depots by x_{ij}, which is 1 if customer i is supplied by depot j and else 0. With these variables we could model the inventory constraint as follows:

$$\sum_{i=1}^{N} d_{ip} x_{ij} \leq I_{jp} \qquad \forall j = 1 \ldots M, \;\; p = 1 \ldots P \tag{1}$$

However, it might happen that a customer orders two products p_1 and p_2, and no depot has sufficient inventory to deliver both products simultaneously. In this case, the customer needs to be supplied by two different depots, e.g., p_1 is delivered by depot d_1 and p_2 by d_2. Routing problems rarely model this possibility of visiting a customer more than once. We model it by allocating a customer's demand to a depot, rather than the customer itself. More specifically, let x_{ijp} denote whether the demand of customer i for product p is satisfied by depot j, then

$$\sum_{i=1}^{N} d_{ip} x_{ijp} \leq I_{jp} \qquad \forall j = 1 \ldots M, \;\; p = 1 \ldots P. \tag{2}$$

Hereby, a customer's demand for one specific product is satisfied by one tour only, and cannot be split among depots. Additionally, we require that the inventory levels are sufficient to satisfy all demands. As in the MDVRP we look for routes that minimize the total distance traveled. Thereby, as in traditional routing problems, each vehicle starts and ends its delivery routes at the same depot, and has a certain capacity limit. The complete mixed-integer programming problem can be found in the appendix.

3. Design of a heuristic for the MPMDVRPI

The MPMDVRPI is a special case of the MDVRP, and, thus, it is an NP-hard problem for which commercial MIP solvers can find the optimal solution only for small instances within a reasonable time. In the last decades, heuristics have been proven to be a more tractable approach to obtain high-quality solutions for instances of realistic size (Laporte, 2009). Therefore, we introduce a heuristic which can solve MPMDVRPI instances in a short time. The heuristic extends a previously-developed algorithm for the MDVRP.

3.1. A heuristic for the MDVRP

In the following, we present a brief outline of our MDVRP heuristic, and for more details we refer the interested reader to (Arnold and Sörensen, 2016). The heuristic is based upon variable neighborhood search (VNS) (Mladenović and Hansen, 1997), one of the most successful metaheuristic frameworks for various types of routing problems (Sörensen et al., 2008). A VNS alternates between a local search stage and a perturbation stage.

In the local search stage, different local search operators are called iteratively. These operators try to improve the current solution by triggering small local changes (e.g., change the route position of a customer or swap two customers). We use an ejection chain (Glover, 1996), the CROSS-exchange operator (Taillard et al., 1997), and the Lin-Kernighan heuristic (Lin and Kernighan, 1973) as local search operators. While the first two operators look for changes between routes, the latter one tries to keep routes optimal in themselves.

Our perturbation stage introduces small but controlled – or guided – changes to the current solution. This concept is based on guided local search (GLS) (Voudouris and Tsang, 2003). GLS tries to remove those edges in a solution that are classfied as 'bad', by adding penalties to their cost value. A self-evident idea is to penalize the longest edges as in (Mester and Bräysy, 2007). In a data-mining study (**?**) we investigated other characteristics of 'bad' edges and condensed the results in a penalty-function that measures that 'badness' of each edge. The edge with the highest function value is penalized and potentially removed in the succeeding local search stage. These two steps of local search and edge penalization are then iterated thousands of times.

An outline of the heuristic is presented in Algorithm 1. It performs comparable to the best heuristics in literature for a wide range of benchmark instances. More specifically, it solves the MDVRP instance benchmark from Cordeau with an average gap to the best known solutions of 0.28 % in less than a minute computation time.

Algorithm 1 Outline of the MDVRP heuristic

```
Solution S ← Clark-Wright
while not stopping condition do
    Local search:
    while improvements found do
        S ← Ejection chain
        S ← Lin-Kernighan
        S ← CROSS-exchange
        S ← Lin-Kernighan
    end while
    Perturbation:
    Penalize edge in S with highest value in penalty function
end while
```

3.2. Extension of MDVRP to MPMDVRPI

We integrate inventory levels and product distinction into the heuristic above. The simplicity of the heuristic design allows to incorporate these additional restrictions with only a few adaptations.

The starting solution is generated in a greedy fashion, by firstly allocating customers to depots and then computing routes with the Clark and Wright heuristic. For each customer we compute the difference between the distances to the second-closest and the closest depot, and store these potential savings in a list. We sort the list in descending order, and starting from the top, iteratively allocate a customer to its closest depot, if it still has sufficient inventory. Otherwise, we recompute all savings as the difference in the distances between the two closest depots with sufficient inventory, and reinserted them into the list. In this way we allocate each customer to a depot. This approach, also called regret heuristic, minimizes the opportunity cost that result when the closest depot does not have sufficient inventory to serve the customer. Finally, we compute, for each depot separately, initial routes with the Clark and Wright heuristic.

Multiple products are distinguished by assigning a demand value for each customer and product. The used capacity of a vehicle is then computed by aggregating the demand per product over all visited customers. If a customer has demand for more than one product, he might be delivered by two different depots as noted above. We handle this special case by defining a new customer for each additional ordered product, e.g. if a customer demands two products, we split it into two customers (with the same location) that both order one product. Even though this approach might increase the number of customers and, therefore, the problem complexity, it can be readily applied in the heuristic above.

Likewise, we consider inventory restrictions by assigning inventory levels to each depot. The restrictions themselves are validated within the local search operators. Our heuristic maintains feasible solutions at each stage of the search, and the local search operators only execute those changes that satisfy all restrictions. For each depot and product, we keep track of the sum of demand over all customers that are assigned to routes from this depot. Thus, at each stage of the heuristic we know how much inventory is available in each depot. Only then can customers from other depots be relocated or swapped into a route from this depot, if there is still sufficient inventory available.

4. Experimental setup

The motivation behind the following analysis is that the theoretical best routing plan is not always feasible in practice when considering inventory levels. Given multiple depots and multiple products, inventory can be either sparse or badly allocated and, thus, result in non-optimal routing decisions. Therefore, our hypothesis is that tighter inventory constraints (we will define below what we mean with this) worsen the quality of the routing. The goal of our analysis is to quantify these dependencies. Firstly, we motivate why non-optimal inventory levels can occur in practice, and then we describe our methodological design. We use the term 'non-optimal inventory levels' in the sense that they do not allow optimal delivery routes.

4.1. Practical reasons for non-optimal inventory levels

Traditionally, inventory management and distribution planning are two independent problems in logistic optimisation. With vertical supply chain integration on the rise, the interdependencies between both entities are increasingly more explored. In practice however, inventory levels might still interfere with the realization of optimal distribution plans for the following reason.

- Inventory and distribution are planned independently of each other, or there is a lack of information exchange between both planning levels. This happens when the logistic chain is not sufficiently vertically integrated.

- Inventory planning is usually a tactical decision that is taken over a period of several days or weeks. In contrast, distribution decisions are operational and taken on a day-by-day basis. Synchronization between both decision levels can therefore be difficult, since the inventory plan needs to anticipate the demand and distribution for a longer time period.
- Inventory levels fluctuate and are generally not accurately predictable. Lead times are usually stochastic, customers can cancel or change orders and products can break and spoil. These dynamic components make it difficult to optimise inventory levels on their own, let alone the synchronization with distribution.
- Inventory holding is expensive (capital lockup, storage costs, costs of spoilage, depot costs,…). Thus, even though most companies keep a safety stock, the amount of stored inventory will be kept at a minimum.
- Not all products might be stored in all depots, for instance, because of special storage conditions for a product. We will investigate related effects in Section 5.2.

In summary, inventory plans might not always be synchronized with anticipated distribution plans, and even if they are, inventory levels are usually low and subject to stochastic effects. Consequently, it might not always be possible to select the best routing plan since it is infeasible with respect to the available inventory in the depots.

4.2. Experimental setup

We conduct a simulation study with the goal to quantify the effect that inventory restrictions have on routing costs. The outline of the study is as follows: (1) we generate random MPMDVRPI instances, (2) for each instance we vary inventory levels, and (3) compute high-quality routing solutions. By comparing the results for different inventory levels per instance we can then derive interdependencies between routing and inventory levels.

We conduct different experiments in which we vary the number of depots $M \in \{2, 3, 4\}$ and products $P \in \{1, 2, 3\}$. In each experiment we generate 100 MPMDVRPI instances with with $N = 80$ customers (this is small enough to be computational feasible and large enough to be realistic). Both customers and depots are located randomly on a squared plane. We set the capacity limit of vehicles in such a way that a solution has 9 routes on average. For each customer

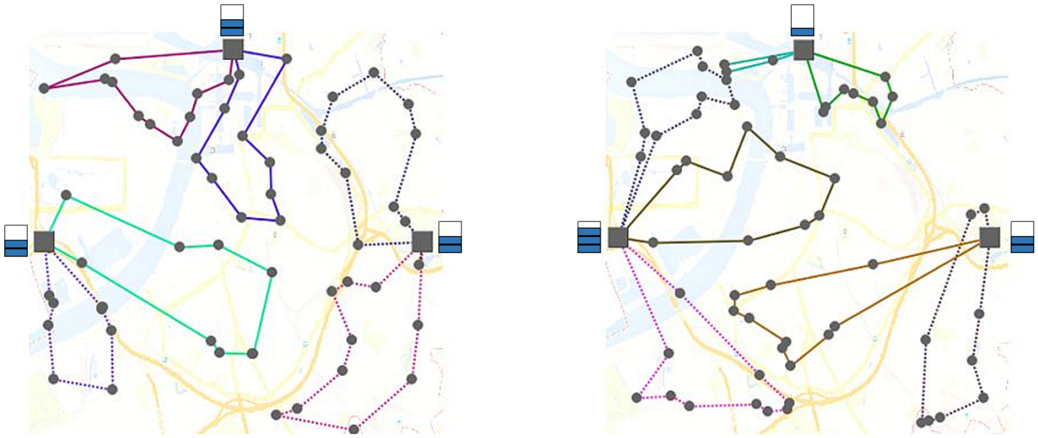

Figure 1: *Solutions for MPMDVRPI instances with different inventory allocation. (left-hand-side) Inventory levels allow an almost optimal routing solution. (right-hand-side) Inventory allocation results in non-optimal routes.*

i we assign a random demand $d_{ip} = \{1, \ldots, 10\}$ for one or two products p; In the case $P = 1$ we can only assign demand for a single product, whereas in the cases $P = 2$ and $P = 3$ we assume that 50% of the customers have demand for a single random product and 50% for two random products.

The inventory levels in the depots can be varied in two ways. We can vary the total amount of inventory available in all depots $I_p^g = \sum_j I_j p$ (global constraint), and we can vary how this total amount is allocated amongst the depots, by changing $I_j p$ (local constraint). Both changes affect the routing solutions. If there is more inventory globally available, the routing solutions should be better, however, if the available inventory is not well-allocated (e.g. in the extreme case that one depot has all inventory), the solutions become worse. Fig. 1 visualizes this interaction. On the other hand, if inventory is optimally allocated, i.e. according to the optimal routing solution, inventory levels only need to cover customer demand. We want to remark again that we do not optimise inventory allocation, but rather study the effect that a different allocation has on routing costs.

More formally, let $D_p = \sum_{i=1}^{N} d_{ip}$ be the sum of demand over all customers for a certain product p. If $I_p^g = D_p$ for all products p, then the available inventory is just sufficient to satisfy the demand, and if $I_p^g > D_p$ we have some inventory to spare. In the experiments we investigate $D_p \leq I_p^g \leq 2D_p$ (global constraint). For the allocation of inventory to depots (local constraint) we consider three

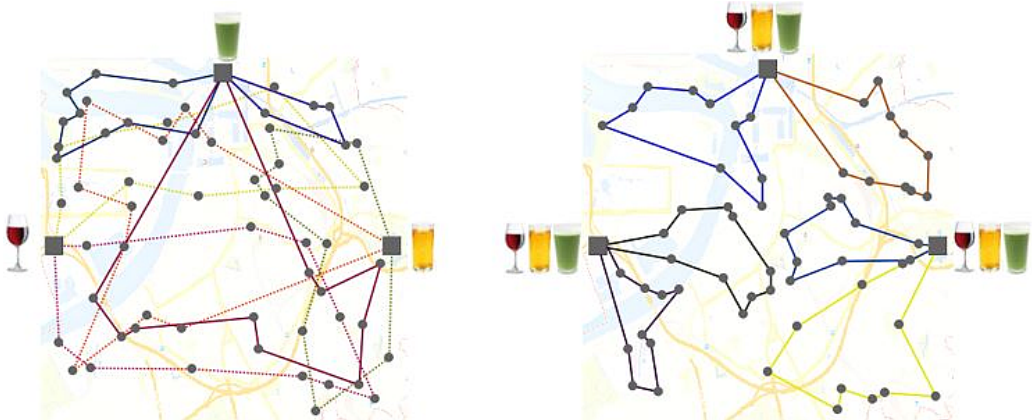

Figure 2: *Example of the effect of not storing all products in all depots. (left-hand-side) Three products are delivered from three different depots. (right-hand-side) Each depot stocks all products.*

scenarios. (1) In the first scenario, the inventory is allocated equally, i.e. $I_{jp} = \frac{I_p^g}{M}$. In the other scenarios, we consider a better and a worse allocation. For each depot j, we count the number of customers c_j who are closer to this depot than to any other. (2) In the better allocation scenario, we set $I_{j_p} = \frac{c_j I_p^g}{N}$, so the inventory levels increase linearly with the number of closest customers. (3) Likewise, in the worse allocation scenario we inverse this value, so that the inventory level in a depot decreases if it has more closer customers.

Note that these allocation scenarios imply that all products are stocked in all depots. Moreover, all products are allocated amongst the depots according to the same distribution, i.e. each depot stocks the same percentage of the total amount of each product. We have explained above that this assumption does not necessarily hold, and, thus, in a second set of experiments we consider the scenario that not all products are available in all depots. Fig. 2 illustrates an example.

In summary, our methodological approach is the following. For each experimental setup, defined by the number of depots and products, we generate 100 random instances. For each instance, we generate 15 different inventory scenarios (five different global inventory levels from D_p to $2D_p$, paired with three allocation scenarios). Finally, for each inventory scenario we compute a routing solution s^I with the heuristic described in Section 3. Additionally, we compute a (optimal) routing solution s^O in which inventory constraints are neglected, and

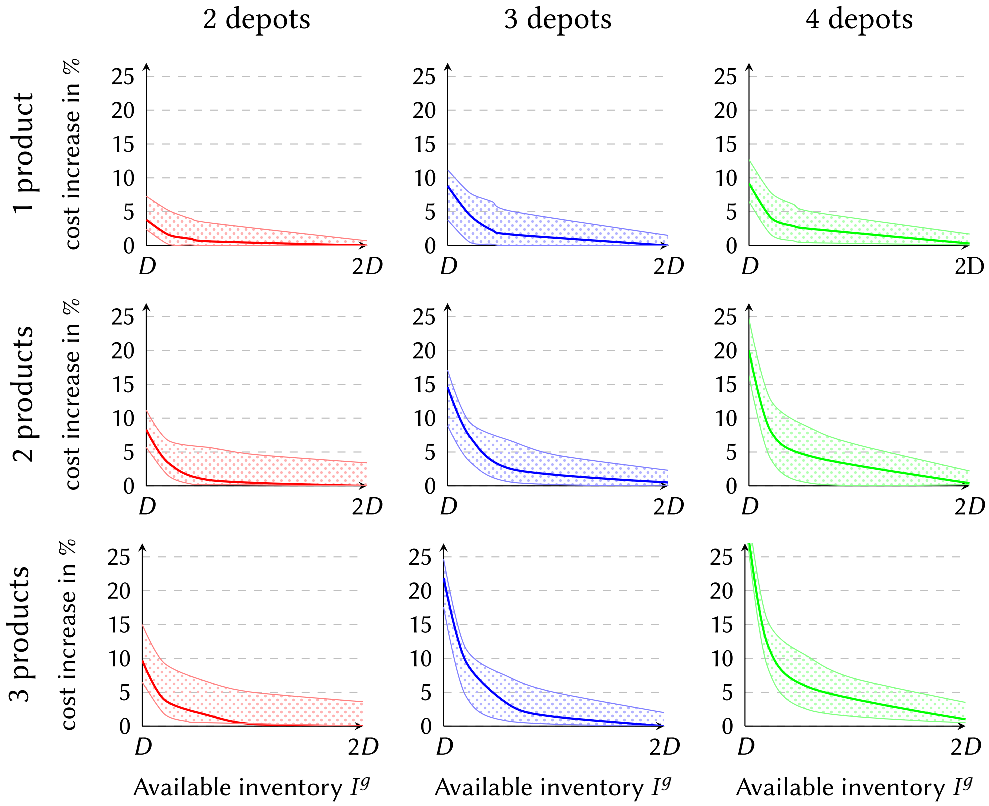

Figure 3: *Impact of different inventory levels on routing costs for different number of depots and products.*

which should therefore be always equally good or better. We then compute for each inventory scenario the additional routing costs $\frac{s^I}{s^O}$ with respect to the unrestricted solution, and average this cost increase over all 100 instances.

5. Experimental Results

5.1. All depots stock all products

The results are visualized in nine graphs in Fig. 3. Each graph presents one experimental setup in terms of the number of depots and products. While the horizontal axis depicts the amount of globally available inventory I^g, the vertical axis highlights the average cost increase with respect to the non-restricted solutions. Note that for all products the ratio $\frac{I^g_p}{D_p}$ is always the same, and, hence, we drop the index p and simply write I^g . Each line presents a different allocation scenario. The thick line depicts the equal-allocation scenario, whereas the lines above and below present the worse and better allocation scenario, respectively. The corridor between the better and worse allocation scenario reflect the impact – or sensitivity – of inventory allocation on routing costs. We can make three major observations from these results:

(1) If there is sufficient inventory globally available, inventory constraints have only a minor impact on routing costs, independent of the allocation scenario. However, if the global inventory levels fall below a certain threshold, the additional routing costs increase significantly. This threshold seems to be around $I_p^g = 1.2D_p$ and can be observed in all experimental setups, i.e. the number of products and depots. Hereby, the experimental setup moderates the magnitude of the increase in additional routing costs, which reach between 4% and 33% for $I_p^g = D_p$.

(2) With more depots, inventory constraints have a higher impact on routing solutions. More precisely, the additional routing costs due to inventory constraints in MPMDVRPIs are higher, if the problem considers more depots. Intuitively, with more depots available inventory becomes more scattered, i.e. all depots stock a little of everything. This scattering of inventory results in a more inflexible routing, since not all depots have sufficient inventory to realize the most cost-efficient routes.

(3) Analogous to the previous observation, inventory constraints also have a higher impact on routing if more products are considered. The possible explanation is also quite similar. With more considered products, the number of inventory constraints grows as well (one for each product and depot). Overall, the problem becomes constrained and, thus, solutions tend to become worse.

(4) The way that globally available inventory is allocated among the depots impacts the increase in routing costs by up to 9%. More precisely, for $I_p^g = D_p$ the cost increase in the worse allocation scenario is between 4% (one product, two depots) and 9% (two products, three depots) higher than in the better allocation scenario. Surprisingly, this difference in cost increases is quite stable for higher values $I_p^g > D_p$. In other words, the width of the 'cost increase corridor' in the graphs remains relatively constant.

The main implication of these results is that routing costs increase significantly, if inventory is sparse and not well-allocated. A possible quick fix to this problem is to simply store more inventory in the depots. However, higher inventory levels come at the expense of higher holding costs, even if they might result in better routing decisions. This results in a trade-off between inventory

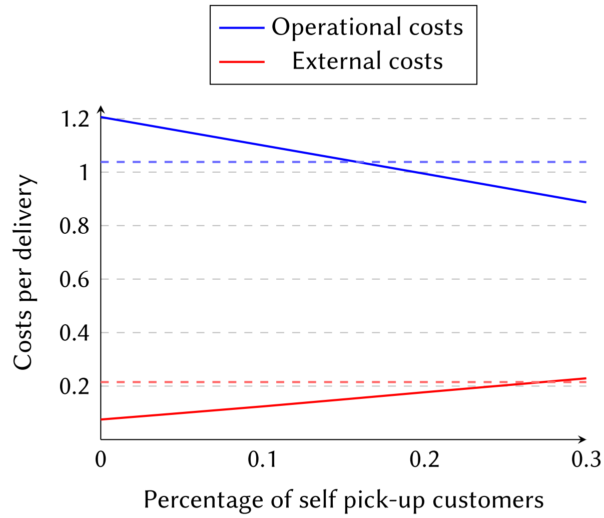

Figure 4: *Routing and inventory holding costs as a function of global inventory* I_p^g *for* $M = 3$ *and* $P = 2$ *(left-hand-side), as well as the sum of both costs (right-hand-side).*

and routing costs.

We analyse this cost trade-off using the example of the setup three depots and two products. Assuming that holding costs decrease linearly in the amount of stored inventory, we can compute the sum of storage and routing costs as a function of available inventory and in dependece of allocation. The results of such a hypothetical computation are visualized in Fig. 4.

The joint costs of distribution and inventory processes could then be minimized, by choosing those global inventory levels that minimize this function for a certain allocation scenario. However, the exact inventory levels, and thereby the allocation, are difficult to forecast due to stochastic effects. One could account for this uncertainty by weighting the cost values for certain inventory levels with the likelihood that this inventory levels occur. This example should illustrate that inventory and routing decisions are heavily intertwined and require a careful analysis.

5.2. Not all depots stock all products

In Section 5.1 we assumed that each depot has all products in stock. With the following analysis we investigates the effects in case that this assumption does not hold. Practical reasons for why a certain depot does not stock all offered products are the following:

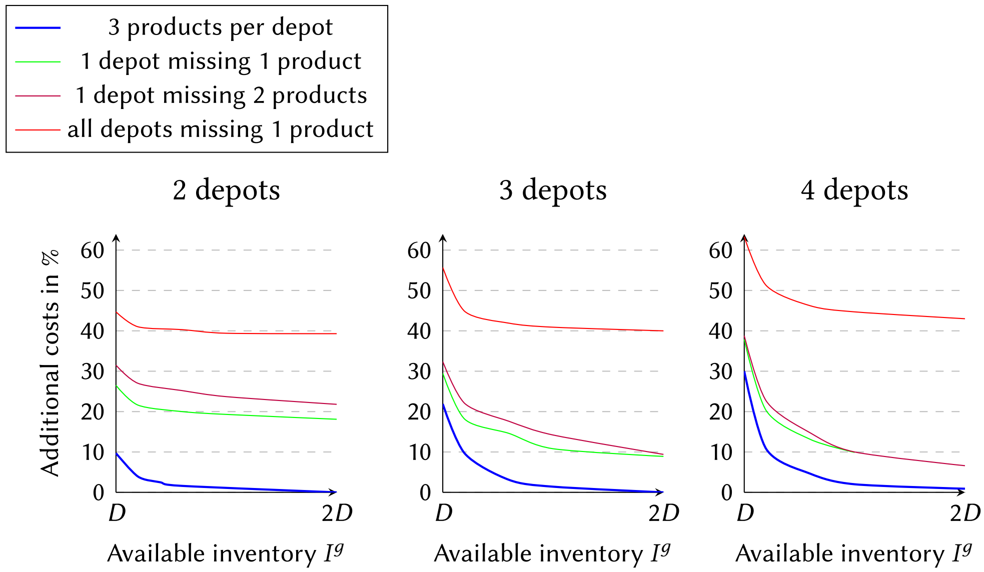

Figure 5: *Additional routing costs when not all of the three products are stocked in all depots, and the available inventory is distributed equally among depots.*

1) Certain products require specialized storage conditions. For instance, many food products need to be cooled and require a cooling box, and building material such as poles might require long corridors. Thus, products can only be stored in a depot that meets the requirements.

2) It might be cheaper to stock only certain products in some depots, for instance because of supply costs or storage planning.

More formally, let s_j denote the number of products that are stocked in depot j. If $s_j = P$ for all depots j, then we have the same design as in Section 5.1. We analyse the effect of setting $s_j < P$ for one or more depots, choosing $P = 3$. Thus, one experimental setup is defined by a tuple $(s_1, s_2, \ldots, s_M)$ where each depot j is assigned s_j random products (under the condition that each product is stocked at least once). For each tuple we compute the additional routing costs for various global inventory levels as above, and, for simplicity, we only consider the allocation scenario where each product is distributed equally among all depots that stock it. We consider the tuples (1,3,3,. . .), i.e. one depot stocks only one random product and the other depots stock all products, (2,3,3,. . .), and (2,2,. . .) and compare it with the case in Section 5.1 (3,3,. . .).

Fig. 5 depicts the results. If only one depot does not stock all products, the additional cost increase due to inventory constraints is about 7% (4 depots) to 16% (2 depots) higher, independent of the amount of global inventory. Hereby, it does not matter much, whether the depot's stock is missing one or two products. Finally, we observe a drastic increase in additional costs if all depots are not stocking all products, reaching additional costs of 50% and more.

6. Conclusions and future research

In this paper we extended the classical MDVRP to consider multiple products and inventory limitations and formulated the resulting planning problem MPMDVRPI. We further introduced an effective heuristic to solve instances of this problem type. At the heart of this work, we analysed the impact that different inventory constraints have on routing costs. Depending on the number of considered products and depots, we observed cost increases in routing between 4% and 30%. Hereby, both the amount of inventory that is available globally as well as its allocation to the depots play a crucial role. Especially when inventory is sparse or badly allocated do inventory levels affect routing costs significantly. In our experiments, differences in allocation amounted to up to 9% in additional routing costs. Finally, we found that routing costs increase drastically by up to 60% if not all products are stocked in all depots.

The experiments are meant to produce generalizable results in the sense that we averaged over many different instances, rather than investigating a particular case study. Naturally, given a particular case study, these numbers might not hold anymore, however, our results give some intuition as to how inventory constraints impact routing plans in general. Furthermore, we assumed that all considered products are demanded in equal measure. In reality, this assumption might not hold, and it might be interesting to investigate whether different effects can be observed when, for instance, one product is rarely asked for.

In conclusion, these results are some of the first to shed light on the interdependencies between inventory management and distribution activities. They can be used to further develop insights into optimal vertically-integrated logistic chains. In this respect, one of the most interesting questions is as to how inventory processes and routing can be optimised jointly. The presented work lays the foundation by quantifying the effects that a certain inventory situation has on routing. A next step could be to shift the focus towards the inventory level, and optimise inventory decisions (restock policy, safety stocks,…) in such a way that the costs of both inventory and corresponding routing is minimal.

F. Arnold and K. Sörensen. A knowledge based local search heuristic for routing problems. *Working paper, University of Antwerp*, 1(1):1, 2016.

Pankaj Chandra and Marshall L Fisher. Coordination of production and distribution planning. *European Journal of Operational Research*, 72(3):503–517, 1994.

Awi Federgruen and Paul Zipkin. A combined vehicle routing and inventory allocation problem. *Operations Research*, 32(5):1019–1037, 1984.

Fred Glover. Ejection chains, reference structures and alternating path methods for traveling salesman problems. *Discrete Applied Mathematics*, 65(1):223–253, 1996.

Gilbert Laporte. Fifty years of vehicle routing. *Transportation Science*, 43(4):408–416, 2009.

Shen Lin and Brian W Kernighan. An effective heuristic algorithm for the traveling-salesman problem. *Operations research*, 21(2):498–516, 1973.

David Mester and Olli Bräysy. Active-guided evolution strategies for large-scale capacitated vehicle routing problems. *Computers & Operations Research*, 34 (10):2964–2975, 2007.

Nenad Mladenović and Pierre Hansen. Variable neighborhood search. *Computers & Operations Research*, 24(11):1097–1100, 1997.

N.H. Moin and S. Salhi. Inventory routing problems: a logistical overview. *Journal of the Operational Research Society*, 58(9):1185–1194, 2007.

J. Montoya-Torres, J.L. Franco, S.N. Isaza, H.F. Jiménez, and N. Herazo-Padilla. A literature review on the vehicle routing problem with multiple depots. *Computers & Industrial Engineering*, 79:115–129, 2015.

Marc Reimann, Roberto Tavares Neto, and Elisabeth Bogendorfer. Joint optimization of production planning and vehicle routing problems: a review of existing strategies. *Pesquisa Operacional*, 34(2):189–214, 2014.

Said Salhi and Graham K Rand. The effect of ignoring routes when locating depots. *European journal of operational research*, 39(2):150–156, 1989.

Kenneth Sörensen, Marc Sevaux, and Patrick Schittekat. fimultiple neighbourhoodfi search in commercial vrp packages: Evolving towards self-adaptive methods. In *Adaptive and multilevel metaheuristics*, pages 239–253. Springer, 2008.

Éric Taillard, Philippe Badeau, Michel Gendreau, François Guertin, and Jean-Yves Potvin. A tabu search heuristic for the vehicle routing problem with soft time windows. *Transportation science*, 31(2):170–186, 1997.

Christos Voudouris and Edward PK Tsang. *Guided local search.* Springer, 2003.

R Wilson. 17th annual state of logistics report. *Council of Supply Chain Management Professionals*, 17(1), 2006.

V_c	$\{v_1, v_2, ..., v_N\}$: Set of customers
V_d	$\{v_{N+1}, v_{N+2}, ..., v_{N+M}\}$: Set of depots
V	$V_c \bigcup V_d$: Set of nodes
P	$\{p_1, p_2, ..., p_P\}$: Set of products
K	Number of vehicles
C	Capacity limit of each vehicle
d_{ip}	Demand of customer i for product p
I_{jp}	Inventory level of product p at depot j
c_{ij}	Travel cost between vertex i and vertex j (symmetric)

$$x_{ijk} = \begin{cases} 1 & \textit{if a visit to vertex } i \textit{ is followed by a visit to vertex } j \textit{ in the tour of vehicle } k \\ 0 & \textit{otherwise} \end{cases}$$

$$s_{ikp} = \begin{cases} 1 & \textit{if vehicle } k \textit{ delivers product } p \textit{ to customer } i \\ 0 & \textit{otherwise} \end{cases}$$

$$a_{jk} = \begin{cases} 1 & \textit{if vehicle } k \textit{ is allocated to depot } j \\ 0 & \textit{otherwise} \end{cases}$$

$$\min \left[\sum_{i=1}^{N+M} \sum_{j=1}^{N+M} \sum_{k=1}^{K} c_{ij} x_{ijk} \right] \tag{.1}$$

Subject to

1.) Tour continuation

$$\sum_{i=1}^{N+M} x_{imk} = \sum_{j=1}^{N+M} x_{mjk} \qquad \forall m = 1 \dots N+M, \forall k = 1 \dots K \tag{.2}$$

2.) Tours start and end at a depot

$$\sum_{i=N+1}^{N+M} \sum_{j=1}^{N} x_{ijk} = 1 \qquad \forall k = 1 \dots K \tag{.3}$$

$$\sum_{i=1}^{N} \sum_{j=N+1}^{N+M} x_{ijk} = 1 \qquad \forall k = 1 \dots K \tag{.4}$$

3.) Synchronization of decision variables

$$\sum_{i=1}^{N} x_{ijk} = a_{jk} \qquad \forall j = N \dots N+M, k = 1 \dots K \tag{.5}$$

$$\sum_{p=1}^{P} s_{ikp} \leq P \sum_{j=1}^{N+M} x_{jik} \qquad \forall i = 1 \dots N, k = 1 \dots K \tag{.6}$$

$$\sum_{p=1}^{P} s_{ikp} \geq \sum_{j=1}^{N+M} x_{jik} \qquad \forall i = 1 \dots N, k = 1 \dots K \tag{.7}$$

$$\mathbb{I}_{d_{ip}>0} = \sum_{k=1}^{K} s_{ikp} \qquad \forall i = 1 \dots, N, p = 1 \dots P \tag{.8}$$

4.) Vehicle capacity constraint

$$\sum_{i=1}^{N} \sum_{p=1}^{P} d_{ip} s_{ikp} \leq C \qquad \forall k = 1 \dots K \tag{.9}$$

5.) Inventory restrictions

$$\sum_{k=1}^{K} a_{jk} \sum_{i=1}^{N} d_{ip} s_{ikp} \leq I_{jp} \qquad \forall j = N+1 \dots N+M, p = 1 \dots P \tag{.10}$$

$$x_{ijk}, s_{ikp}, a_{jk} \in \{0, 1\} \tag{.11}$$