

Volling, Thomas; Grunewald, Martin; Spengler, Thomas S.

## Article

# An Integrated Inventory-Transportation System with Periodic Pick-Ups and Leveled Replenishment

BuR - Business Research

## Provided in Cooperation with:

VHB - Verband der Hochschullehrer für Betriebswirtschaft, German Academic Association of Business Research

*Suggested Citation:* Volling, Thomas; Grunewald, Martin; Spengler, Thomas S. (2013) : An Integrated Inventory-Transportation System with Periodic Pick-Ups and Leveled Replenishment, BuR - Business Research, ISSN 1866-8658, VHB - Verband der Hochschullehrer für Betriebswirtschaft, German Academic Association of Business Research, Göttingen, Vol. 6, Iss. 2, pp. 173-194, <https://doi.org/10.1007/BF03342748>

This Version is available at:

<https://hdl.handle.net/10419/103725>

### Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

### Terms of use:

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*



# An Integrated Inventory-Transportation System with Periodic Pick-Ups and Leveled Replenishment

Thomas Volling, Institute of Automotive Management and Industrial Production, Technische Universität Braunschweig, Germany,  
E-Mail: t.volling@tu-braunschweig.de

Martin Grunewald, Institute of Automotive Management and Industrial Production, Technische Universität Braunschweig, Germany,  
E-Mail: m.grunewald@tu-braunschweig.de

Thomas S. Spengler, Institute of Automotive Management and Industrial Production, Technische Universität Braunschweig, Germany,  
E-Mail: t.spengler@tu-braunschweig.de

## Abstract

In this paper we develop a combined inventory-transportation system. The general idea is to integrate a simple replenishment policy with a routing component to derive operationally consistent standard routes as a basis for milk run design. The most interesting feature of the approach is that we combine stochastic vehicle routing with a replenishment policy which makes use of inventory to level the variability propagated into transportation operations. To evaluate the approach, we compare its performance with stochastic vehicle routing as well as sequential vehicle routing and replenishment planning. With respect to these approaches, substantial gains are achieved.

JEL-classification: C61, M11

Keywords: inbound logistics, milk run design, inventory routing

Manuscript received June 8, 2012, accepted by Karl Inderfurth (Operations and Information Systems) April 14, 2013.

## 1 Introduction

The inbound logistics function is concerned with the provision of parts sourced from geographically dispersed suppliers to facilitate a planned production program. From the perspective of manufacturers, inbound logistics operations result in the material collection problem (Qu, Bookbinder, and Iyogun 1999). As a result of this problem, the following question must be answered: How should the required material be supplied to minimize costs?

In analogy with vendor-managed inventory concepts, which have been widely implemented to improve distribution operations, there is currently a trend in inbound logistics of manufacturers taking responsibility for the planning of transportation operations (Baudin 2004). The objective is to make use of synergies by concentrating all replenishment and transportation decisions at a single centralized point in the supply chain (Disney, Potter, and Gardner 2003; Andersson, Hoff, Christiansen, Hasle,

and Løkketangen 2010). In contrast to distribution logistics, where the divergent structure of the material flow suggests concentrating all decisions at suppliers serving multiple customers, we typically face convergent structures in inbound logistics. To make best use of synergies, it is thus reasonable to concentrate decisions at the manufacturer.

One of the most widespread concepts used in implementing centrally coordinated inbound logistics are milk runs. The concept of milk runs adopts the principles of lean production for transportation operations. The idea is to define standard routes, which are repeated on a regular basis (e.g., daily or twice a week) for a defined planning cycle (e.g., 3 weeks). Every standard route consolidates replenishment quantities from a cluster of suppliers in a geographically concentrated region and defines a periodic pick-up schedule for each supplier (Chuah and Yingling 2005). Compared to directly shipping material from each supplier, this may result in sig-



nificantly reduced distances travelled and fixed transportation costs, in particular if medium demand items are regarded. Milk runs may thus facilitate the frequent shipping of less-than-full-truckload volumes and, as a consequence, contribute towards reduced lead times. The underlying effect is referred to as economies of proximity (Larson 1988). By defining standards for pick-up and routing, milk runs additionally provide logistic providers and suppliers with an improved basis for planning. This allows for a series of process improvements which should, in theory, translate into reduced transportation and procurement costs for the manufacturer (Gaur and Fisher 2004).

In implementing milk runs, manufacturers struggle with two important challenges. A first challenge concerns the definition of milk runs which are to be operated within a mid-term planning horizon. This relates to the questions of how to consolidate the replenishment quantities of multiple parts into capacitated standard routes and how to define frequencies for these routes. A second challenge is related to the choice of replenishment policies. These convert the material requirements derived from market demand for final products, i.e., the secondary demand, into replenishment orders for every single part and, therefore, generate the input for the definition of milk runs.

Both decisions are closely interrelated. Given the mid-term planning horizon, material requirements are typically not known with certainty. In order to hedge against random swings in material requirements, there are two fundamental principles. A first principle is to synchronize replenishment quantities with material requirements and to install additional transportation capacity into the milk runs. As a consequence, milk runs remain operationally consistent, even if replenishment quantities exceed expectations. Yet, by introducing additional capacity into the milk runs, the average vehicle utilization is reduced. This results in the need for additional milk runs and vehicles respectively. In addition, unfavorable routes and thus additional distances travelled may result. Both effects lead to an increase in transportation costs and greenhouse gas emissions.

A second principle is to introduce inventory to eliminate the variability propagated into the transportation system. If the required replenishment quantity is larger than the expectation, inventory is used to reduce transportation demand. If the required replenishment quantity is low, inventory is recovered.

In doing so, using inventory allows for decoupling replenishment quantities from part requirements. While resulting in additional holding costs, there is a chance to define efficient milk runs and reduce transportation costs respectively.

Against this background, we introduce an inventory-transportation system which combines both principles in the following. The idea is to solve the trade-off between transportation and holding costs by integrating a replenishment policy with leveling property into a stochastic vehicle routing model. The replenishment policy makes use of inventory to control the variability propagated into transportation operations. An approach is presented to simultaneously determine milk runs and part-specific parameters of the leveling replenishment policy.

The contribution of the paper is fourfold: (1) We provide a new modeling approach for the integration of vehicle routing and replenishment. (2) To solve the resulting non-linear stochastic model we propose an approximate deterministic reformulation for normally distributed material requirements and develop an exact decomposition approach to solve the reformulated model. (3) We provide a numerical analysis to compare our approach with two other approaches presented in the literature (stochastic vehicle routing and vehicle routing with safety stock). Based on the results of the numerical analysis we (4) derive operating conditions for the approach.

The remainder of the paper is organized as follows: After the review of prior works on the topic in section 2 we introduce the generic model formulation in section 3. A solution approach, which transforms the stochastic model into an approximate deterministic model, is given in section 4. We conclude with a numerical study in section 5 and some concluding remarks in section 6.

## 2 Literature review

While there is a rich body of literature on the independent planning tasks of replenishment planning and routing, there is considerably less research devoted to the combined problem. On a high level, two streams of literature can be distinguished: the first has its roots in replenishment planning, while the second has emerged from routing.

Works on **replenishment planning** share the common characteristic that transportation costs (e.g., as a result of milk run planning) are assumed

to be exogenously given. The interface to transportation operations is addressed in two groups of works. The idea of the first group is to reduce the variability of replenishment quantities; for the second group it is to better approximate transportation costs.

Within the first group, works on smoothing inventory control seek to improve the basis for the planning of downstream processes (Cannella and Ciancimino 2010). The principle idea is to constrain the replenishment quantities to a certain share of the factual requirements. The share to be replenished is regulated by proportional controllers. While downstream processes benefit, smoothing replenishment comes at the cost of a deteriorated service level (Chen and Disney 2007). This strongly limits the applicability for the material collection problem (MCP), where high service levels are required to facilitate production.

A similar idea, which overcomes the limitations of reduced service levels, was developed in Ormeci, Dai, and Vate (2008). The inventory balance is modeled using a stochastic Brownian-motion process. To control the process, a control band policy is developed. Control action is taken every time the inventory drops below a given level or exceeds a maximum level. Every control action is associated with fixed and variable costs. These may be interpreted as additional transportation costs. The authors showed that for the assumption of zero replenishment times the control band policy minimizes average holding and inventory adjustment costs. Since routing aspects are not modeled explicitly, the question remains open, to what extent the reduced adjustment costs contribute towards reduced transportation costs. This in particular becomes relevant in supply networks, which are characteristic for the MCP.

A more explicit model of transportation costs is used in the second stream of works on replenishment policies. The idea is to coordinate replenishment decisions across multiple items. The associated planning task is known as (stochastic demand multi-item) joint replenishment problem (JRP) (Khouja and Goyal 2008). Two kinds of transportation costs are distinguished, major and minor costs. While the former are fixed for every shipment, the latter depend on the shipment size. Since the structure of the optimal replenishment policy is generally unknown, most works focus on the development of heuristics. Relevant for the MCP are in particular

works, which consider transportation capacity. In these works policies are developed, which generate replenishment orders either every time the cumulated reorder point is reached or if the inventory level of every single item drops to its reorder point. The remaining capacity of the shipment is then allocated to the remaining items (Miltenburg 1985; Kiesmüller 2010). Both approaches share the common characteristic that full-truck-load shipments are generated. The transportation costs which are used for the JRP are assumed to be independent of replenishment decisions. If routing aspects are disregarded, works of the JRP meet the requirements of the MCP. However, due to the geographically dispersed network of suppliers, there is a strong link with routing (i.e., milk run planning) for the problem considered in this paper. The applicability of JRP is therefore limited.

Works on the **inventory routing problem (IRP)** seek to compute transportation costs endogenously. The idea is to simultaneously determine the quantity and timing of replenishments as well as routes. There are diverse possibilities to structure works on the IRP (Andersson, Hoff, Christiansen, Hasle, and Løkketangen 2010; Moin and Salhi 2007; Bertazzi, Savelsbergh, and Speranza 2008). For the MCP we in particular distinguish works in whether resulting plans are repeated periodically and in the way demand is modeled (deterministic/stochastic).

Works on the classical IRP focus on the determination of integrated replenishment and transportation plans, which minimize costs for a given (typically dynamic) demand. Both, deterministic and stochastic demand models have been proposed. The resulting plans may cover one (Golden, Assad, and Dahl 1984) or multiple periods (e.g., Dror, Ball, and Golden 1985; Bard, Huang, Jaillet, and Dror 1998; Trudeau and Dror 1992; Yu, Chu, Chen and Chu 2012). A repeated execution in the sense of a periodic schedule is not intended. Due to the complexity of the problem, the focus is on the development of heuristics. Most works rely on a two-step procedure. In a first step, information on the demand and the inventory position is used to allocate replenishment items (parts) to days (possibly only the first day, for the case of a single period approach). In a second step, routes are planned for each day of the planning horizon. In recent works also integrated solution approaches have been proposed (e.g., Hemmelmayr, Doerner, Hartl, and Savelsbergh 2010). Works on the IRP generally match well with the

requirements of the MCP. However, there are two important limitations. (1) Given the fixed planning horizon, myopic solutions may result. This is the case if only the most urgent or best-fitting material requirements are included in the schedule, leaving all others for future plans. Possible consequences are infeasibilities if future material requirements exceed capacities, or inefficiencies if opportunities to include future demand in current shipments are neglected. To reduce this effect, planning on rolling horizons was suggested (Campbell and Savelsbergh 2004). More severely, (2) the missing definition of stable mid-term plans countervails opportunities to introduce standards into the MCP, which are required for implementing lean practices such as milk runs.

The missing stability of plans derived from the classical IRP has been a motivation for a second class of works regarding the periodic inventory routing problem (PIRP). These works most suitably reflect the situation of the MCP. The objective is to determine a  $p$ -periodic schedule of tours which may be repeated in the medium to long term, possibly indefinitely (Francis, Smilowitz, and Tzur 2008). In industry this approach is known as common frequency routing (Chuah and Yingling 2005). The structure of the resulting problem resembles that of the periodic vehicle routing problem, PVRP (Beltrami and Bodin 1974; Russell and Igo 1979; Christofides and Beasley 1984).

The MCP was interpreted as PIRP in a series of works. Amongst the first were Chuah and Yingling (2005). The objective is to define periodic milk runs, which minimize transportation, holding and handling costs. Pick-up frequencies (and thus cycle stock) are derived from the solution. In addition to that, time windows are modeled. Deterministic and stationary demand rates are assumed. In order to accommodate for short-term variations, the authors suggested inflating these rates, yet without given details on how to inflate them. A set partitioning, column generation approach is combined with a tabu-search procedure to heuristically generate solutions.

A similar version of the problem is analyzed in Alegre, Laguna, and Pacheco (2007). The objective is to establish periodic milk runs for a comparably long planning horizon of 90 days. In minimizing transportation costs the authors assumed deterministic demand. Since pick-up frequencies are set exogenously, holding costs are disregarded. A two-step

heuristic solution approach is developed: In the first step, pick-up schedules are developed for each material such that frequency requirements are fulfilled. In the second step independent vehicle routing problems are solved for each day. Scatter search is used to identify good solutions.

A third work in the line of research is Stacey, Natarajathinam, and Sox (2007). In contrast to Chuah and Yingling (2005), frequencies are not chosen from a discrete set of values but determined endogenously, based on a heuristic two-step approach. First, possible routes are determined and optimal frequencies are computed for each route. This is done by solving a joint economic lot-sizing problem. In the second step, a set-partitioning problem is solved to select those routes, which minimize total costs.

All three approaches assume deterministic demand. This assumption may be valid for short planning horizons for which market demand and / or material requirements are known with certainty. For mid-term planning horizons considered in this paper, material requirements will most likely vary due to unexpected swings in market demand and production planning decisions at the manufacturer. The outcome of deterministic planning is thus prone to infeasibility and inefficiency.

A stochastic version of the problem was studied in Qu, Bookbinder, and Iyogun (1999). The objective was to derive periodic milk runs which minimize major- and minor-ordering, holding, backlogging, and transportation costs. To this end, the authors integrated a periodic order-up-to policy with a common base period and a separate travelling salesman problem for each period of the schedule. In order to solve the problem, first the inventory (master-) problem is solved. The result is used to compute transportation costs from the traveling salesman (sub-) problems in a second step. Feeding these costs back into the inventory model, the procedure is iterated until a stopping criterion is reached. As such, the model has strong analogies with the uncapacitated JRP, yet extended by endogenously computed transportation costs.

Approaches for the stochastic PIRP which take capacity constraints into account have not been developed for the MCP yet. However, the first attempts in this direction have been presented for the deterministic PIRP. In order to hedge against demand uncertainty, the idea is to either make use of inventory or additional transportation capacity. The usage of



inventory to assure the feasibility of the solution derived from a capacitated PIRP in the light of stochastic demand was discussed in Gaur and Fisher (2004) and Custódio and Oliveira (2006). In the first step, deterministic demand information is used to determine a periodic routing plan. In the second step, initial inventory is set for every part such that a certain cycle service level is obtained for the replenishment frequency which results from the periodic plan. One way to assure sufficient transportation capacity is to inflate demand rates and to solve a deterministic PIRP for these inflated rates (e.g., Chuah and Yingling 2005). A more rigorous approach to compute the required extra capacity is to make use of the stochastic vehicle routing problem (SVRP) (Stewart and Golden 1983). Adoptions to the IRP were proposed by Bell, Dalberto, Fisher, Greenfield, Jaikumar, Kedia, Mack, and Prutzman (1983) and Bard, Huang, Jaillet, and Dror (1998). However, both works consider the case of non-periodic IRP.

From this analysis we conclude that a comprehensive approach for the MCP is missing which takes into account geographically dispersed suppliers, transportation capacities and costs, the definition of operationally consistent standard routes and the stochastic nature of demand.

In the following we combine ideas from smoothing replenishment and stochastic vehicle routing to develop an integrated inventory-transportation system with periodic pick-ups and leveled replenishment.

### 3 Problem formulation

#### 3.1 General concept

In the remainder we formulate and integrate two sub-models to determine (1) part-specific replenishment policies and (2) standard routes (milk runs) which are to be repeated on a regular basis for a defined number of repetitions (in the remainder: one planning cycle). Given information on each part's stochastic material requirements, the objective is to minimize the average expected transportation and holding costs per period while assuring (a) a given cycle service level and (b) a given probability for feasible milk runs.

(1) A leveling replenishment policy is proposed which defines a rule to convert the material requirements of a part (i.e., the secondary demand of

that part) into replenishment quantities. The main property of this policy is that the variability of the replenishment quantities is reduced relatively to material requirements. To set the degree of leveling, a specific parameter is introduced for every part. Depending on the value of that parameter, initial inventory is required to avoid material shortages. We determine the initial inventory level such that the probability of material shortages within the planning cycle is constrained.

(2) A chance-constrained SVRP is developed to assign each part to one milk run. Milk runs are set in a manner such that for each milk run the anticipated cumulated replenishment quantities, which result from applying the leveling replenishment policy, do not exceed transportation capacity with a defined probability.

By integrating both sub-models, we are able to simultaneously determine part-specific replenishment parameters and milk runs which minimize average expected total costs per period.

Please note that our cycle service level concept differs from the replenishment cycle concept that is typically used in the literature. The principle idea of the proposed leveling policy is to partially decouple replenishment quantities from material requirements in order to reduce the variability that is propagated into the supply chain. This decoupling is based on inventory. The initial inventory level is set such that it is possible to apply the leveling replenishment policy for a defined planning cycle with a certain probability, the proposed cycle service level. After the planning cycle the milk runs and leveling parameters are re-optimized and inventory is re-initialized. To constrain the frequency of re-planning and to increase stability in accordance with the principles of lean production (section 1), we refer to a planning cycle that comprises multiple replenishment cycles.

The approach is based on the following sequence of events.

- I. *Initialization before the planning cycle:*  
 Determination of milk runs and part-specific leveling parameters based on stochastic material requirements
- II. *For each period  $t$  within the planning cycle:*
  - (a) Application of the leveling policy based on deterministically known material requirements
  - (b) Execution of the milk runs based on the leveled replenishment quantities; material is available at the beginning of period  $t$



- (c) Consumption of material in period  $t$
- (d) Update of the inventory balance at the end of period  $t$

In the following, we first define the replenishment policy with leveling property. Second, we integrate this replenishment policy into an SVRP.

### 3.2 Assumptions

In the remainder we will make use of the following assumptions:

- A1. Common frequency routing: We assume that all parts are replenished in the same frequency. This for instance holds true for parts which are supplied on a daily basis but can likewise be adapted to other replenishment frequencies. As a consequence, replenishment orders for each part are placed every time a milk run is operated.
- A2. Stationary distributions of the material requirements: The probability distribution of material requirements is part specific, stationary, and independently distributed. The parameters of the probability distribution are known.
- A3. Fixed partition policy: Suppliers are partitioned into disjoint clusters (e.g., based on their geographical location). As a consequence, the milk runs of every cluster can be planned independently of the other clusters.
- A4. Material availability: We assume that material requirements become deterministically known for a constant number of periods in advance of production based on frozen master production schedules (e.g., 3 days before production). Based on this information, the replenishment orders for a planning period are placed such that material is available at the beginning of that period. Without loss of generality, we assume that lead times are zero in the following. We thus do not distinguish physical inventory and net stock but exclusively refer to physical inventory. Material is assumed to be available at the supplier level.
- A5. Unconstrained number of homogeneous vehicles: We consider an unlimited supply of vehicles. Each vehicle has the same fixed capacity.
- A6. Single sourcing: Every part is supplied by one supplier. One supplier may be in charge of multiple parts.

- A7. Manufacturer focus: Since we assume the perspective of the manufacturer, we do not consider costs from holding inventory at the suppliers.
- A8. Closed routes: All tours start and end at the manufacturer.
- A9. Event management: Even if the standard routes and the initial inventory level are designed to accommodate for swings in material requirements within the planning cycle, there may be situations of impending stock-outs. In this case event management is necessary to supply additional parts (e.g., express delivery). The maximum probability of such an event is exogenously defined by management.

### 3.3 Leveling replenishment policy

The goal of the leveling policy is the reduction of the variability of the replenishment quantities as compared to material requirements. A leveling replenishment policy therefore defines a functional relationship between the random material requirements  $D_{i,t}$  of a part  $i$  in period  $t$  and the associated replenishment quantity  $\Gamma_{i,t}$ . To moderate this relationship, we make use of real-valued variables, the part-specific degree of leveling  $\eta_i$  ( $\eta_i \in [0,1]$ ). From this we derive the following conditions for the replenishment quantity  $\Gamma_{i,t} = \Gamma_{i,t}(D_{i,t}, \eta_i)$ :

$$(1) \quad E[\Gamma_{i,t}(D_{i,t}, \eta_i)] = E[D_{i,t}] = \mu_i$$

$$(2) \quad \text{Var}[\Gamma_{i,t}(D_{i,t}, \eta_i)] = (1 - \eta_i)\text{Var}[D_{i,t}] = (1 - \eta_i)\sigma_i^2$$

Equation (1) ensures that the expected value of the replenishment quantity and the expected value of the material requirements are equal. As a result, there is no effect from leveling on the inventory balance in the long run. Equation (2) describes the leveling of the replenishment quantity. By definition, the variance of the replenishment quantity is  $(1 - \eta_i)$  times the variance of the material requirements.

#### Lemma 1

*The replenishment policy*

$$(3) \quad \Gamma_{i,t}(D_{i,t}, \eta_i) = \sqrt{1 - \eta_i}D_{i,t} + (1 - \sqrt{1 - \eta_i})\mu_i$$

*fulfills equations (1) and (2).*

### Proof

Condition (1): With respect to the expected value  $E[\Gamma_{i,t}(D_{i,t}, \eta_i)]$  we can state:

$$\begin{aligned}
 (4) \quad E[\Gamma_{i,t}(D_{i,t}, \eta_i)] &= E[\sqrt{1-\eta_i}D_{i,t}] + E[(1-\sqrt{1-\eta_i})\mu_i] \\
 &= \sqrt{1-\eta_i}E[D_{i,t}] + (1-\sqrt{1-\eta_i})\mu_i \\
 &= \sqrt{1-\eta_i}\mu_i + (1-\sqrt{1-\eta_i})\mu_i \\
 &= \mu_i
 \end{aligned}$$

Condition (2): From rearranging terms we find for the variance  $\text{Var}[\Gamma_{i,t}(D_{i,t}, \eta_i)]$ :

$$\begin{aligned}
 (5) \quad \text{Var}[\Gamma_{i,t}(D_{i,t}, \eta_i)] &= \text{Var}[\sqrt{1-\eta_i}D_{i,t}] \\
 &= (1-\eta_i)\text{Var}[D_{i,t}] = (1-\eta_i)\sigma_i^2
 \end{aligned}$$

Equation (3) can be interpreted as a convex combination of the actual material requirements and the expected requirements. To scale the influence of both terms, the degree of leveling  $\eta_i$  is used. For this variable, values between 0 and 1 are feasible. If the degree of leveling is set to zero, the policy orders the part's actual material requirements. Accordingly, there is no leveling. If the degree of leveling is set to one, the policy orders the constant expected value of the part's material requirements. This may be interpreted as complete leveling.

The leveling replenishment policy in (3) follows a myopic concept, because the replenishment quantity of a period  $t$  only depends on the demand of that period. While differing from commonly used replenishment policies, this emancipation is necessary to obtain the desired leveling property.

The leveling replenishment policy will be used to extend an SVRP in the following section.

### 3.4 Chance constrained stochastic vehicle routing problem

For the modeling of the SVRP we adopt the chance constrained formulation of [Stewart and Golden \(1983\)](#). Due to the assumption of common frequency routing (A1), there is only one set of milk runs, which is repeated every period. As a result it is sufficient to formulate a static model.

Without loss of generality we assume that every part  $i$  is sourced from one supplier  $i$  (A6). In total  $n$  parts/suppliers are considered. These correspond to the given cluster of suppliers (A3). We have costs  $c_{ij}$

for traveling from supplier  $i$  to supplier  $j$ , which may include direct transportation costs as well as stop-over and dispatching costs. We consider  $m$  vehicles with the same capacity  $Q$  (A5). The number of vehicles is set sufficiently high (e.g., equivalent to the number of suppliers).  $T_m$  is the set of all feasible solutions for the multiple traveling salesman problem ( $m$ -TSP). Decision variables  $x_{ijk}$  equal one if vehicle  $k$  travels from supplier  $i$  to supplier  $j$  and zero otherwise. The probability that material requirements exceed capacity is constrained to the level  $\delta$ . With this notation we can define the SVRP as given in (6) – (8).

$$(6) \quad \min_x Z = \sum_k \sum_{i,j} c_{ij}x_{ijk}$$

s.t.

$$(7) \quad \Pr\left(\sum_{i,j} D_{i,t}x_{ijk} \leq Q\right) \geq 1 - \delta, \forall k, t$$

$$(8) \quad x = [x_{ijk}] \in T_m.$$

The objective function (6) minimizes the transportation costs over all milk runs. Chance constrained formulation (7) limits the probability that the uncertain material requirements of all parts assigned to route  $k$  are larger than the available capacity  $Q$  on that route to the level  $\delta$ . Constraints (8) state that the milk runs  $x$  must be a feasible solution of an  $m$ -TSP. As a result, only those solutions are considered where every supplier is served exactly once. For the sake of brevity we skip a formal description of the  $m$ -TSP. For a comprehensive overview of  $m$ -TSP formulations and solution procedures please refer to [Bektas \(2006\)](#). In section 4 we will present a solution approach which relies on a TSP for which we provide a formal model description.

### 3.5 Model integration

In order to integrate the leveling replenishment policy (3), the SVRP (6) – (8) has to be modified in two terms. This includes (a) the substitution of the material requirements  $D_{i,t}$  by the leveled replenishment quantities  $\Gamma_{i,t}$  and (b) the extension of the objective function with respect to inventory costs.

We apply a cyclical approach. The idea is to minimize the average expected costs per period of a cycle with  $T$  milk run executions. Due to common fre-



quency routing (A1), we have the same transportation costs in every period. The average expected inventory costs of a part  $i$  per period depend on the holding cost rate  $h_i$  and the average expected inventory level ( $1/T \cdot E[I_{i,t}]$ ). More specifically, the inventory level is given as a function of the degree of leveling  $\eta_i$  and a given (cycle) service level ( $1 - \alpha$ ). With these additional parameters and variables the following model results:

$$(9) \quad \min_{x,\eta} Z = \sum_k \sum_{i,j} c_{ij} x_{ijk} + \frac{1}{T} \sum_t \sum_i h_i E[I_{i,t}(\eta_i, 1 - \alpha)]$$

s.t.

$$(10) \quad \Pr \left( \sum_{i,j} \Gamma_{i,t}(D_{i,t}, \eta_i) x_{ijk} \leq Q \right) \geq 1 - \delta, \forall k, t$$

$$(11) \quad x = [x_{ijk}] \in T_m, \eta = \{\eta_i \in [0,1], \forall i\},$$

$$t \leq T, t \in \mathbb{N}$$

The objective of model (9) – (11) is to minimize the average transportation and expected inventory costs per period. The model combines two decisions: the decision on the routes  $x_{ijk}$  and the parameter  $\eta_i$  of the leveling replenishment policy for every part (3). Based on this, there are two options to deal with the variability of material requirements. The first is to include additional capacity in the routes such that varying replenishment quantities can be satisfied to the level ( $1 - \delta$ ). This will, however, result in additional transportation costs, as expressed by the first term of the objective function (9). The alternative is to make use of leveling. Leveling reduces the variability of the material requirements and therefore allows for improved routing and lower transportation costs respectively. Yet, more inventory is required to compensate for swings in material requirements. This is reflected by the second term of the objective function (9).

#### 4 Solution approach for normally distributed material requirements

Model (9) – (11) is valid for any stationary probability distribution of the parts' material require-

ments  $D_{i,t}$ . Due to constraints (10) it is both, non-linear and stochastic. Since it embeds an  $m$ -TSP, it is generally NP-hard.

To solve the model, we need to specify the probability distribution of the parts' material requirements. Based on that, we can compute the expected inventory level and provide an explicit formulation of the capacity constraints (10).

We present a solution approach in the following which is based on the assumption of normally distributed material requirements. This approach allows for simultaneously determining the part-specific degree of leveling and the milk runs. It is based on an approximate deterministic reformulation of the problem (9) – (11) and a simple exact decomposition scheme.

##### 4.1 Deterministic reformulation

For the deterministic reformulation we will specify the expected inventory level and present a reformulation for the stochastic capacity constraints (10).

###### 4.1.1 Approximation of the average expected inventory level

Our procedure to compute the average expected inventory level is based on two approximations: (a) In the first step, the initial inventory is used to approximate the average expected inventory. (b) Secondly, we propose an approximate approach to compute the initial inventory.

(a) For the leveling replenishment policy (3) the unconstrained inventory level  $I_{i,t}^{uc} = I_{i,t}^{uc}(\eta_i, 1 - \alpha)$  at the end of period  $t$  is given by the sum of the inventory level at the end of period  $t-1$  and the difference between the replenishment quantity  $\Gamma_{i,t} = \Gamma_{i,t}(D_{i,t}, \eta_i)$  of part  $i$  in period  $t$  and the part's material requirements  $D_{i,t}$  in that period (12).

$$(12) \quad I_{i,t}^{uc} := I_{i,t-1}^{uc} + (\Gamma_{i,t} - D_{i,t}), t \in \mathbb{N}, I_{i,0} \geq 0$$

Assuming normally distributed material requirements ( $D_{i,t} \sim N(\mu_i, \sigma_i), \forall t$ ), both, the replenishment quantities and the differences between the replenishment quantity and the actual material requirements follow a normal distribution (13). Note, that both figures are not stochastically independent as the replenishment quantities depend on the material requirements. For details see Appendix A2.

$$(13) \quad (\Gamma_{i,t} - D_{i,t}) \sim N(0, (1 - \sqrt{1 - \eta_i}) \sigma_i)$$

As a result, we can interpret the unconstrained inventory level (12) as a time discrete stochastic process with normally distributed increments and without a drift. This corresponds to a time discrete Brownian motion, which varies around the initial inventory level  $I_{i,0} = I_{i,0}(\eta_i, 1 - \alpha)$ , i.e., the inventory level at the beginning of the cycle before replenishment.

If we allowed for negative inventory levels, the expected (physical) inventory level  $I_{i,t}$  at the end of period  $t$  and the initial inventory level derived from the stochastic process (12) would be identically the same. However, since the physical inventory level  $I_{i,t}$  is only defined for non-negative numbers while the inventory level  $I_{i,t}^{uc}$  in (12) is theoretically unconstrained, there is a discrepancy between both figures. The initial inventory level will hence be less than or equal to the expected (physical) inventory per period. With the approximation error due to non-negativity  $\varepsilon_i^{NN}$  we have the average expected inventory level:

$$(14) \quad \frac{1}{T} \sum_{t=1}^T E[I_{i,t}] = I_{i,0} + \varepsilon_i^{NN}, \varepsilon_i^{NN} \geq 0, \forall i$$

(b) To determine the initial inventory level  $I_{i,0}$  such that the given cycle service level  $(1 - \alpha)$  is met, we use a second approximation. This approximation is based on the time continuous representation  $\tilde{I}_{i,t}$  ( $t \in \mathbb{R}^+$ ) (i.e., a time continuous Brownian motion) of the time discrete stochastic process (12).

Let  $M_{i,T}$  be the minimum inventory level of a cycle with  $T$  periods and let  $\tilde{M}_{i,T}$  be the minimum inventory level of the time continuous approximation of that cycle (15).

$$(15) \quad M_{i,T} := \min_{t \leq T, t \in \mathbb{N}_0} I_{i,t}, \tilde{M}_{i,T} := \min_{t \leq T, t \in \mathbb{R}^+} \tilde{I}_{i,t}$$

By the definition of the cycle service level, the probability of a stock-out is constrained to  $\alpha$  as given in (16). For the time continuous approximation  $\tilde{I}_{i,t}$  we know that the probability of a stock-out is greater than or equal to that of the discrete process (17). With the reflection principle it follows for the time continuous Brownian motion  $\tilde{I}_{i,t}$  that the probability of a stock-out in a cycle is equal to two times the probability of a stock-out in the last period (17) (Protter 2005).

$$(16) \quad \Pr(M_{i,T} \leq 0) \leq \alpha$$

$$(17) \quad \Pr(M_{i,T} \leq 0) \leq \Pr(\tilde{M}_{i,T} \leq 0) = 2 \cdot \Pr(\tilde{I}_{i,T} \leq 0)$$

Based on this, we make the following approximation.

$$(18) \quad 2 \cdot \Pr(\tilde{I}_{i,T} \leq 0) \leq \alpha$$

$$(19) \quad \Leftrightarrow \Pr(\tilde{I}_{i,0} + (\tilde{I}_{i,T} - \tilde{I}_{i,0}) \leq 0) \leq \alpha/2$$

Since  $\tilde{I}_{i,t}$  is a Brownian motion, it follows that  $(\tilde{I}_{i,T} - \tilde{I}_{i,0})$  is normally distributed with expected value 0 and a standard deviation which is  $\sqrt{T}$  times the standard deviation of a discrete increment (13). Applying the one-dimensional quantile rule we get (20). The initial inventory level is the product of the  $(1 - \alpha/2)$  quantile of the standard normal distribution and the standard deviation of the cumulative increments.

$$(20) \quad \Leftrightarrow \tilde{I}_{i,0} \geq z_{1-\alpha/2}^{N(0,1)} \cdot (1 - \sqrt{1 - \eta_i}) \cdot \sigma_i \cdot \sqrt{T}$$

Due to the continuous approximation in (17), (20) is an upper boundary on the required initial inventory  $I_{i,0}$ . With the error due to the continuous approximation  $\varepsilon_i^{CP}$  we have (21).

$$(21) \quad I_{i,0} = \tilde{I}_{i,0} - \varepsilon_i^{CP}, \varepsilon_i^{CP} \geq 0, \forall i.$$

The quality of the approximation increases with the number of planning periods. For a large number of planning periods the time discrete process eventually converges to the time continuous one.

The two approximations have opposing effects. While the first approximation tends to underestimate the average expected inventory level, there is an opposing bias resulting from the second approximation.

As a result of (14) and (21) we can write for the average expected inventory:

$$(22) \quad \frac{1}{T} \sum_{t=1}^T E[I_{i,t}] = I_{i,0} + \varepsilon_i^{NN} \\ = \tilde{I}_{i,0} - \varepsilon_i^{CP} + \varepsilon_i^{NN}, \forall i, \varepsilon_i^{NN}, \varepsilon_i^{CP} \geq 0$$

From inserting (20) into (22) we obtain for the approximation of the average expected inventory level:

$$(23) \quad \frac{1}{T} \sum_{t=1}^T E[I_{i,t}] \approx z_{1-\alpha/2}^{N(0,1)} (1 - \sqrt{1 - \eta_i}) \sigma_i \sqrt{T}, \forall i.$$

#### 4.1.2 Reformulation of the capacity constraints

In the next step, we will present the deterministic reformulation of the stochastic constraints (10). This reformulation is based on Lemma 2 below:

##### Lemma 2

If a constant  $z_k$  exists such that  $\forall t \leq T$

$$(24) \Pr \left( \frac{\sum_{i,j} \Gamma_{i,t}(D_{i,t}, \eta_i) x_{ijk} - \sum_{i,j} \mu_i x_{ijk}}{(\sum_{i,j} (1 - \eta_i) \sigma_i^2 x_{ijk}^2)^{1/2}} \leq z_k \right) = 1 - \delta,$$

then the stochastic capacity constraint (10) can be replaced by the following deterministic constraint

$$(25) \sum_{i,j} \mu_i x_{ijk} + z_k \left( \sum_{i,j} (1 - \eta_i) \sigma_i^2 x_{ijk}^2 \right)^{1/2} \leq Q.$$

##### Proof:

Since  $D_{i,t}$  is independent identically normally distributed for every period  $t$  and the replenishment quantities  $\Gamma_{i,t}(D_{i,t}, \eta_i)$  fulfill equation (1) and (2) (Lemma 1), we can state that  $\Gamma_{i,t}(D_{i,t}, \eta_i)$  is normally distributed with  $\Gamma_{i,t} \sim N(\mu_i, \sqrt{1 - \eta_i} \sigma_i)$ . For a given routing decisions  $x_{ijk}$ , the left side in (24) is also normally distributed. For the continuous normal distribution there is a constant  $z_k$  such that equation (24) is fulfilled. Therefore, we can apply the quantile rule for decision variables of Vajda (1972) (transformation from (24) to (25)). Note that  $x_{ijk}^2 = x_{ijk}$  because  $x_{ijk}$  is binary.  $\square$

The quantile rule for decision variables differs from the general one-dimensional quantile rule. In general, the distribution function of the cumulative replenishment quantities on a milk run depends on the routing decisions  $x_{ijk}$ . However, there is a constant  $z$  which is independent of  $x_{ijk}$ , if:

$$(26) \frac{\sum_{i,j} \Gamma_{i,t}(D_{i,t}, \eta_i) x_{ijk} - \sum_{i,j} \mu_i x_{ijk}}{(\sum_{i,j} (1 - \eta_i) \sigma_i^2 x_{ijk}^2)^{1/2}}$$

has the same probability distribution as

$$(27) \frac{\Gamma_{i,t}(D_{i,t}, \eta_i) - \mu_i}{(1 - \eta_i) \sigma_i}$$

(Vajda 1972).

This applies to convolution-invariant probability distribution. Assuming normally distributed material requirements, we can interpret  $z_k$  in inequality (25) as the  $(1 - \delta)$  quantile of the standard normal distribution for every route  $k$ .

The first summand in inequality (25) is independent of the degree of leveling. It describes the sum of the expected values of the replenishment quantities on a particular milk run. The second summand is dependent on the degree of leveling and represents the variation of the replenishment quantity. With increasing degree of leveling, the second summand in inequality (25) decreases so that the required transportation capacity decreases, too. To make sure that the sum of the expected replenishment quantities sourced from the suppliers on every tour is lower than or equal to the vehicle capacity, the second summand has to be positive. Accordingly,  $(1 - \delta)$  is greater than 0.5.

The approximation of the average expected inventory level (23) and the deterministic reformulation of the capacity constraints (25) result in the following model:

$$(28) \min_{x, \eta} Z = \sum_k \sum_{i,j} c_{ij} x_{ijk} + \sum_i h_i z_{1-\delta}^{N(0,1)} (1 - \sqrt{1 - \eta_i}) \sigma_i \sqrt{T}$$

s.t.

$$(29) \sum_{i,j} \mu_i x_{ijk} + z_{1-\delta}^{N(0,1)} \left( \sum_{i,j} (1 - \eta_i) \sigma_i^2 x_{ijk}^2 \right)^{1/2} \leq Q, \forall k$$

$$(30) x = [x_{ijk}] \in T_m, \eta = \{\eta_i \in [0,1], \forall i\}.$$

The objective is to minimize the average costs per period of a cycle of  $T$  periods. Note, that even though the model is static, there is an influence of the length of the planning cycle on the inventory costs in (28).

Like the basic model (9) – (11), model (28) – (30) is NP-hard, but in addition to that, the objective function and the constraints are non-linear. Therefore, this problem is difficult to solve even for small instances. In the following we develop a solution approach which is based on decomposition.

## 4.2 Problem decomposition

The following decomposition scheme generates an optimal solution to the model (28) – (30). It is based on the idea to distinguish one sub-model for the determination of the degree of leveling and one sub-model for routing. The optimal solution is derived from integrating both sub-models.

For a given set of suppliers/parts  $S$  on one milk run we can determine the **degrees of leveling** which result into the minimal inventory costs  $IC$  while satisfying the capacity constraint (25) by solving the following sub-model:

$$(31) \quad IC = \min_{\eta} \sum_{i \in S} h_i z_{1-\alpha/2}^{N(0,1)} (1 - \sqrt{1 - \eta_i}) \sigma_i \sqrt{T}$$

s.t.

$$(32) \quad \sum_{i \in S} (1 - \eta_i) \sigma_i^2 \leq \tilde{Q} = \left( \frac{Q - \sum_{i \in S} \mu_i}{z_{1-\delta}^{N(0,1)}} \right)^2,$$

$$(33) \quad \eta = \{\eta_i \in [0,1], \forall i \in S\}$$

This model is a simplification of model (28) – (30). In the objective function (31) the inventory costs  $IC$  are minimized for a given route serving all suppliers in  $S$ . Accordingly, the degrees of leveling  $\eta_i$  are the only decision variables left. By rearranging terms in (32) we define the new (constant) capacity  $\tilde{Q}$ . It can easily be seen that inequality (32) is linear. The objective function of the sub-problem (31) – (33) is still non-linear. However, it can be solved to optimality with a standard solution algorithm for non-linear optimization problems, because the objective function is convex and the decision variables are continuous.

The routing decision is still NP-hard. However, due to the fixed partition policy (A3), we can restrict the analysis to small (sub) networks of geographically concentrated suppliers (i.e., a small number of milk runs). For these networks we enumerate all set partitions of suppliers, such that every supplier is included in exactly one element of a partition. Each partition may be interpreted as the basis for defining milk runs. Suppliers are aggregated to groups (the elements of a partition) which are to be served on one milk run. However, partitions do not give details on the milk runs and the associated transportation costs. To determine the transportation costs  $TC$  of each  $r$ -th partition  $P_r$ , we determine **optimal routes** for each element  $S \in P_r$  of the partition, starting and ending at the manufacturer's

site "0" (A8). Towards this end, we solve the following TSP. The binary decision variables  $y_{ij}$  indicate whether the vehicle travels from  $i$  to  $j$  ( $y_{ij} = 1$ ). Otherwise  $y_{ij}$  is set to zero.

$$(34) \quad TC = \min_y \sum_{i,j \in S \cup \{0\}} c_{ij} y_{ij}$$

s.t.

$$(35) \quad \sum_{j \in \{S \cup \{0\}\} \setminus \{i\}} y_{ij} = 2, \forall i \in \{S \cup \{0\}\}$$

$$(36) \quad \sum_{i \in \tilde{S}, j \notin \tilde{S}} y_{ij} \geq 2, \forall \tilde{S} \subset \{S \cup \{0\}\}, \tilde{S} \neq \emptyset$$

$$(37) \quad y_{ij} \in \{0,1\}, \forall i, j \in \{S \cup \{0\}\}.$$

Objective function (34) minimizes the transportation costs. Constraints (35) and (36) ensure that every supplier will be served and that there are no sub-cycles.

The **integrated solution approach** is given in Algorithm 1. After determining the set of partitions in the initializing step, we compute the transportation and holding costs for each element of each partition. This is done by solving models (31) – (33) and (34) – (37).

---

### Algorithm 1

---

$k := 1$

Determine set of all partitions of suppliers  $P$

**for**  $r = 1$  to  $|P|$  **do**

**if**  $(\sum_{i \in S} \mu_i \leq Q, \forall S \in P_r)$  **then**

    Solve model (31) – (33),  $\forall S \in P_r$

    Solve model (34) – (37),  $\forall S \in P_r$

$Costs(n) := \sum_{S \in P_r} (TC(S) + IC(S))$

**else if**

$Costs(n) := BigM$

**end for**

$MinPartition := \arg \min_r Costs(r)$

$\eta_i^* := \eta_i(MinPartition)$

**for all**  $S \in P_{MinPartition}$  **do**

$x_{ijk}^* := y_{ij} \forall i, j \in S$

$k = k + 1$

**end for**

---

To improve the performance of the algorithm, we only consider partitions if the cumulated expected material requirements do not exceed the vehicle

capacity for every element of the partition. In every other case, the partition does not allow for feasible milk runs. To exclude these partition from the analysis, we assign the prohibitively high costs  $BigM$ . Despite this effort to reduce complexity, the number of partitions grows exponentially with the number of suppliers.

The optimal partition is the one that minimizes total costs. From this partition we obtain both, the optimal milk runs  $x_{ijk}^*$  and the optimal degree of leveling  $\eta_i^*$ . Note that  $x_{ijk}^*$  and  $\eta_i^*$  are optimal solutions to model (28) – (30), which is based on an approximation of the average holding costs.

## 5 Numerical analysis

### 5.1 Overview and data

The following analysis has three objectives. In the first step, we investigate the performance of the proposed approach as compared to two benchmarks from the literature. In the second step, we take a closer look at how effectively the approach works in consolidating material requirements into milk runs and in the third step, we identify operating conditions for which the approach is in particular well suited.

To incorporate different operating conditions into the analysis, we consider three factors: cost structure, variability of material requirements, and network structure.


Different cost structures are obtained from varying the holding cost rate  $h_i$ , while keeping the transportation costs  $c_{ij}$  constant. The same holding cost rate is used for each supplier ( $h_i \equiv h$ ). In the analysis we consider values of 0, 0.1, 0.2, and 0.3.

To analyze the effect of variability in the material requirements, we include different values for the coefficient of variation  $CV$ . Values of 0.1 and 0.2 are considered. Again identical values are used for each part.

Problems of different network structure are obtained from three idealistic network topologies (Figure 1). To this end, we locate the manufacturer's site in the middle of a square. Network structures are generated from combining different quadrants of this square. Starting from the baseline structure 1 (NS1), where suppliers are evenly spread all over the four quadrants, we consider a structure 2 (NS2), where the suppliers are concentrated into one quadrant. This results in the manufacturer's site being

located at the boundary of the cluster. For structure 3 (NS3) we consider the case that the suppliers are separated into two remote clusters.

**Figure 1: Network structures**



For the analysis we consider 10 suppliers. This number may be representative for the suppliers located in one geographical region. We assume that one part is procured from each supplier.

The parameters and their values are summarized in Table 1.

**Table 1: Parameters used for the analysis**

Parameter	Values
Holding cost rates	$h$ 0 0.1 0.2 0.3
Coefficient of variation	$CV$ 0.1 0.2
Network structure	- NS1 NS2 NS3

From the combination of the values of the three parameters we obtain 24 scenarios (problem classes). To evaluate the performance of the proposed approach we randomly generate 100 problem instances for each scenario and compute the mean of the costs and the number of routes, which result from applying the proposed approach and its benchmarks. For this number of problem instances we found ample convergence in the results.

Each problem instance is defined by distinct supplier locations (x- and y- coordinates) and corresponding transportation cost as well as part-specific stochastic material requirements. The supplier locations are generated from drawing random numbers for the x- and y-coordinates of each supplier. The range of possible coordinates is chosen in correspondence with the network structure. Uniform distributions are used for both coordinates. Given the suppliers' locations, we set the transportation costs in correspondence to the Euclidean distances. In line with other works on milk run planning we assume costs for dispatching and stop-overs to be negligible.



We use a uniform distribution to generate the expected material requirements  $\mu_i$  for every part  $i$ . Regardless of the scenario and problem instance we consider values between 0 and 10. Please note that the range of values has been chosen for the sake of brevity. Since we do not specify units; the range may as well be interpreted as any other range which occurs in industry (e.g., 0...10,000 parts per period). To compute the standard deviation  $\sigma_i$  of each part's material requirements, we use the randomly drawn expected material requirements and the scenario-specific value for  $CV$  according to the following formula.

$$(38) \sigma_i = \mu_i CV, \forall i$$

The vehicle capacity is set to 21, with each part consuming one capacity unit. For the assumptions taken above (10 parts, 5 units expected material requirements per part and period) this results in an expected number of three milk runs and an expected capacity utilization of 80%.

We consider a cycle of 20 periods with milk runs which are operated every period. The approach is implemented in AIMMS. For non-linear optimization we use CONOPT 3.14V, for linear optimization CPLEX 12.3.

## 5.2 Benchmarks

### 5.2.1 Overview of the benchmark policies

To compare the performance of our inventory-transportation system (ITS), we use two benchmarks. The first is defined by an approach which uses information on average material requirements for vehicle routing (i.e., milk run planning) and separately computes the initial inventory (safety stock) which is required to assure a defined cycle service level. Such an approach was for instance proposed in [Gaur and Fisher \(2004\)](#) and [Custódio and Oliveira \(2006\)](#) and has been widely adopted in industry. We will refer to the approach as "VRP+SS". (20) is used to compute the inventory level.

As a second benchmark we consider the approach where the SVRP is solved to determine milk runs (e.g., [Stewart and Golden 1983](#)). Initial inventory is disregarded. The approach is labeled "SVRP" in the following and may be interpreted as the consequent adaption of lean principles to milk run planning.

### 5.2.2 Establishing a common service level definition

In order to compare the ITS with the benchmarks it is necessary to establish a common definition for the service level. Each approach uses different concepts to hedge against swings in material requirements and, as a result, uses different service level definitions. In case of the SVRP a transportation service level is used, which is related to the period-to-period milk run operations. Accordingly, transportation capacity is set such that the probability of the cumulated replenishment quantities exceeding the milk run's transportation capacity is constrained to a defined level. The VRP+SS concept uses a cycle service level which is defined for every part. The cycle service level quantifies the probability that there is no material shortage in the entire planning cycle (i.e., for the defined number of milk runs). In our approach we combine both service levels. To make the approaches comparable, we convert the service levels to the effective cycle service level per part ( $1 - \alpha_{result}$ ). A target level of 95 % is used in the analysis.

To convert the different service level concepts we make the following assumptions. For the SVRP we assume that a stock-out occurs for each part on a milk run, every time transportation demand exceeds transportation capacity. For a single part, only the first stock-out per cycle is considered. Based upon that, the resulting effective cycle service level can be determined analytically (39). An effective cycle service level ( $1 - \alpha_{result}$ ) of 95 % is obtained for a transportation service level ( $1 - \delta$ ) of approximately 99.75 %.

$$(39) (1 - \alpha_{result}) = \sum_{t=1}^T \delta(1 - \delta)^{t-1}$$

The challenge in evaluating the ITS is that we, on the one hand, use approximations to compute the initial inventory and, on the other hand, combine different service level concepts. To assure a fair comparison we make use of simulation. The idea is to numerically identify those parameter values for  $(1 - \alpha)$  and  $(1 - \delta)$  which result into the required effective cycle service level per part of 95 %.

The following procedure is used. For each problem instance introduced in section 5.1 we solve model (28) – (30) with varying values of  $(1 - \alpha)$  and  $(1 - \delta)$ . From this we obtain milk runs and part-specific degrees of leveling for each problem in-

stance. For each problem instance we then use simulation to determine the mean effective cycle service level  $(1 - \alpha_{result})$  which results from applying the milk runs and part-specific degrees of leveling. 1,000 cycles of 20 days' length are considered. For the comparison with the benchmarks we choose the combination of values for  $(1 - \alpha)$  and  $(1 - \delta)$ , where the average effective cycle service level is greater than or equal to 95 % at minimal costs. An approach similar to that applied to the SVRP is used to convert the transportation service level into the material service level. The approach differs in that inventory is used to avoid a stock-out in case of transport capacity being scarce. Using inventory the replenishment quantities of all parts on the milk run are reduced proportionally until the cumulative replenishment quantity matches transportation capacity. Stock-outs for one or multiple parts only occur if there is not sufficient inventory available to allow for the reduction of replenishment quantities. By definition of the leveling concept used, stock-outs also occur, if the inventory level is not sufficiently high to apply the replenishment policy. Only the first stock-out per part and cycle is considered. We vary the values of  $(1 - \alpha)$  and  $(1 - \delta)$  based on the step-size of 0.5 ( $\alpha$ ) and 0.05 ( $\delta$ ) percentage points. For  $(1 - \delta)$  the adjusted transportation service level of 99.75 % is used as a lower boundary. For  $(1 - \alpha)$  we incorporate a range of values which varies around the target value. The same simulation approach is used to establish the effective service level for the VRP+SS. Here, the target value of 95 % is used as an upper boundary when determining that value of  $(1 - \alpha)$ .

### 5.2.3 Dealing with the approximation errors

In reformulating the model we make use of an approximate formulation to compute the expected inventory level. As a result we encounter two types of errors (section 4.1): (a) the error which is due to the continuous approximation  $\varepsilon_i^{CP}$  and (b) the error which is due to non-negativity  $\varepsilon_i^{NN}$ .

(a) From the continuous approximation it follows that the initial inventory level which results from (20) is set too high with respect to the material cycle service level  $(1 - \alpha)$ . However, due to the simulation procedure which is used to establish the

effective cycle service level, it is assured that the required effective cycle service levels of 95 % is met for each part at lowest costs. As a result, we compensate for any bias from  $\varepsilon_i^{CP}$ .

(b) The following approach is used to account for the error  $\varepsilon_i^{NN}$ . Given that we use the reformulated objective function (28) to assess the performance of the ITS and the VRP+SS, we systematically underestimate the average expected physical inventory level and, therefore, the average expected inventory holding costs. The error is a function of the required material cycle service level  $(1 - \alpha)$  and the cycle length. For the analysis, we numerically determine the maximum error and correct the costs of the ITS and the VRP+SS correspondingly.

For the case of the VRP+SS we have a constant value for  $(1 - \alpha)$  and, therefore have a constant error  $\varepsilon_i^{NN}$ . In compensating for this error we report the corrected costs in the following analysis.

The ITS follows a differentiated leveling concept, which is based on instance specific values for  $(1 - \alpha)$ . It is therefore not possible to compute a constant error. In using the maximum error to adjust the performance of the ITS, we generally overestimate the average expected costs in the following. The reported results thus give an upper boundary on the costs of the ITS.

Note, that this upper boundary is valid for both, the approximate model (28) – (30) and the original model (9) – (11). In most cases, the milk runs are determined by transportation costs. The degree of leveling is then used as an additional degree of freedom to assure the required service levels  $(1 - \alpha)$  and  $(1 - \delta)$  at minimal average expected costs. As a consequence, we have the same solutions for the original model and the approximate model. The solutions do deviate, if the underestimated average expected holding costs result in differing milk runs being optimal. The upper boundary (i.e., the total costs of the approximate formulation corrected by the maximum error) is, however, greater than or equal to the true average expected costs, which would result if the solution of the approximate model was inserted into the original model, which, in turn, is by definition greater than or equal to the optimal average expected costs of the original model. The upper boundary is therefore valid for the base model as well.

**Table 2: Average total costs across all instances**

Holding cost rate	CV	Network structure 1			Network structure 2			Network structure 3		
		ITS	VRP+SS	SVRP	ITS	VRP+SS	SVRP	ITS	VRP+SS	SVRP
0	0.1	77.99	77.99	83.39	57.50	57.50	62.85	69.44	69.44	76.42
0	0.2	77.99	77.99	90.82	57.50	57.50	71.69	69.44	69.44	84.98
0.1	0.1	79.19	82.64	83.39	59.13	62.00	62.85	70.95	74.21	76.42
0.1	0.2	81.75	87.30	90.82	62.15	66.50	71.69	73.73	78.98	84.98
0.2	0.1	79.97	87.30	83.39	60.22	66.50	62.85	72.14	78.98	76.42
0.2	0.2	84.46	96.61	90.82	65.22	75.51	71.69	76.72	88.52	84.98
0.3	0.1	80.57	91.96	83.39	62.85	71.00	62.85	76.42	83.75	76.42
0.3	0.2	86.16	105.93	90.82	71.69	84.51	71.69	84.98	98.06	84.98

### 5.3 Results

#### 5.3.1 Analysis of the average total costs

The average total costs of the instances which result from applying the ITS and its benchmarks are reported in Table 2.

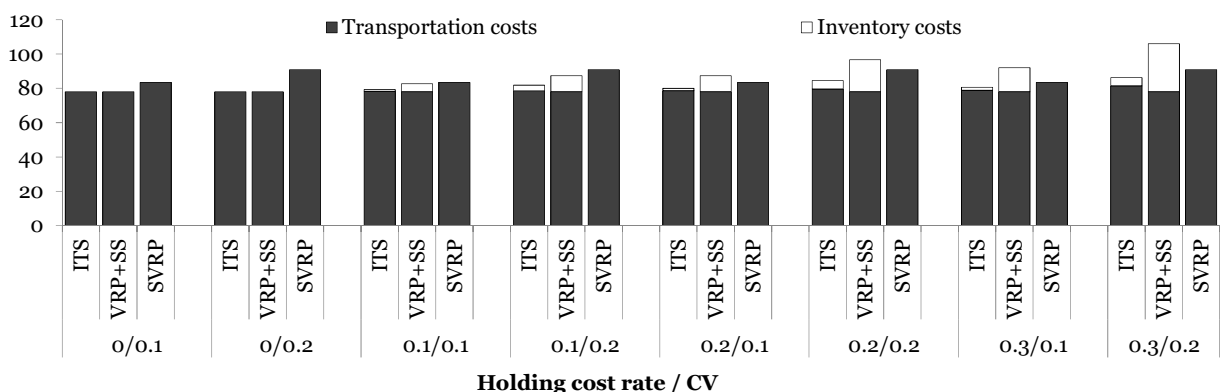
With adjusted service levels the ITS performs at least as well as both benchmarks in terms of costs. Only if holding costs are disregarded, we find the same cost performance as for VRP+SS. This is to be expected, since it will always be optimal to make full use of leveling if holding costs are zero. Accordingly, both approaches result into the same solution.

The impact of different values for  $CV$  and holding cost rates on the structure of the average costs is depicted in Figure 2. If the value of  $CV$  is increased, the costs of the SVRP also increase, since more transportation capacity is required to hedge against swings in material requirements. As a result, transportation costs increase as well. The same effect with respect to  $CV$  applies for the VRP+SS. Here, the transportation costs remain constant, since only the mean material requirements are considered for the routing decision. Instead, the inventory costs

increase with increasing variability of material requirements. In the case of the ITS there is also an increase in transportation and inventory costs depending on  $CV$ . However, this increase is smaller as compared to the VRP+SS and the SVRP. The reason is that the ITS uses transportation capacity and inventory to hedge against the volatility of material requirements.

With increasing holding cost rates, the inventory costs increase. This effect is most prominent for the VRP+SS, which solely relies on inventory, but can likewise be seen for the ITS. The relatively largest inventory cost (26 % of the total costs) result for the VRP+SS for  $CV = 0.2$  and  $h = 0.3$ . As expected, the performance of the SVRP is independent of the holding cost rate. For holding cost rates larger than or equal to 0.2, the SVRP performs better than the VRP+SS. At the same time it can be observed that the ITS solves the trade-off between transportation and inventory costs. While higher transportation costs are encountered as compared to the VRP+SS, these additional costs allow for dampening the increase in inventory costs and therefore for reduced total costs. Overall, the share of inventory costs is

**Figure 2: Average absolute costs for NS1**



rather small; it does not exceed 6 % in any of the scenarios.

For NS2 and NS3 we in principle find the same effects. But because of the higher local concentration of suppliers the transportation costs are lower as compared to NS1. The lowest transportation costs are found for NS2, which has the highest geographical concentration. The inventory costs remain on the same level, and therefore gain in relative importance. Details are given in Table 2.

### 5.3.2 Analysis of routes

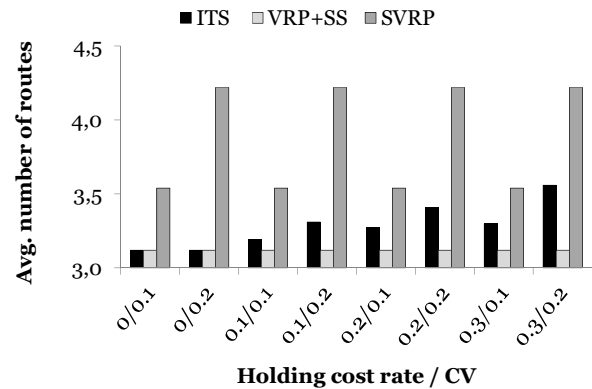
The mean number of routes is depicted in Figure 3, again for the case of different values of  $CV$  and  $h$ . As expected, the mean number of routes of VRP+SS is constant for every scenario. This is due to the fact that expected material requirements are used for routing. These are the same for all parameter combinations.

The SVRP incorporates the variance of material requirements into routing. Since holding inventory is not made use of, the results are independent of holding costs rates. The SVRP therefore generates the same milk runs for all values of  $h$  but depends on the variability as expressed by  $CV$ .

When the ITS is applied, the average number of routes increases with increasing values for the holding cost rates and  $CV$ . This effect is significantly stronger with respect to  $CV$ . The results show that the increase in routes can be effectively reduced as compared to the SVRP. However, given the costs of holding, it is reasonable to build more routes with less suppliers as compared to the VRP+SS.

Again, we find structurally very similar results for the different network structures.

**Figure 3: Average numbers of routes for NS1**



### 5.3.3 Analysis of the impact of different parameter settings

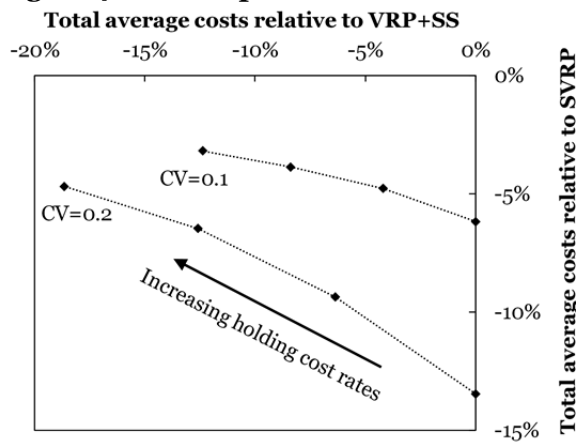
The average relative cost performance of the ITS as compared to the VRP+SS and the SVRP is given in Table 3. For NS1 the results are illustrated in Figure 4. Depicted are isoquants for each value of  $CV$ . The single points of each isoquant differ in holding cost rates.

With increasing values for  $CV$  the relative performance of the ITS increases with respect to the VRP+SS and the SVRP. As compared to the VRP+SS the best average performance is obtained for high holding costs rates, topping at -18.7 % for  $CV = 0.2$  and  $h = 0.3$ . The best average performance with respect to the SVRP is obtained for  $CV = 0.2$  and  $h = 0$  (-13.5 %). For medium holding costs rates of 0.1 and 0.2, the ITS allows for the reduction of transportation and inventory holding costs between 4 % and 12.5 % compared to either benchmark.

**Table 3: Relative performance of the ITS as compared to the benchmarks**

Holding cost rate	CV	Network structure 1		Network structure 2		Network structure 3	
		VRP+SS	SVRP	VRP+SS	SVRP	VRP+SS	SVRP
0	0.1	0.0%	-6.2%	0.0%	-8.5%	0.0%	-8.9%
0	0.2	0.0%	-13.5%	0.0%	-19.3%	0.0%	-17.5%
0.1	0.1	-4.2%	-4.8%	-4.7%	-5.9%	-4.4%	-6.9%
0.1	0.2	-6.4%	-9.4%	-6.6%	-12.7%	-6.7%	-12.5%
0.2	0.1	-8.4%	-3.9%	-9.5%	-4.2%	-8.7%	-5.4%
0.2	0.2	-12.6%	-6.5%	-13.6%	-8.5%	-13.3%	-9.1%
0.3	0.1	-12.4%	-3.2%	-14.2%	-3.1%	-12.8%	-4.3%
0.3	0.2	-18.7%	-4.7%	-20.7%	-6.1%	-19.8%	-6.9%


**Figure 4: Relative performance NS1**



The other network structures show very similar results. However, the level of the relative performance differs between the structures. This is illustrated in Figure 5. Depicted are the average total costs relative to the benchmarks. We averaged results for all scenarios and instances considered.

As can be seen, there is a positive effect from the geographical concentration of the suppliers on the relative performance. With respect to both benchmarks, better results are obtained for NS2 and NS3. Despite this, there are fundamental differences with respect to the underlying reasons.

**Figure 5: Overview of relative performance**



In the case of the SVRP the relative performance for NS2 and NS3 increases by a third as compared to NS1. This is due to the fact that the SVRP tends to build more and smaller routes as compared to the other approaches. As a consequence, relatively poor plans result if synergies from combining material requirements into larger milk runs are high, as is the case for the geographically dense supplier loca-

tions of NS2 and NS3 and the decentralized depot of NS2.

For VRP+SS the relative cost performance of NS2 is slightly better than that of NS3, which is again better than that of NS1. This effect can be traced back to the relative importance of the transportation costs. For NS2 the travelled distances and in turn the transportation costs are lower than for NS3 and NS1. With decreasing transportation costs, there is an increasing relative importance of the inventory costs. Since VRP+SS has the highest inventory costs of all approaches, we see an increase in the relative performance of the ITS as compared to the VRP+SS for the geographically dense topologies NS2 and NS3.

Results on the minimum and maximum performance of the ITS are given in the Appendix. As expected, there is an increasing range of the relative performance figures for increasing values of the average relative performance. The minimal performance figures show that in the worst case the ITS is not able to strictly outperform its benchmarks. Since we adjust the ITS results according to the maximum error (section 5.2) there are even problem instances where the ITS performs worse than the SVRP as indicated by positive relative performance figures. As compared to the SVRP the maximum relative performance of the ITS is up to 31.7 % (NS1,  $CV = 0.2$ ,  $h = 0$ ). This performance is obtained for problem instances for which the leveling of material requirements results in milk runs that provide large reductions in transportations costs. As compared to the VRP+SS the relative performance tops at 26.3 % (NS1,  $CV = 0.2$ ,  $h = 0.3$ ). This performance, in turn, results for problem instances for which there is a small potential to reduce transportation costs by combining suppliers to milk runs. Both effects are the same for NS2 and NS3, yet, on a slightly higher level.

## 6 Conclusions

Given the trend towards centralizing inbound logistics operations at manufacturers, there are emerging opportunities to better coordinate decisions on replenishment and transportation. This in particular becomes relevant if milk runs are used. While making use of economies of proximity to allow for reliable pick-ups at high frequencies, reduced lead times and process improvements along the supply chain, milk runs require for standard routes to be repeated



on a regular basis. This stands in sharp contrast with volatile material requirements, which is typical for many industries. Therefore the question has to be answered of how to coordinate milk run planning and replenishment policy design in the mid-term.

The modeling approach proposed in this paper was developed with the intention of providing decision support for this question. The principle idea is to integrate a simple replenishment policy with a routing component to derive operationally consistent standard routes as a basis for the definition of milk runs. The contribution of the paper is fourfold. (1) We provide a new integrated modeling approach which combines stochastic vehicle routing and a replenishment policy that makes use of inventory to level the variability propagated into transportation operations. (2) To solve the resulting non-linear stochastic model we propose an approximate deterministic reformulation for normally distributed material requirements and develop an exact decomposition approach to solve the reformulated model. (3) We compare our approach with two other approaches which have been presented in the literature (SVRP and VRP+SS). For positive holding cost rates our approach provides the best average performance over all parameter settings considered. Even if holding costs are disregarded, the approach performs at least as well as the benchmarks. (4) From the numerical analysis we also derive operating conditions for the approach. The approach works best for high variability, low to medium holding costs rates and geographically concentrated network structures. From a managerial point of view, these requirements best fit to the characteristics of A/BY and A/BZ parts.

There are five main avenues for future work.

(1) The solution approach works fine for regionally concentrated supply networks, for which the number of suppliers to be considered simultaneously is small. This is in line with current practice in industry, where milk runs typically do not serve more than three suppliers from a single region. The findings from our numerical analysis provide further support for this practice, as milk runs are in particular suited for such types of networks (NS<sub>2</sub>). If, however, larger networks are to be considered, the applicability of the exact solution approach will be limited. For 15 suppliers, the number of possible partitions increases to 1,382,958,545 (from 115,975 for 10 suppliers). While the resulting problem can still be solved to optimality in a few minutes on a

standard PC, it requires efficient heuristic, if larger instances are considered. The development of such heuristics is a promising field for future research. Along this line, there are two main approaches: (a) increasing the performance of the proposed solution approach by reducing the number of partitions and/or incorporating sophisticated TSP heuristics (e.g., [Rego, Gamboa, Glover, and Osterman 2011](#)) and (b) adopting ideas which have been proposed for finding solutions for the VRP ([Laporte 1992](#); [Toth and Vigo 2002](#); [Potvin 2009](#)).

(2) A second field for future work arises from the assumption of common frequency routing. Typically a variety of parts has to be considered within inbound logistics. In the automotive industry these may be up to 10,000 per vehicle model. These parts vary in terms of value, material requirements pattern (i.e., frequency and volume), and physical properties as well as in terms of the geographical region they are supplied from. This gives rise to the question of how to define clusters of parts and/or suppliers, which should be served in a common frequency.

(3) The proposed ITS combines two service level concepts. This results in two parameters, which have to be set by management. The challenge is that there are interdependencies between the parameters. To facilitate a fair comparison with the benchmarks, we use simulation to identify parameter values which result into a defined effective cycle service level. In this, the analytical work could be used to narrow down the range of parameter values. The explicit modeling of the interdependencies between the parameters and its integration into the analytical model is a promising field for future work. Based on this it will be possible to further reduce the required extend of managerial input.

(4) The determination of adequate service levels gains additional importance due to the leveling concept used in this paper and the respective need for event management in case of impending stock-outs. Following the leveling concept, impending stock-outs are accepted, even if transportation capacity is not fully utilized. While being counterintuitive, this is in line with what we see in industry. A central motivation for the introduction of milk runs (and lean production instruments in general) is that process improvement can be reaped along the supply chain. This requires stable plans, suppliers can rely on. In case of deviations to these plans, e.g., in case of impending stock-outs, event management



becomes necessary in collaboration with the affected suppliers. Since replenishment requirements typically become known, once the manufacturer's production plan has been frozen (e.g., 1 week before production), there is sufficient time available for event management to avoid the occurrence of factual out-of-stock situations. The decision as to whether available milk run capacity is used for the transportation of additional replenishment quantities is subject to event management. In order to constrain the probability of event management the proposed approach incorporates service levels. However, future work is necessary to determine the adequate service level, given the implications of event management along the supply chain.

(5) In contrast to most works from the literature, we use a service level concept that is based on planning cycles and therefore several replenishment cycles. This has two main implications: From an economic point of view, the longer uncertainty interval results in the need for additional flexibility (i.e., inventory) and therefore increased costs. However, in industry, implementing plans that are stable for a certain period of time can be seen as a fundamental prerequisite for leveling and exploiting (continuous) process improvements in the context of lean production initiatives. The intention is to reduce costs on a broader scale. From a mathematical point of view, the incorporation of planning cycles results in an increased complexity. If it was feasible to refer to a single replenishment interval, it would be possible to simplify the problem greatly based on decomposition. However, this simplification would require refraining from the principle ideas of implementing stable plans and would result into a large effort for frequent re-planning. A highly relevant question for future work is to investigate the potential of leveled replenishment quantities and stable milk runs on a system level. Case study research seems to be appropriate to tackle this issue.

## Appendices

### A1: List of symbols

#### Decision variables:

$\eta_i$	Degree of leveling of part $i$
$x_{ijk}$	1 if vehicle $k$ travels from supplier $i$ to supplier $j$ , 0 otherwise

$y_{ij}$	1 if a vehicle travels from supplier $i$ to supplier $j$ , 0 otherwise
----------	--

#### Parameters:

$m$	Number of vehicles
$n$	Number of suppliers/parts
$T$	Number of periods in a cycle
$D_{i,t}$	Material requirements of part $i$ in period $t$
$\varepsilon_i^{CP}$	Error due to continuous approximation
$\varepsilon_i^{NN}$	Error due to non-negativity
$\mu_i$	Expected value of material requirements of part $i$
$\sigma_i$	Standard deviation of material requirements of part $i$
$\Gamma_{i,t}$	Replenishment quantity of part $i$ in period $t$ ; available at the beginning of $t$
$CV$	Coefficient of variation
$c_{ij}$	Transportation costs to travel from supplier $i$ to supplier $j$
$h_i$	Holding costs rate of part $i$
$I_{i,t}$	Inventory level of part $i$ at the end of period $t$
$I_{i,t}^{uc}$	Unconstrained inventory level of part $i$ at the end of period $t$
$\tilde{I}_{i,t}$	Inventory level of the continuous approximation of part $i$ at the end of period $t$
$M_{i,T}$	Minimum inventory level of part $i$ until the end of period $T$
$\tilde{M}_{i,T}$	Minimum inventory level of the continuous approximation of part $i$ until the end of period $T$
$T_m$	Set of all solutions of an $m$ -TSP
$Q$	Vehicle capacity
$\tilde{Q}$	Modified vehicle capacity
$(1 - \alpha)$	Inventory cycle service level
$(1 - \alpha_{result})$	Effective cycle service level
$(1 - \delta)$	Transportation service level
$z_{\zeta}^{N(0,1)}$	$\zeta$ quantile of standard normal distribution
$P$	Set of all partitions of suppliers
$P_r$	$r$ -th partition of suppliers
$S$	Subset of suppliers, $S \in P_r$
$s$	Cardinality of Subset $S$
$TC$	Transportation costs
$IC$	Inventory costs

### A2: Variance of increments

$$\begin{aligned}
 (40) \text{Var}(D_{i,t} - \Gamma_{i,t}) &= \text{Var}(D_{i,t}) + \text{Var}(\Gamma_{i,t}) - 2 \cdot \text{Cov}(D_{i,t}, \Gamma_{i,t}) \\
 &= \sigma_i^2 + (1 - \eta_i) \cdot \sigma_i^2 - 2 \cdot \text{Cov}(D_{i,t}, \Gamma_{i,t}) \\
 &= \sigma_i^2 + (1 - \eta_i) \cdot \sigma_i^2 - 2 \cdot \sqrt{1 - \eta_i} \cdot \sigma_i^2 \\
 &= (1 + (1 - \eta_i) - 2 \cdot \sqrt{1 - \eta_i}) \cdot \sigma_i^2 \\
 &= (1 - \sqrt{1 - \eta_i})^2 \cdot \sigma_i^2
 \end{aligned}$$

with

$$\begin{aligned}
 (41) \text{Cov}(D_{i,t}, \Gamma_{i,t}) &= \text{Cov}(D_{i,t}, \sqrt{1 - \eta_i} D_{i,t} + (1 - \sqrt{1 - \eta_i}) \mu_i) \\
 &= \text{Cov}(D_{i,t}, \sqrt{1 - \eta_i} D_{i,t}) \\
 &= \sqrt{1 - \eta_i} \cdot \text{Cov}(D_{i,t}, D_{i,t}) \\
 &= \sqrt{1 - \eta_i} \cdot \sigma_i^2
 \end{aligned}$$

### A3: Additional Results

**Table 4: Relative performance of the ITS as compared to the benchmarks (NS1)**

Holding cost rate	CV	VRP+SS			SVRP		
		Min.	Avg.	Max.	Min.	Avg.	Max.
0	0.1	0.0%	0.0%	0.0%	0.0%	-6.2%	-21.3%
0	0.2	0.0%	0.0%	0.0%	0.0%	-13.5%	-31.7%
0.1	0.1	-0.7%	-4.2%	-6.5%	0.2%	-4.8%	-16.1%
0.1	0.2	-0.7%	-6.4%	-10.1%	0.0%	-9.4%	-29.3%
0.2	0.1	-1.7%	-8.4%	-12.3%	0.4%	-3.9%	-14.2%
0.2	0.2	-1.2%	-12.6%	-19.1%	1.2%	-6.5%	-26.9%
0.3	0.1	-2.4%	-12.4%	-17.8%	0.3%	-3.2%	-12.8%
0.3	0.2	-10.1%	-18.7%	-26.3%	0.2%	-4.7%	-24.4%

**Table 5: Relative performance of the ITS as compared to the benchmarks (NS2)**

Holding cost rate	CV	VRP+SS			SVRP		
		Min.	Avg.	Max.	Min.	Avg.	Max.
0	0.1	0.0%	0.0%	0.0%	0.0%	-8.5%	-23.3%
0	0.2	0.0%	0.0%	0.0%	-4.6%	-19.3%	-32.5%
0.1	0.1	-0.8%	-4.7%	-8.5%	0.0%	-5.9%	-19.4%
0.1	0.2	-0.7%	-6.6%	-14.2%	0.2%	-12.7%	-25.4%
0.2	0.1	-1.4%	-9.5%	-15.7%	0.5%	-4.2%	-16.6%
0.2	0.2	-2.1%	-13.6%	-24.6%	1.8%	-8.5%	-23.1%
0.3	0.1	-3.5%	-14.2%	-21.9%	0.5%	-3.1%	-15.4%
0.3	0.2	-8.5%	-20.7%	-32.5%	0.2%	-6.1%	-20.8%

**Table 6: Relative performance of the ITS as compared to the benchmarks (NS3)**

Holding cost rate	CV	VRP+SS			SVRP		
		Min.	Avg.	Max.	Min.	Avg.	Max.
0	0.1	0.0%	0.0%	0.0%	0.0%	-8.9%	-25.7%
0	0.2	0.0%	0.0%	0.0%	-0.2%	-17.5%	-32.3%
0.1	0.1	-1.1%	-4.4%	-7.3%	0.2%	-6.9%	-22.8%
0.1	0.2	-1.4%	-6.7%	-10.5%	0.0%	-12.5%	-26.5%
0.2	0.1	-2.1%	-8.7%	-14.1%	0.3%	-5.4%	-19.8%



**Table 6 continued: Relative performance of the ITS as compared to the benchmarks (NS3)**

Holding cost rate	CV	VRP+SS			SVRP		
		Min.	Avg.	Max.	Min.	Avg.	Max.
0.2	0.2	-2.6%	-13.3%	-21.1%	0.1%	-9.1%	-24.7%
0.3	0.1	-3.0%	-12.8%	-19.8%	0.0%	-4.3%	-16.9%
0.3	0.2	-6.6%	-19.8%	-31.1%	0.2%	-6.9%	-23.3%

### Acknowledgements

The authors wish to express their thanks to the department editor and the two anonymous referees for their valuable comments and suggestions that helped to improve this paper.

### References

Alegre, Jesús, Manuel Laguna, and Joaquín Pacheco (2007): Optimizing the Periodic Pick-Up of Raw Materials for a Manufacturer of Auto Parts, *European Journal of Operational Research*, 179 (3): 736-746.

Andersson, Henrik, Arild Hoff, Marielle Christiansen, Geir Hasle, and Arne Løkketangen (2010): Industrial Aspects and Literature Survey: Combined Inventory Management and Routing, *Computers and Operations Research*, 37 (9): 1515-1536.

Bard, Jonathan F., Liu Huang, Patrick Jaillet, and Moshe Dror (1998): A Decomposition Approach to the Inventory Routing Problem with Satellite Facilities, *Transportation Science*, 32 (2): 189-203.

Baudin, Michel (2004): *Lean Logistics: The Nuts and Bolts of Delivering Materials and Goods*, Productivity Press: New York, NY.

Bektas, Tolga (2006): The Multiple Traveling Salesman Problem: An Overview of Formulations and Solution Procedures, *Omega*, 34 (3): 209-219.

Bell, Walter J., Louis M. Dalberto, Marshall L. Fisher, Arnold J. Greenfield, R. Jaikumar, Pradeep Kedia, Robert G. Mack, and Paul J. Prutzman (1983): Improving the Distribution of Industrial Gases with an On-Line Computerized Routing and Scheduling Optimizer, *Interfaces*, 13 (6): 4-23.

Beltrami, Edward J. and Lawrence D. Bodin (1974): Networks and Vehicle Routing for Municipal Waste Collection, *Networks*, 4 (1): 65-94.

Bertazzi, Luca, Martin Savelsbergh, and Maria G. Speranza (2008): Inventory Routing, in: Bruce Golden, S. Raghavan, and Edward Wasil (eds.): *The Vehicle Routing Problem: Latest Advances and New Challenges*, Springer: Boston, MA, 49-72.

Campbell, Ann M. and Martin W. P. Savelsbergh (2004): A Decomposition Approach for the Inventory-Routing Problem, *Transportation Science*, 38 (4): 488-502.

Cannella, Salvatore and Elena Ciancimino (2010): On the Bull-whip Avoidance Phase: Supply Chain Collaboration and Order

Smoothing, *International Journal of Production Research*, 48 (22): 6739-6776.

Chen, Youhua F. and Stephen M. Disney (2007): The Myopic Order-Up-To Policy with a Proportional Feedback Controller, *International Journal of Production Research*, 45 (2): 351-368.

Christofides, Nicos and John E. Beasley (1984): The Period Routing Problem, *Networks*, 14 (2): 237-256.

Chuah, Keng H. and Jon C. Yingling (2005): Routing for a Just-in-Time Supply Pickup and Delivery System, *Transportation Science*, 39 (3): 328-339.

Custódio, Ana L. and Rui C. Oliveira (2006): Redesigning Distribution Operations: A Case Study on Integrating Inventory Management and Vehicle Routes Design, *International Journal of Logistics: Research and Applications*, 9 (2): 169-187.

Disney, Michael S., Andrew T. Potter, and Bernard M. Gardner (2003): The Impact of Vendor Managed Inventory on Transport Operations, *Transportation Research Part E: Logistics and Transportation Review*, 39 (5): 363-380.

Dror, Moshe, Michael Ball, and Bruce Golden (1985): A Computational Comparison of Algorithms for the Inventory Routing Problem, *Annals of Operations Research*, 4 (1): 1-23.

Francis, Peter M., Karen R. Smilowitz, and Michal Tzur (2008): The Period Vehicle Routing Problem and its Extensions, in: Bruce Golden, S. Raghavan, and Edward Wasil (eds.): *The Vehicle Routing Problem: Latest Advances and New Challenges*, Springer: Boston, MA, 73-102.

Gaur, Vishal and Marshall L. Fisher (2004): A Periodic Inventory Routing Problem at a Supermarket Chain, *Operations Research*, 52 (6): 813-822.

Golden, Bruce, Arjang Assad, and Roy Dahl (1984): Analysis of a Large Scale Vehicle Routing Problem with an Inventory Component, *Large Scale Systems*, 7 (2-3): 181-190.

Hemmelmayr, Vera, Karl F Doerner, Richard F. Hartl, and Martin W. P. Savelsbergh (2010): Vendor Managed Inventory for Environments with Stochastic Product Usage, *European Journal of Operational Research*, 202 (3): 686-695.

Khouja, Moutaz and Suresh Goyal (2008): A Review of the Joint Replenishment Problem Literature: 1989-2005, *European Journal of Operational Research*, 186 (1): 1-16.

Kiesmüller, Gudrun P. (2010): Multi-Item Inventory Control with Full Truckloads: A Comparison of Aggregate and Individual Order Triggering, *European Journal of Operational Research*, 200 (1): 54-62.

Laporte, Gilbert (1992): The Vehicle Routing Problem: An Overview of Exact and Approximate Algorithms, *European Journal of Operational Research*, 59 (3): 345-358.



Larson, Richard C. (1988): Transporting Sludge to the 106-Mile Site: An Inventory/Routing Model for Fleet Sizing and Logistics System Design, *Transportation Science*, 22 (3): 186-198.

Miltenburg, G. John (1985): Allocating a Replenishment Order among a Family of Items, *IIE Transactions*, 17 (3): 261-267.

Moin, Noor H. and Said Salhi (2007): Inventory Routing Problems: A Logistical Overview, *Journal of the Operational Research Society*, 58 (9): 1185-1194.

Ormeçi, Melda, J. G. Dai, and John V. Vate (2008): Impulse Control of Brownian Motion: The Constrained Average Cost Case, *Operations Research*, 56 (3): 618-629.

Potvin, Jean-Yves (2009): A Review of Bio-Inspired Algorithms for Vehicle Routing, in: Francisco B. Pereira and Jorge Tavares (eds.): *Bio-Inspired Algorithms for the Vehicle Routing Problem*, Springer: Berlin et al. 1-34.

Protter, Philip (2005): *Stochastic Integration and Differential Equations*, Springer: Berlin et al.

Qu, Wendy W., James H. Bookbinder, and Paul Iyogun (1999): An Integrated Inventory-Transportation System with Modified Periodic Policy for Multiple Products, *European Journal of Operational Research*, 115 (2): 254-269.

Rego, César, Dorabela Gamboa, Fred Glover, and Colin Osterman (2011): Traveling Salesman Problem Heuristics: Leading Methods, Implementations and Latest Advances, *European Journal of Operational Research*, 211 (3): 427-441.

Russell, Robert A. and Wayne Igo (1979): An Assignment Routing Problem, *Networks*, 9 (1): 1-17.

Stacey, Jennifer, Malini Natarajarathinam, and Charles Sox (2007): The Storage Constrained, Inbound Inventory Routing Problem, *International Journal of Physical Distribution & Logistics Management*, 37 (6): 484-500.

Stewart, William R. and Bruce L. Golden (1983): Stochastic Vehicle Routing: A Comprehensive Approach, *European Journal of Operational Research*, 14 (4): 371-385.

Toth, Paolo and Daniele Vigo (2002): Models, Relaxations and Exact Approaches for the Capacitated Vehicle Routing Problem, *Discrete Applied Mathematics*, 123 (1-3): 487-512.

Trudeau, Pierre and Moshe Dror (1992): Stochastic Inventory Routing: Route Design with Stockouts and Route Failures, *Transportation Science*, 26 (3): 171-184.

Vajda, Steven (1972): *Probabilistic Programming*, Academic Press: New York, NY.

Yu, Yuang, Chengbin Chu, Haoxun Chen, and Feng Chu (2012): Large Scale Stochastic Inventory Routing Problems with Split Delivery and Service Level Constraints, *Annals of Operations Research*, 197 (1): 135-158.

## Biographies

**Thomas Volling** studied Industrial Engineering and Management Science at the Technische Universität Braunschweig and the University of Rhode Island. He obtained his PhD in 2008 and holds the position of an Assistant Professor at the Institute of Automotive Management and Industrial Production. His main area of research is make-to-order order fulfillment systems. Further interests include the development of quantitative and qualitative operational templates.

**Martin Grunewald** studied Business Mathematics at the University of Magdeburg. He currently works as a research assistant at the Institute of Automotive Management and Industrial Production. His research interests are in- and outbound logistics in the automotive industry, with a special interest in leveled logistics concepts. He worked in different logistics projects with an automotive manufacturer.

**Thomas S. Spengler** is Professor of Production and Logistics Management and Director of the Institute of Automotive Management and Industrial Production at the Technische Universität Braunschweig. As well as logistics his research interests include the development and implementation of techno-economic models for decision support. He has researched and consulted with a wide range of industries including iron and steel, automotive, and consumer goods. He has published widely in the field of Production and Logistics Management.