Logistics Research (2017) 10:2
DOI 10.23773/2017_2

Original Article

Allocation of Bulk Tanks to Improve Industrial Gas Distribution to Customers with Time-Varying Demand

Leily Farrokhvar[1] | Kimberly P. Ellis[2] | Jeffrey E. Arbogast[3]

Received: 22 March 2016 / Accepted: 21 March 2017 / Published online: 17 January 2018

ABSTRACT

In this paper, we describe the results from an academic and industry collaboration to address the bulk tank allocation problem for industrial gas distribution systems where customer demand varies over time. The bulk tank allocation problem determines the preferred size of bulk tanks to assign to customer sites to minimize recurring gas distribution costs and initial tank installation costs. The problem is modeled as a mixed-integer programming model, and three solution approaches are presented. In the first two approaches, the problem is decomposed and a restricted master problem is solved. The third approach is a two phase periodically restricting heuristic approach. The results demonstrate the opportunities for substantial improvements in resource allocation and reductions in operational costs.

KEYWORDS strategic tank allocation · inventory routing · large-scale optimization · vehicle routing with split deliveries · logistics planning

Leily Farrokhvar
leily.farrokhvar@mail.wvu.edu

Kimberly P. Ellis
kpellis@vt.edu

Jeffrey Arbogast
Jeffrey.Arbogast@airliquide.com

[1] Assistant Professor, Center for Excellence in Logistics and Distribution, West Virginia University, Morgantown, West Virginia 26506, USA

[2] Associate Professor, Center for Excellence in Logistics and Distribution, Virginia Tech, Blacksburg, Virginia 24061, USA

[3] AL International Expert, Computational & Data Science R&D, Air Liquide, Newark, Delaware, USA

1 INTRODUCTION AND BACKGROUND

The global market for the industrial gas sector was approximately $77 billion in 2014 and is projected to increase to $116.6 billion in 2020 (Industry Experts, 2014). Industrial gas companies typically have significant investments in their distribution systems and strive to improve operational efficiency throughout their supply chains.

Industrial gases can be distributed via pipelines, delivered to customers as packaged gases, or transported in bulk via truck, rail, or ship. The focus of this research is bulk distribution, which accounts for approximately 34% of revenue in this industry (Baker and Garvey, 2004). For gases that are transported by cryogenic trailers to bulk tanks at customer sites, industrial gas producers often implement a vendor managed inventory (VMI) system with their customers. The producers are responsible for inventory management as well as gas delivery to customer sites. Customers assume ownership of the gas inventory once it is delivered to the tanks located at their sites. The industrial gas producer generally owns and monitors the tanks and replenishes the tanks using cryogenic trailers. These trailers and storage tanks are generally high value assets and their effective utilization is important for the company. Also, industrial gas producers must meet strict customer service levels while coping with demand fluctuations. Thus, efficient distribution is a main driver to achieve lower costs and remain competitive (Chowhan, 2013).

When distributing gases, tractor-trailers (referred to by trailers in the remainder of this work) typically depart their depots, obtain liquefied products at gas sources, and deliver the gas product to bulk tanks at one or more customer sites. The trailers may either return to the depot and conclude the route or may obtain additional product by refilling at the same or another source and continuing delivery, resulting in a *continuous delivery route*. With a continuous delivery route, a trailer has the opportunity to serve additional customers before concluding a route. Common route construction approach, however, primarily consider depots and customer nodes. Thus, existing route construction methods are extended to accommodate depots, sources, customer nodes and continuous deliver.

The customer tank size is a determining factor for the quantity and volume of deliveries that a customer requires. Generally, customers with small tanks and high demand require more frequent replenishments of smaller volumes. Allocating a larger tank to such customers may improve distribution efficiency by requiring fewer replenishment deliveries in larger quantities. Assigning large tanks to all customers, however, is not feasible when considering limited tank resources, associated costs, and budgets restrictions. When determining the preferred tank for a customer site, industrial gas producers consider estimated consumption patterns, safety stock requirements of each customer, and proximity to other customers. Swapping tanks between customers involves removing a tank from one customer site, moving it to a tank warehouse, refurbishing it, transporting it to another customer location, and installing it at the new location. This is an expensive process which is justified only if the benefits of a more efficient distribution plan exceed the costs of tank exchanges. Therefore, considering the significant investment in bulk tanks, industrial gas companies face an important strategic level decision of bulk tank allocation which directly affects the operational efficiency of their distribution networks.

In practice, bulk tanks are typically allocated to customer sites for multiple years. The bulk tank allocation (BTA) decision is further complicated by customer demand that frequently varies by time period (e.g. quarterly). Thus, it is important (and challenging) to allocate tanks that effectively accommodate time varying demand. Installing proper size tanks at customer locations not only provides improved utilization of their assets (bulk tanks and vehicle fleet) but also better service quality to their customers. This research is motivated by Air Liquide, a leading international gas company, that produces and distributes industrial gases such as hydrogen, oxygen, argon, helium, carbon dioxide, and nitrogen for several business segments in over eighty countries.

The main contributions of this paper include: 1) the formulation of a mixed-integer programming model of the bulk tank allocation problem with time varying demand, 2) development of a solution approach to address the bulk tank allocation problem with time varying demand that also generates a gas distribution plan; 3) incorporation of multiple sources and continuous routes where tractor trailers visit the same or a different source for refilling during gas delivery to customers, which is used in practice, into our methodology; 4) validation of the approach and analysis of solutions using industry representative data sets; and 5) managerial insights on potential total cost savings and distribution savings as well as internal rate of return on tank investments.

In the following section, we briefly review the related literature and formally describe a mixed-integer programming (MIP) model for the bulk tank allocation with time varying (BTATVD) problem. Then, we evaluate the model formulation and explore modifications to improve the solvability of the model. In section 4, we describe a route generation procedure to generate a set of high quality delivery routes quickly and discuss the modifications that we made to the previously developed routing subproblem. The generated routes are then used in three solution approaches proposed in section 5. The three solution approaches include solving a BTATVD restricted master problem (RMP), BTATVD with recurring demand assumption, and a two phase periodically restricting heuristic (PRH). The developed methodologies are demonstrated using modified data sets provided by Air Liquide in section 6, followed by a scalability analysis in section 7. We close with conclusions and our future research directions in section 8.

2 LITERATURE SURVEY

The bulk tank allocation (BTA) problem involves the allocation or reallocation of bulk tanks to customer sites in order to minimize the sum of the net present value of investment costs in the logistics network and distribution costs over a time horizon. The BTA problem involves resource allocation with an underlying split delivery vehicle routing problem (SDVRP). SDVRP is a generalized vehicle routing problem (VRP) in which a customer’s demand can be satisfied from multiple vehicle visits rather than just one. The SDVRP can also be viewed as a special case of the inventory routing problem (IRP) with known demand where no back orders are allowed.

In a classic VRP, a set of customers with known demand is served by a fleet of vehicles. The objective is to minimize the total travel cost while ensuring that each customer’s demand is satisfied by exactly one vehicle visit. In a SDVRP, the restriction of one visit for each customer is relaxed, and a customer’s demand can be fulfilled by several deliveries on multiple vehicles routes. SDVRP provides the set of delivery routes and determines the amount to deliver to each customer on each vehicle route. This problem was first proposed in the literature by Dror and Trudeau (1989), where they show the potential cost savings through split deliveries in a VRP.

Archetti and Speranza (2008) and Archetti and

Speranza (2012) provide surveys of the current research on the SDVRP problem. They provide a description of the mathematical formulation along with analysis of the properties and complexity of the problem. They review exact solution approaches for solving a SDVRP for small instances of the problem including a dynamic programming model (Lee et al., 2006), a two-stage solution approach including iteratively solving assignment problems and traveling salesman problems (Jin et al., 2007), and an approach based on a set covering formulation of the problem and a column generation approach (Feillet et al., 2004). Archetti et al. (2011) propose a branch-and-price-and-cut scheme for solving SDVRP. They use column generation approach to generate routes and determine delivery amounts in the subproblem using a labeling algorithm and then master problem selects the optimal routes.

Archetti and Speranza (2008) also review some effective heuristics that have been designed and tested on larger data sets. They compare the performance of major heuristics available in the literature on several instances ranging from 50 to 199 customers. More recently, several metaheuristic approaches are proposed by Aleman and Hill (2010), Aleman et al. (2010), and Derigs et al. (2010) to solve the SDVRP. Most of the SDVRP literature reports solving instances with less than 200 customers. The additions of tank allocation decisions along with time varying customer demand introduce additional complexity to the problem. Therefore, an efficient heuristic approach is necessary for solving the problem.

In general, the IRP involves the integration and coordination of two components of the logistics systems: inventory management and vehicle routing. It is concerned with the distribution of products from facilities to a set of customers over a given planning horizon. Customers consume the product and can maintain an inventory of the product up to a certain level. A fleet of vehicles of known capacities are available for the distribution of the product. The objective is to minimize the distribution and inventory costs during the planning period. If there are no back orders allowed in the IRP, the problem will be similar to the SDVRP. The problem studied in this research shares similarities with the IRP. Thus, literature on IRP is reviewed to provide insights.

In a survey paper, Moin and Salhi (2006) present a logistical overview of the IRP and classify the papers into single-period, multi-period, infinite time horizon, and problems with stochastic demand. Single-period models are popular in the literature, as they are simpler to solve and often provide a basis for solving multi-period models. The multi-period models are often used for long term planning as they better capture the trade-offs between short-term and long-term decisions. While multi-period models provide a more realistic representation of the problem, they are more complex and therefore most of the papers consider deterministic demand for the customers and solve the problem using heuristic methods. Some research has been published on stochastic IRP (see Bard et al. (1998); Kleywegt et al. (2002, 2004)). Campbell et al. (1998) discuss the complexity and practical issues of the IRP with stochastic demand and present different solution approaches.

In a more recent paper, Bertazzi et al. (2008) review the inventory routing literature and discuss the characteristics and complexity of the problem and describe the challenges encountered when trying to simultaneously minimize the inventory holding and routing costs. Andersson et al. (2010) also provide a comprehensive survey of over 90 scientific papers for IRP where one party in the supply chain is responsible for transportation and inventory management. The papers are categorized by time horizon into the following three categories: instant time horizon, finite time horizon, and infinite time horizon. The papers are compared with respect to problem description, assumptions, and solution approaches proposed in major papers within each category. In most of the cases, the complexity of the problem makes it difficult to solve to optimality and heuristic or decomposition approaches are often employed. For most industry-sized problems, a computational time limit is enforced for exact approaches and then the best integer solution found is reported or a heuristic has been implemented. They emphasize the importance of using advanced decision support tool for inventory management and routing decisions for businesses with complex systems, but found no commercial decision support tools available for the combined inventory management and routing problem. Thus, optimization-based decision support systems are needed for the IRP. Aksoy and Derbez (2003) survey the available software systems for supply chain management. They identify 160 software companies providing supply chain management software. Among these they find several software systems that have separate modules for inventory management and routing, but no software systems where it is combined.

While there is abundant research on the SDVRP and relevant cases of IRP, minimal research was found that integrates this problem with tank allocation for the industrial gas industry. Ellis et al. (2014), You et al. (2011a), and You et al. (2011b) are among the few researchers addressing this problem.

Ellis et al. (2014) develop a mixed-integer programming (MIP) model which minimizes the net present value of gas distribution over a horizon, tank installation, tank reallocation, tank refurbishment, and tank purchasing costs. The model captures complex trade-offs between strategic level investment decisions and operational level routing decisions. They also develop a decomposition-based heuristic solution approach for this problem where a restricted master problem is solved using a promising subset of routes that are generated using a sweep-based algorithm. The estimates of distribution costs using the sweep-based heuristic compared closely to the actual distribution costs in practice. This approach was evaluated using several data sets ranging from a small data set involving a single depot with 5-10 customers to a larger national data set with 18 depots and 1287 customers provided by Air Liquide. The result of the solution approach is an allocation of tanks to customer sites and an estimate of the investment and distribution costs.

You et al. (2011a) address an inventory-distribution planning problem in industrial gas distribution, including bulk tank allocation. The bulk tank allocation problem is modeled as a large scale MIP model to minimize investment costs, change-out costs, and distribution costs. They propose heuristic solution approaches, with a primary focus on a continuous approximation approach. With the continuous approximation approach, detailed vehicle routing parameters and variables are approximated by functions to represent the distribution of customer demands and locations, thus simplifying some aspects of the problem. When included in the model, however, the functions result in a mixed integer nonlinear problem (MINLP). The problem is then linearized and applied to select tank sizes. A comparison of the continuous approximation distribution costs with detailed routing distribution costs is not provided.

You et al. (2011b) continue exploration of inventory-distribution problems in industrial gas distribution. In this paper, the bulk tank allocation problem considers uncertain customer demand and the loss or addition of new customers with an objective to minimize investment costs, tank change-out costs, and distribution costs. For this problem, customer demand fluctuations are assumed to follow a normal distribution. Customer demand can vary by year, but seasonal demand within a year is not considered. They present a stochastic MINLP model, which relies on a continuous approximation approach. The paper primarily focuses on the computational aspects of the stochastic MINLP model, with problem instances varying from 4 to 200 customers. A comparison of the continuous approximation distribution costs with the detailed routing distribution costs is not provided.

The approaches presented by Ellis et al. (2014) and You et al. (2011a) assume constant demand for customers. In reality, customer demand frequently varies by time period, often with seasonal variations. If the tank allocation decisions were made each time period then the problem would decompose by time period. In practice, however, bulk tanks that are allocated to customer sites typically remain in place for multiple years. Throughout this paper, the study time horizon is the time period in which customers' allocated tanks are assumed to remain fixed for the purpose of distribution cost savings. Given that customer demand tends to vary, the challenge is to allocate tanks effectively to accommodate time varying demand during the study time horizon. Considering the significant costs involved in exchanging tanks, a tank exchange is only justifiable if the net present value of the resulting distribution cost savings over the study time horizon outweighs the tank allocation costs. Therefore, we extend the previous model of Ellis et al. (2014) to consider changing demands by period for each customer, resulting in the bulk tank allocation model for time varying demand (BTATVD). Introducing this new dimension to the problem increases complexity, and therefore, we investigate three approaches for addressing the problem including a periodically restricting heuristic (PRH) solution approach. The three approaches are compared and their effectiveness is evaluated.

3 PROBLEM DESCRIPTION

We consider an industrial gas distribution network for a single gas product distributed by trailers from multiple gas sources. The industrial gas company assigns bulk tanks to the customers and deliver the products to these tanks. The customers assume ownership of the delivered product but the tank itself is owned by the industrial gas company. Tank allocation is an expensive process that involves the costs of purchasing new tanks or refurbishing existing tanks, transporting the tanks, and installing them at customer location. The tanks, which are assumed to be allocated at the beginning of the time horizon, directly affect distribution planning. Each customer must be assigned a tank that is large enough to hold each delivery plus the safety stock requirement of that customer. Smaller tanks generally require more frequent but smaller deliveries. In different time periods, a customer may be replenished on different routes. For example, in a high demand period, a customer may be visited on a direct route if the customer has a tank large enough to receive a full trailer load. Delivery to the same customer may be combined with others such that the customer is replenished on a multiple customer route in a time period with lower demand. Note that the high demand period for one customer may be the low demand period for another.

The industrial gas company is assumed to own a homogeneous feet of trailers with a given capacity to distribute products to customers, and sufficient trailers are available to deliver gas products. Trailers start at depots, obtain product from gas sources, and then deliver the gas to customers assigned to those depots and sources. Trailers can provide continuous delivery by visiting sources in between deliveries to refill their tanks. In this context, the entire tour from depot departure to arrival back at the depot is referred to as a *route*, and the set of customers served between source visits is a single *trip*. The location of customers, depots, and bulk tank warehouse are known and customers are assigned to sources and depots. The initial number of tanks available in the tank warehouse is known as well as a budget for purchasing new tanks and refurbishing existing tanks located at customer locations. We assume that customer demand for each period is given and each customer is replenished from the depot to which it is assigned.

The model minimizes the sum of the initial cost of tank installation, tank transportation from the warehouse, tank refurbishment, and new tank purchase cost and the present value of the periodic distribution costs in each time period. Other costs associated with production, inventory holding, and back order costs are not considered in this study as our main focus is on tank allocation and product transportation costs involved in distribution planning. To account for customer demand variability the time horizon is divided into shorter, equal, independent, time periods (e.g. months or quarters). The challenge is to allocate an expensive resource over a long time horizon (e.g., 7 years) while accommodating customer demand that potentially fluctuates over each time period (e.g., months or quarters). Although the focus of this study is for a single gas type, the model is adaptable for multiple gas types. The objective is to assign tanks to customer sites to minimize the tank installation, tank reallocation, tank refurbishment, and tank purchasing costs over a time horizon and the present value of gas distribution costs. The primary assumptions of the model are highlighted below.

Assumptions:

- Locations of customers, depots, and tank warehouses are known.
- Customer demands and required safety stocks are known for each time period (e.g., quarterly) in the study time horizon (e.g., seven years).
- Customer demand is constant during a time period but may vary across time periods.
- Customer demand must be satisfied in each time period.
- The industrial gas producer owns the tanks that are allocated to customer sites. Customers assume ownership of the gas product upon delivery.
- The capacity and quantity of each tank type, which is either assigned to a customer or available in the tank warehouse, is known.
- Tanks that are currently at customer sites or available in the tank warehouse can be allocated or reallocated, and new tanks can be purchased if necessary at the beginning of the time horizon. Tanks remain in place throughout the time horizon.
- Swapping tanks between customers involves removing a tank from one customer site, moving it to a tank warehouse, refurbishing it, transporting it to another customer location, and installing it at the new location.
- Allocating a tank from the warehouse to a customer also involves refurbishment costs, transportation costs, and tank change out costs.
- The capacity and quantity of the homogeneous delivery trailers are known.
- The relevant costs are known, including delivery cost per distance, fixed access cost at customer sites, tank installation costs, refurbishment costs, and purchase costs.
- The duration of each route must comply with the mandatory maximal working hours of drivers, and the aggregated duration of all of the selected routes must not exceed the product of the length of each time period and the number of trailers.
- The budget limits for tank procurement and tank refurbishment are known.

3.1 Notation

The notation used in the proposed model is summarized as follows:

Sets and indices:

P	set of depots indexed by p.
Q	set of periods indexed by q.
T	set of tank types, indexed by t.
R	set of possible routes, indexed by r.
S	set of trips, indexed by s.
S_r	subset of trips composing route r.
R_p	subset of routes pertaining to depot p.
I	set of customers, indexed by i, j.
I_p	subset of customers assigned to depot p.
T_i	subset of tanks considered for customer i.
0_p	index for depot p, $\forall p = 1, \ldots, \lvert P \rvert$.

Depot and trailer parameters:

m	distribution cost per unit distance traveled for trailers.
g_p	volume capacity of each trailer in the fleet of depot p.
v_p	average travel speed of trailers in depot p area.
k_p	number of trailers available at the depot p.

Tank parameters:

n_t	the number of tanks of type t available at the warehouse.
v_t	volume of tank type t.
v_{max}	maximum tank volume where $v_{max} = \max_{t \in T} v_t$.
c_t	cost to purchase tank type t.
a_{it}	indicates if tank type t is initially assigned to customer site i.
b_{wt}	cost to change from tank type w to tank type t at a customer site.
ϕ_t	cost to refurbish tank type t.
λ_i	cost per distance to transport a tank to customer i.

Customer parameters:

δ_{iq} mass demand for customer i during period q.
σ_i mass amount of safety stock for customer i.

d_{ij} distance from site i to site j, where $i, j \in I \cup \{0\}, i \neq j$.
ρ_i working density for customer i (determined by working pressure and product type).
f_i fixed cost to visit customer i.
h_i fixed time for delivery at customer i.
w_i distance to customer i's assigned warehouse.

Route Parameters:

φ_r distribution cost of route r.
t_r time for route r to be executed by a trailer.
y_{is} binary indicator that equals 1 if customer i is visited on trip s, and 0 otherwise.
z_{ijs} binary indicator that equals 1 if customer i is immediately followed by customer j on trip s, and 0 otherwise.

Economic parameters:

ι periodical discount rate.
β_{purch} total budget allocated to purchase new tanks during the time horizon.
β_{refurb} total budget allocated for refurbishment during the time horizon.

Time parameters:

η length of the planning horizon in time periods.
τ_{max} maximum allowable time for a route.
κ length of each period.
ξ number of periods in one year.

Decision variables:

X_{it} binary variable that equals 1 if customer i is allocated a tank of type t, and 0 otherwise.
N_{it} binary variable that indicates if a new tank of type t is purchased for customer i.
M_{it} binary variable that indicates if customer i receives a tank of type t from the warehouse, and 0 otherwise.
Ψ_{rq} binary variable that equals 1 if route r is selected during period q, and 0 otherwise.
D_{isq} amount of gas delivered to customer i on trip s during period q (continuous).

3.2 Model Formulation for BTATVD

The bulk tank allocation for time varying demand (BTATVD) model captures complex trade-offs between strategic level investment decisions and operational level routing decisions. The model simultaneously allocates tank types, selects distribution routes from a set of candidate routes, and determines delivery amounts. For a set of all potential routes, the problem is formulated as follows:

BTATVD: Minimize

$$\sum_{p\in P}\sum_{r\in R_p}\sum_{q\in Q}\varphi_r\Psi_{rq}/(1+\iota)^q+\sum_{p\in P}\sum_{i\in I_p}\sum_{\{\omega,t\}\subseteq T}a_{i\omega}b_{\omega t}X_{it}+\sum_{p\in P}\sum_{i\in I_p}\sum_{t\in T}2\lambda_i\omega_i(M_{it}+N_{it})$$
$$+\sum_{p\in P}\sum_{i\in I_p}\sum_{t\in T}\phi_t M_{it}+\sum_{p\in P}\sum_{i\in I_p}\sum_{t\in T}C_t N_{it} \tag{1}$$

Subject to:

$$\sum_{i\in I}M_{it}\le n_t+\sum_{i\in I}a_{it}(1-X_{it}), \qquad \forall t\in T, \tag{2}$$

$$X_{it}(1-a_{it}) = M_{it}+N_{it}, \qquad \forall t\in I, \forall t\in T, \tag{3}$$

$$\sum_{i\in I}\sum_{t\in T}\phi_t M_{it}\le \beta_{refurb}, \tag{4}$$

$$\sum_{i\in I}\sum_{t\in T}C_t N_{it}\le \beta_{purch}, \tag{5}$$

$$\sum_{t\in T}X_{it}=1, \qquad \forall i\in I, \tag{6}$$

$$D_{isq} + \sigma_i \leq \sum_{t\in T} \rho_i v_t X_{it}, \qquad \forall i \in I, s \in S, q \in Q, \quad (7)$$

$$\sum_{i\in Ip} y_{is} D_{isq} \leq g_p \Psi_{rq}, \qquad \forall r \in R_p, s \in S_r, p \in P, q \in Q \quad (8)$$

$$\sum_{s\in S_r} \sum_{r\in R_p} y_{is} D_{isq} = \delta_{iq}, \qquad \forall i \in I, p \in P, q \in Q, \quad (9)$$

$$D_{isq} \leq min\{g_p, \rho_i v_{max} - \sigma_i\} y_{is}, \qquad \forall i \in I, s \in S, q \in Q, \quad (10)$$

$$\sum_{s\in S_r} \sum_{r\in R_p} y_{is} \Psi_{rq} \geq \sum_{t\in T} max\left\{\left\lceil \frac{\delta_{iq}}{g_p} \right\rceil, \left\lceil \frac{\delta_{iq}}{\rho_i v_t - \sigma_i} \right\rceil\right\} X_{it}, \qquad \forall i \in I, p \in P, q \in Q, \quad (11)$$

$$\sum_{r\in R_p} t_r \Psi_{rq} \leq \kappa k_p, \qquad \forall p \in P, q \in Q, \quad (12)$$

$$t_r \Psi_{rq} \leq \tau_{max}, \qquad \forall r \in R, q \in Q, \quad (13)$$

$$X, \Psi, M, N \; binary, D \geq 0, \quad (14)$$

The objective function (1) minimizes the net present value of the recurring distribution costs and one-time tank installation, tank transportation from the warehouse, tank refurbishment, and tank purchase costs. Constraint set (2) ensures that the number of tanks of type t moved from the warehouse does not exceed the warehouse inventory and the tanks returned to the warehouse. For each customer i, with a new assigned tank type t, the tank can be either obtained from the warehouse or newly purchased, as stated by constraint set (3). Constraint sets (4) and (5) ensure that the budget limits for tanks refurbishment and procurement are not exceeded. Each customer i must be assigned a tank, as stated by constraint set (6). Constraint set (7) ensures that the amount of each delivery to customer i on trip s in period q plus the safety stock requirements for that customer does not exceed the allocated tank mass capacity. Constraint set (8) ensures that the amount delivered to all the customers on trip s in each period is less than or equal to the trailer capacity, g, when the route containing trip s is selected, or zero otherwise. The delivery amounts to customer i across all routes r in each period q must meet the requirements for customer i in that period, as enforced by constraint set (9). Constraint set (10) ensures that the amount of gas delivered to customer i on trip s in period q is less than or equal to the minimum of the trailer capacity or the largest tank size less the safety stock for customer i, or zero if customer i is not visited on trip s in period q. The minimum number of visits required to the customer site in period q, is determined by the requirements for customer i divided by the trailer capacity or the requirements for customer i divided by the tank capacity less the safety stock (all rounded up), as stated in constraint set (11). Constraint set (12) ensures that the time required for all the selected routes in each time period does not exceed the length of the time period κ multiplied by the number of available vehicles, k_p, for each depot p. Constraint set (13) ensures that routes with duration longer than the maximum allowable time for a route are not selected. Constraint set (14) represents the logical restrictions on the decision variables.

Moreover, the cost for route $r \in R$ to be executed by a trailer includes the cost to access customer sites and the cost to transport the gas on the distribution route. This cost is defined as follows:

$$\varphi_r \equiv \sum_{i\in I} \sum_{s\in S_r} f_i y_{is} + \sum_{i,j\in I \cup O_p, i\neq j} \sum_{s\in S_r} m\, d_{ij} z_{ijs} \quad (15)$$

The time for a route, $r \in R$, to be executed by a trailer is defined as follows:

$$t_r \equiv \sum_{i\in I} \sum_{s\in S_r} h_i y_{is} + \sum_{i,j\in I \cup O_p, i\neq j} \sum_{s\in S_r} \frac{d_{ij} z_{ijs}}{v_p} \quad (16)$$

4 ROUTE GENERATION

In the BTATVD model formulation, set R includes all the possible routes that start at a depot, visit a source, deliver to at least one customer, and return to the depot. The number of routes for each period increases exponentially as the number of customers increases, which results in an intractable problem even for a small to moderate sized problem. Thus, we strive to create a set of high quality realistic routes that are sufficiently large to ensure solution quality while maintaining tractability.

Two types of routes are generated: single-customer routes that contain a direct trip from a depot to a customer and back to the depot, and multiple-customer routes in which more than one customer is visited on a route. To generate the multiple-customer routes, the customers are clustered and ordered using a sweep heuristic. The algorithm is an adaptation of the sweep and petal methods used to solve VRP (Gillett and Miller, 1974; Ryan et al., 1993). For each depot, the customers are first sorted by corresponding polar coordinates. Then the customers are added to a cluster according to the polar coordinates using either a clockwise or counter clockwise direction until the total residual partial demands in a cluster reaches the trailer capacity. A traveling salesman problem (TSP) can be solved to generate the sequence of customer visits in each cluster. For this implementation, however, the sequence of customer visits is based on the order the customers are added to the cluster. The cluster and route generation process is repeated for all of the depots and sources. The output is a set of routes with their corresponding cost and duration which are input to the restricted bulk tank allocation problem to determine the allocation of bulk tanks to customer sites. To ensure feasibility, multiple copies of the generated routes are input to the bulk tank allocation model.

Estimating the product delivery cost plays a pivotal role in the bulk tank allocation decisions. The more accurate the estimate of the cost to serve a customer, the easier it is to determine the effect of a tank change on the overall distribution network. To more thoroughly represent the industrial gas distribution practices, we extend the routing approach of Ellis et al. (2014) by differentiating between depots and gas sources and considering the industry practice of continuous delivery.

- ***Depots and sources***: The routing subproblem in Ellis et al. (2014) generates candidate routes assuming depots and sources are co-located such that a full trailer departs from a depot, delivers product to up to five customers, and then returns to the depot. In reality, a depot is the location in which trailers park and receive maintenance. In order to deliver product, trailers must first obtain liquefied product from a source, which may or may not be co-located with the depot. Customers may have additional product requirements such as purity of product provided by certain sources. The allocation of customers to sources is provided as input to ensure that these requirements are met.
- ***Continuous delivery:*** The consideration of sources allows for the possibility for a trailer to refill product during a delivery route. Under certain circumstances, a trailer may refill at a source and serve additional customers before returning to the depot. Modeling this continuous delivery practice provides a more reflective estimate of the distribution costs in practice. Two common delivery practices considered in the network are shown in Figure 1. A trailer following the route (a) fills at source S1, visits two customers, refills at source S1, then delivers to three more customers before returning to the depot. A trailer following route (b) fills at source S1, visits two customers, refills at source S2, then delivers to three customers before returning to the depot. Note that in route (b), only customers assigned to receive product from this different source may be considered after refill.

When generating routes for the model to evaluate, we incorporate this continuous delivery practice. A trailer obtains product from a source, then it begins a polar sweep to add additional customers to the trip. Eventually, the trailer may consider refilling on product if the product level is below a threshold.

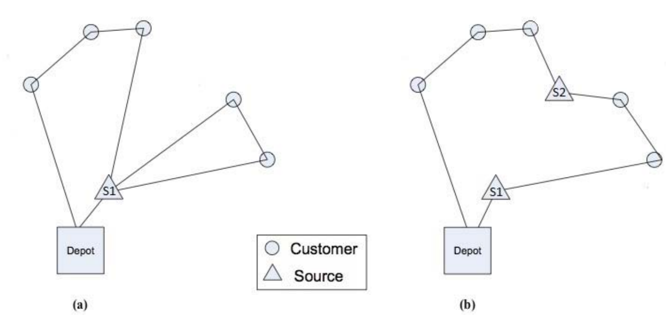

Figure 1: Two examples of routes with sources

The algorithm searches the set of sources allocated to serve customers assigned to the current depot for the one nearest the trailer's current location. If this source is different from the last one, the trailer refills on product and proceeds to serve customers assigned to this new source.

The objective of the route generation module is to generate a set of high quality feasible routes quickly. Details of the route generation approach are presented in Appendix 1. The generated routes along with their associated costs and duration are input to a BTATVD restricted master problem (RMP). The model then assigns tanks to customers, selects delivery routes, and determine the delivery amount to each customer on each route.

5 SOLUTION APPROACHES

The BTATVD problem can be viewed as a variant of a multidimensional knapsack problem which belongs to the category of NP-hard problems, along with additional constraints. Therefore, the BTATVD problem also belongs to the category of NP-hard problems. Considering the budget limits for tank purchase and tank (re-)allocations, the goal is to identify a subset of customers, for which different tanks will contribute the most to the distribution cost savings. Given the complexity of the BTATVD problem, we propose the following three solution approaches for the BTATVD problem:

- BTATVD Restricted Master Problem (BTATVD-RMP),
- BTATVD with Recurring Demand (BTATVD-R), and
- Periodically Restricting Heuristic (PRH).

5.1 BTATVD Restricted Master Problem (BTATVD-RMP)

Based on a restricted set of routes created using the route generation approach, the restricted master problem determines the allocation of bulk tanks to customers while considering their time varying demand. The model, which is described in section 3.2, also selects delivery routes for each time period from the set of potential routes and determines the delivery amount on each route in each time period. The objective function value provides an estimate for total tank allocation costs and the net present value of the distribution costs.

5.2 BTATVD with Recurring Demand (BTATVD-R)

In many practical situations, demand varies within a year but is stationary over multiple years, such that the demand in each period of year one is equal to the demand in corresponding periods of subsequent years. In this case, we can solve the BTATVD model for a single year while considering the overall time horizon. We modify the objective function using an economic conversion factor (PV) to account for the time horizon length and convert recurring costs to present value using effective interest rate. Effective annual interest rate is calculated as $(1+\iota)^{\xi}-1$, where ι represents periodical discount rate and ξ represents number of periods in one year. When values recur annually, their present value is calculated as follows:

$$PV = 1 + \frac{(1+\iota)^{\eta-\xi}-1}{\left((1+\iota)^{\xi}-1\right)(1+\iota)^{\eta-\xi}} \tag{17}$$

Note that we assume that the first year cost occurs at the beginning of the year. The objective function is revised as follows:

$$\begin{aligned}\text{Minimize } PV\text{:}\ & \frac{\sum_{p\in P}\sum_{r\in R_p}\sum_{q\in Q}\varphi_r\Psi_{rq}}{(1+\iota)^q} \\ &+\sum_{p\in P}\sum\nolimits_{i\in I_p}\sum_{\{\omega,t\}\subseteq T} a_{i\omega}b_{\omega t}X_{it} \\ &+\sum_{p\in P}\sum\nolimits_{i\in I_p}\sum_{t\in T} 2\lambda_i w_i(M_{it}+N_{it}) \\ &+\sum_{p\in P}\sum\nolimits_{i\in I_p}\sum_{t\in T}\phi_t M_{it} \\ &+\sum_{p\in P}\sum\nolimits_{i\in I_p}\sum_{t\in T} c_t N_{it},\end{aligned} \tag{18}$$

5.3 Periodically Restricting Heuristic (PRH)

For many industry representative data sets, the problem size becomes intractable for the BTATVD-RMP. Thus, we explore a periodically restricting heuristics (PRH) to address industry representative problem sizes. Our PRH approach for the BTATVD problem includes the following main phases:

- Solve bulk tank allocation period by period to determine candidate tanks for each customer across all time periods, and
- Determine tank allocations for varying demand over multiple periods from the restricted set of tanks.

For the expected demand in each period, the preferred tank size is allocated to each customer by solving the BTA model for expected demand in that period. Then across all periods, each customer has a set of candidate suitable tanks. These candidate tanks are considered when solving the tank allocation problem across the multiple time periods. Thus, the size of the solution space is decreased by restricting the number of tank types that the model considers for each customer. The overall PRH approach is illustrated in Figure 2 and details are provided in the following sections.

5.3.1 Solve bulk tank allocation with constant demand period by period to determine candidate tanks

The objective of phase I in our solution approach is to develop a set of preferred candidate tanks for each customer. In this phase, the bulk tank allocation problem is solved for each period assuming that the monthly demand over the time horizon is equal to the monthly demand in that period. We adapt the decomposition approach presented in Ellis et al. (2014) in which first a set of high quality routes are generated. Using the output from this modified route generation algorithm, we solve the restricted bulk tank allocation problem to determine the allocation of bulk tanks to customer sites for each period while considering customer locations and demand in relationship to other customers in the same proximity. The resulting tank assignments for the individual periods are used to establish a set of candidate tanks, T_i, for each customer i.

5.3.2 Determine tank allocations for varying demand over multiple periods

In phase II, we solve a restricted multiperiod bulk tank allocation problem, where the model selects from one of the tanks in set T_i for each customer i. Potential routes are generated using the sweep algorithm assuming that each customer's demand is equal to their maximum demand across all time periods to ensure sufficient routes are available during each time period. Constraint sets (6) and (7) are revised as follows:

$$\sum_{i \in T} X_{it} = 1, \qquad \forall i \in I, \qquad (19)$$

$$D_{isq} + \sigma_i \leq \sum_{i \in T} \rho_i v_t X_{it}, \qquad \forall i \in I, s \in S, q \in Q, \qquad (20)$$

Constraint set (19) assigns exactly one tank from each customer's set of candidate tanks to that customer, and constraint set (20) ensures that the amount of each delivery to customer i on route r in period q plus safety stock requirements for that customer does not exceed the mass capacity of the allocated tank from the set of candidate tanks. Figure 2 summarizes the overall solution approach.

6 CASE STUDY RESULTS

To evaluate the performance and scalability of the proposed models approaches we conducted analyses based on industrial data sets. We compare the performance of BTATVD-RMP model with the model with recurring demand assumption (BTATVD-R) and the proposed PRH. In these case studies, we use industry representative cost data, omit specific information about customers, and omit monetary units to protect sensitive information from our industrial partner. These data sets range from 50 to 818 customers and the number included in the case name indicates the number of customers in that data set. In the test cases, the customers are assigned to 8 different depots. We consider 21 different tank types with capacities ranging from 1.2 m^3 to 150 m^3, and refurbishment costs are assumed to range from 6,800 to 12,000. For each data set, the routes generated are the same for each solution approach. In scenario I, we consider a 3 year time horizon with 3 month periods (i.e. 4 periods in a year). For scenario II, we extend the length of study time horizon to 7 years with 3 month periods. Table 1 summarizes the parameter values of our test cases. The data sets were analyzed on a workstation with two Intel Xeon 3.10GHz quad-core processors, 32GB RAM, and CPLEX v12.5.

In Table 2, we present the results obtained for each scenario. First we use the BTATVD-RMP model, with the original objective function (1) in which the demand is assumed to be changing for all the periods in the time horizon. Then we use the BTATVD-R model, with the modified objective function (18) where we assume that the demand is recurring. In order for the values to be comparable, we assume that the demand for years 2 and 3 remain the same as the first year. Finally, we test the performance of the PRH by comparing the results obtained from this approach to the other two models.

For the BTATVD-RMP and BTATVD-R approaches, we solve the model to within an optimality gap tolerance of 2% for the set of routes generated. We capture the lower bound at which the optimization model stops. Then we use the maximum of the lower bounds (obtained from BTATVD- RMP and BTATVD-R) as the best known lower bound of the restricted models. The best known lower bound is used to determine the gap between the obtained solution and the lower bound as follows:

$$\text{Gap} = \frac{\text{Estimated Final Cost} - \text{Lower Bound}}{\text{Lower Bound}} \qquad (21)$$

To provide an estimate of the resulting savings yielded by our approaches for the company, we compare the total costs to the distribution costs with the initial tank allocation which we refer to as the *fixed-tank approach*. With the fixed-tank approach, we solve the BTATVD-RMP model but fix the initial tank allocation for each customer to estimate the distribution costs of the initial state.

As shown in Table 2, the recurring demand assumption with the modified objective function results in substantially reduced computational time for model BTATVD-R compared to model BTATVD-RMP. The PRH approach provides comparable objective function values in substantially less computational time than either BTATVD-RMP or BTATVD-R. In all cases, the cost estimates are within 2% of the estimated objective

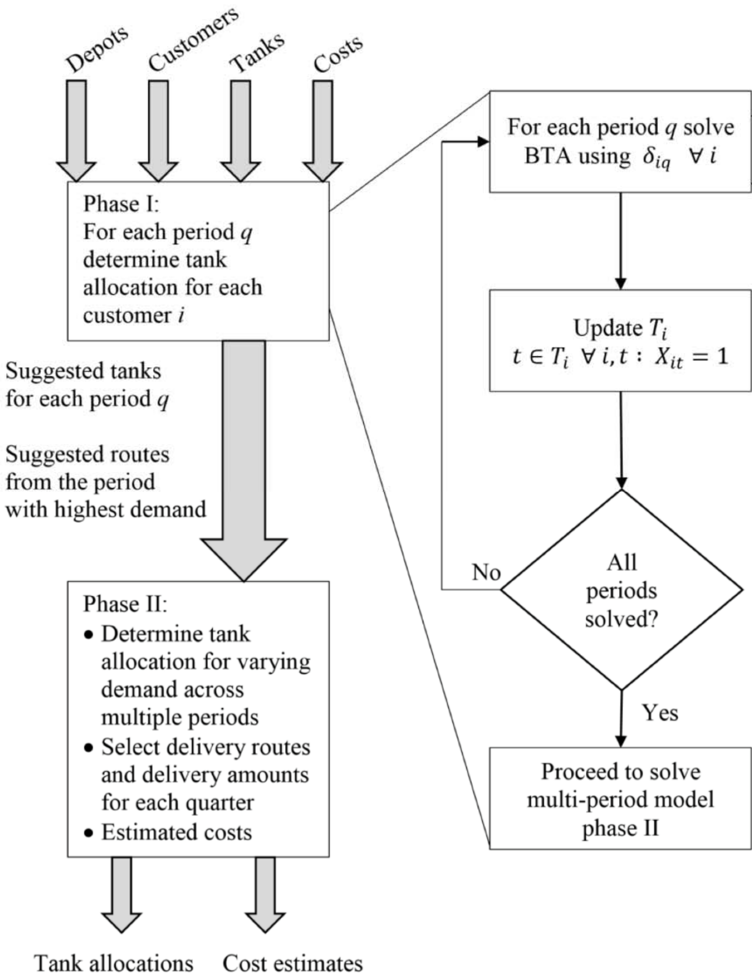

Figure 2: Bulk tank allocation for time-varying demand approach

Table 1: Case study description

Periodic Discount Rate	**Tank Purchase Budget**	**Tank Refurbishment Budget**	**Max Number of Tank Swaps**
3%	1,000,000	300,000	Unrestricted
Trailer Capacity (kg)	**Tank Change-out Cost**	**Tank Transportation Cost (per km)**	**Gas Transportation Cost (per km)**
22,000	8,000	1.8	0.8

Table 2: Performance comparison for Scenario I (3 year time horizon with 4 periods per year)

Case	Model	Estimated Final Cost	Annual Distribution Cost	Tank Exchange Cost	Tank Swaps	Run Time (min)	Gap (%)	Total Savings (%)	Distribution Savings (%)
IA-50	Fixed	1,002,886	374,505	0	0	0.1	0.38	--	--
IA-50	BTATVD-RMP	952,732	325,654	77,614	4	12.5	0.99	5.00	13.04
IA-50	BTATVD-R	947,598	324,876	77,614	4	2.2	0.45	5.51	13.25
IA-50	PRH	954,311	334,752	57,880	3	0.1	1.16	4.84	10.61
IB-100	Fixed	1,537,990	574,328	0	0	3.0	0.41	--	--
IB-100	BTATVD-RMP	1,452,971	498,539	117,898	6	419.4	0.83	5.53	13.20
IB-100	BTATVD-R	1,453,074	498,591	117,898	6	50.9	0.84	5.52	13.19
IB-100	PRH	1,461,597	504,575	110,398	6	4.5	1.43	4.97	12.15
IC-150	Fixed	2,849,469	1,064,070	0	0	3.5	0.55	--	--
IC-150	BTATVD-RMP	2,727,771*	942,269	211,611	11	1,287.5	2.89	4.27	11.45
IC-150	BTATVD-R	2,682,977	921,845	214,372	11	751.0	1.00	5.84	13.37
IC-150	PRH	2,716,825	950,658	171,063	9	3.6	2.27	4.66	10.66
ID-200	Fixed	3,234,234	1,207,752	0	0	14.9	0.46	--	--
ID-200	BTATVD-RMP**	--	--	--	--	--	--	--	--
ID-200	BTATVD-R	3,065,313	1,072,233	193,984	10	1,372.0	1.00	5.22	11.22
ID-200	PRH	3,081,151	1,092,022	156,828	8	7.6	1.15	4.73	9.58

* The incumbent objective value upon premature termination due to memory limitations.
** Solution process was prematurely terminated due to memory limitations before finding a feasible solution.

function obtained from solving the other two models and within 3% of the best-known lower bound.

For scenario I, where the study time horizon is 3 years, the approaches yield between 4%-6% reduction in the overall cost with annual distribution costs reduced by 9%-14% compared to the costs of initial tank allocations (obtained from the fixed-tank approach). For scenario II, the study time horizon is increased to 7 years. The total cost savings are 8%-11%, with 14%-20% reduction in transportation costs. The extended time horizon provides a longer payback period for the initial investments for the tank reallocation. Therefore, compared to a study horizon of 3 years, more tank exchanges are recommended and the total savings are more significant.

For these data sets, the PRH approach is capable of solving larger instances within reasonable computational time in comparison to the other two models and enables the industrial gas production company to consider larger clusters of their customers while capturing their demand fluctuations. Considering that an industrial gas production company typically has thousands of customers around the world, these approaches could improve their operational efficiency and result in significant savings in their distribution network.

For the previous results, the internal rate of return (IRR) is used as a financial metric to provide managerial insights on the potential benefits of a more efficient bulk tank allocation. Table 4 summarizes the IRR for the test cases in scenarios I and II. Using the results from the PRH approach, the tank exchange costs are considered as initial investments (negative cash flow) and annual distribution savings over the fixed case are considered as potential revenue (positive cash flow). As shown, the IRR values vary from 38-53%.

Table 3: Performance comparison for Scenario II (7 year time horizon with 4 periods per year)

Case	Model	Estimated Final Cost	Annual Distribution Cost	Tank Exchange Cost	Tank Swaps	Run Time (min)	Gap (%)	Total Savings (%)	Distribution Savings (%)
IA-50-7	Fixed	1,889,431	374,289	0	0	0.1	0.32	--	--
IA-50-7	BTATVD-RMP	1,683,962	298,743	175,947	9	15.2	0.78	10.87	20.18
IA-50-7	BTATVD-R	1,687,302	299,394	175,947	9	1.3	0.98	10.70	20.01
IA-50-7	PRH	1,706,968	311,253	135,745	7	0.1	2.16	9.66	16.84
IB-100-7	Fixed	2,907,591	575,983	0	0	1.1	0.71	-	-
IB-100-7	BTATVD-RMP	2,592,024	464,190	250,580	13	1,026.4	1.67	10.85	19.41
IB-100-7	BTATVD-R	2,598,255	472,887	211,098	11	7.7	1.91	10.64	17.90
IB-100-7	PRH	2,631,484	473,013	243,690	13	1.8	3.21	9.50	17.88
IC-150-7	Fixed	5,372,422	1,064,256	0	0	3.3	0.54	--	--
IC-150-7	BTATVD-RMP**	--	--	--	--	--	--	--	--
IC-150-7	BTATVD-R	4,859,519	890,488	364,287	19	203.7	1.96	9.55	16.33
IC-150-7	PRH	4,864,245	903,416	303,756	16	49.9	2.06	9.46	15.11
ID-200-7	Fixed	6,099,437	1,208,275	0	0	14.8	0.53	--	--
ID-200-7	BTATVD-RMP**	--	--	--	--	--	--	--	--
ID-200-7	BTATVD-R	5,721,307*	1,037,201	485,462	22	1,267.0	4.83	6.20	14.16
ID-200-7	PRH	5,556,286	1,029,227	360,695	19	8.4	1.81	8.90	14.82

* The incumbent objective value upon premature termination due to memory limitations.
** Solution process was prematurely terminated due to memory limitations before finding a feasible solution.

.

Table 4: Internal rate of return

Scenario	Case	IRR (%)
I - 3 year time horizon	IA-50	47
	IB-100	40
	IC-150	44
	ID-200	53
II - 7 year time horizon	IA-50-7	43
	IB-100-7	38
	IC-150-7	50
	ID-200-7	46

7 SENSITIVITY ANALYSIS

In this section, we further demonstrate the benefits of the PRH approach to gain additional managerial insights. We also perform sensitivity analysis on periodic discount rate (ι) and refurbishment budget (β_{refurb}).

To demonstrate the benefits of the capability to address time varying demand, we compare the results from the PRH approach with those of the BTA approach which assumes constant demand (Ellis et al., 2014). For this analysis, we compare the results of the current allocation (fixed), the results of the BTA approach based in constant demand, and the results of the PRH approach. For the BTA approach, the average demand for each customer is used as input and the preferred tanks are selected. Then the resulting investment and distribution costs of these selected tanks are evaluated using the seasonal demand values for each customer.

As shown in Table 5, the PRH approach outperforms both the initial state and the BTA approach. The results clearly demonstrate the value of using the PRH approach when customer demand varies over time. In most cases, the BTA approach that assumes constant demand actually outperforms the initial case. For the IA-50 case, however, the BTA approach results in less attractive tank allocations particularly with time varying demand compared to the current tank allocations when considering seasonality in customer demand. In general, these results demonstrate the value of having an optimization-based system to evaluate the complex tank allocation decisions.

We also analyze the sensitivity of the results from our approach to the period discount rate. Table 6 shows the results from our PRH model for the cases with periodic discount rates of 1%, 3%, and 5%. To analyze the results, we compare estimated total savings and distribution savings by solving the fixed model (assuming the current tank allocation remains unchanged) for the same data set and calculating the difference in total and distribution costs. As shown, the number of tank exchanges decrease as the periodic discount rate increases. As the periodic discount rate increases, the present values of the distribution costs decrease in comparison to the initial investment costs. Thus, as expected, higher periodic discount rates result in less justification for initial tank-related investments and tank swaps.

We also analyze the effect of more restricting refurbishment budgets on the results. Table 7 displays the results from the PRH model with limited refurbishment budgets. The refurbishment budgets are estimated based on the number of customers. For this analysis, we base the refurbishment budget on the number of customers served (by multiplying the number of customers by 400). In all scenarios, the investments in tank exchanges are less than the original case (with a refurbishment budget of 1,000,000), while the distribution costs increase. For scenario ID-200, the number of the tank swaps is the same as the original scenario, but one of the assigned tank types is different, and the tank exchange costs are lower and distribution costs are higher. In all scenarios, the budgetary limit results in slightly increased distribution costs and overall final costs.

8 SCALABILITY ANALYSIS

Using the PRH approach and the recurring demand assumption, we are able to solve instances for up to 200 customers in less than 10 minutes within 2% of best known lower bound. In reality, an international industrial gas company may have thousands of customers around the world and therefore needs to solve larger instances. To further improve the scalability of our solution approach, the formulation of the mathematical model is revised using a relaxation approach. In this approach, we relax the integrality of the variable for selecting routes, by changing the route selection variables, Ψrq, from binary to continuous, limited between 0 and 1. With this method, a route may be partially selected for fulfilling demand, leading to distribution costs that are potentially under-estimated. Thus solutions from this relaxed model will have objective function values (overall costs) less than or equal to the original models.

Table 5: Constant demand assumption versus time-varying demand
(3 year time horizon with 4 periods per year)

Case	Model	Estimated Final Cost	Annual Distribution Cost	Tank Exchange Cost	Tank Swaps
IA-50	Fixed	1,002,886	374,505	0	0
IA-50	BTA	1,076,318	312,067	240,633	8
IA-50	PRH	954,311	334,752	57,880	3
IB-100	Fixed	1,537,990	574,328	0	0
IB-100	BTA	1,494,739	523,786	92,095	5
IB-100	PRH	1,461,597	504,575	110,398	6
IC-150	Fixed	2,849,469	1,064,070	0	0
IC-150	BTA	2,795,688	992,693	137,358	7
IC-150	PRH	2,716,825	950,658	171,063	9
ID-200	Fixed	3,234,234	1,207,752	0	0
ID-200	BTA	3,162,062	1,122,771	155,397	8
ID-200	PRH	3,081,151	1,092,022	156,828	8

Table 6: Sensitivity analysis for periodic discount rate

Case	Estimated Final Cost	Annual Distribution Cost	Tank Swap Cost	Tank Swaps	Run Time (min)	Total Savings (%)	Distribution Savings (%)
5% Periodic Discount Rate							
IA-50	871,482	326,300	55,880	3	0.11	4.19	10.34
IB-100	1,334,328	496,556	93,164	5	4.33	4.34	11.02
IC-150	2,464,031	923,288	156,234	8	3.44	4.65	10.69
ID-200	2,806,822	1,067,487	138,594	7	34.34	4.22	8.95
3% Periodic Discount Rate (Scenario I)							
IA-50	954,311	334,752	57,880	3	0.08	4.84	10.61
IB-100	1,461,597	504,575	110,398	6	4.51	4.97	12.15
IC-150	2,716,825	950,658	171,063	9	3.61	4.66	10.66
ID-200	3,081,151	1,092,022	156,828	8	7.59	4.73	9.58
1% Periodic Discount Rate							
IA-50	1,058,253	334,084	94,599	5	0.11	4.99	13.49
IB-100	1,623,150	524,448	110,398	6	1.55	5.00	11.46
IC-150	2,974,861	945,383	247,939	13	18.39	6.13	13.95
ID-200	3,409,012	1,108,613	211,260	11	7.04	5.14	11.02

Table 7: Sensitivity analysis for refurbishment budget

Case	Refurbishment Budget	Estimated Final Cost	Annual Distribution Cost	Tank Swap Cost	Tank Swaps	Run Time (min)	Total Savings (%)	Distribution Savings (%)
IA-50	20,000	967,189	346,757	38,609	2	0.07	3.56	7.41
IA-50	1,000,000	954,311	334,752	57,880	3	0.08	4.84	10.61
IB-100	40,000	1,469,246	519,848	77,148	4	1.81	4.47	9.49
IB-100	1,000,000	1,461,597	504,575	110,398	6	4.51	4.97	12.15
IC-150	60,000	2,754,806	969,157	159,504	7	11.39	3.32	8.92
IC-150	1,000,000	2,716,825	950,658	171,063	9	3.61	4.66	10.66
ID-200	80,000	3,088,362	1,094,937	156,234	8	6.74	4.52	9.35
ID-200	1,000,000	3,081,151	1,092,022	156,828	8	7.59	4.73	9.58

Because the focus of the BTATVD model is on a strategic problem of allocating tanks, these estimates of distribution cost may be suitable. If desired, the distribution costs for the resulting tank allocations can be further evaluated. For this purpose, we use the resulting tank allocation decisions and input them to the BTATVD-RMP model as given (referred to as the fixed-tank approach).

We applied this relaxation approach to some of the test instances from Table 2 along with two larger case. As shown in Table 8, the relaxation approach solves the data sets in substantially less computational time. Using the tank allocations obtained from the relaxation method, we solve the fixed-tank approach assuming that the allocated tanks are given and the routes selection variables are binary. This fixed-tank approach provides estimates of distribution costs for the suggested tank allocation consistent with previous results. Table 8 summarizes fixed-tank approach results with new tank allocations for some of the test instances.

For example, case IB-100 was solved using BTATVD-R approach with resulting costs of 1,453,074 and computational time of 50.92 minutes (from Table 2). The relaxed model solves in approximately 0.09 minutes as shown in Table 4. The tank allocations from the relaxed model are then input as fixed solutions in the BTATVD-R model with binary route selection variables to obtain estimated final cost of 1,455,272, which is only 0.15% higher than the original BTATVD-R approach. Thus, with the relaxed approach, we are able to solve larger data sets in reasonable time and the resulting tank allocation and distribution costs are similar to the approaches that use integrated variables.

9 CONCLUSIONS AND FUTURE WORK

The allocation of bulk tanks to customer sites is an important strategic level decision that has direct impact on the operational level distribution planning for industrial gas companies. Through this collaborative university and industry research project, we have developed a mixed integer programming model that addresses the bulk tank allocation problem for time varying demand while capturing the operational characteristics of the actual delivery process. Given the complexity of the problem, we develop a two phase heuristic approach to address industry representative problem instances. In the first phase of this approach, the problem is solved for each period under the assumption that customer demand is constant and equals the demand in that period. The result of this phase is a suggested tank size for each customer in each period and also a set of suggested routes which are generated via a sweep heuristic for the period with highest demand. In phase two, the bulk tank allocation model under varying demand is solved using CPLEX which determines a preferred tank allocation only considering the suggested set of tanks from phase one. This phase also selects delivery routes and amount of gas to be delivered to each customer on each route along with the estimated distribution and investment costs.

The effectiveness of the solution approach was demonstrated using industry representative data sets provided by Air Liquide for both 3-year and 7-year time horizons. For the data sets analyzed, several managerial insights emerged. As the time horizon increases, more tank changes are justified and the potential savings increase. As the discount rate increases, less tank exchange investments are justified. As the refurbishment budget is reduced, fewer tank exchanges are recommended, resulting in increased distribution costs and overall costs.

Table 8: Results with relaxed variables

		Original	**Relaxed Approach**			**Relaxed - Fixed Tanks**	**Original**
Case	**Solution Approach**	**Estimated Cost**	**Tank Swaps**	**Run Time (mins)**	**Estimated Final Cost**	**Estimated Final Cost**	**% Difference**
IA-50	BTATVD-R	940,322	4	0.02	952,560	947,598	0.52
IA-50	PRH	969,136	3	0.01	978,130	954,311	2.43
IB-100	BTATVD-R	1,426,087	6	0	1,455,272	1,453,074	0.15
IB-100	PRH	1,475,237	4	0.02	1,503,454	1,461,597	2.78
IC-150	BTATVD-R	2,660,253	12	0.24	2,706,151	2,682,977	0.85
IC-150	PRH	2,679,351	8	0.05	2,726,305	2,716,825	0.35
ID-200	BTATVD-R	3,015,844	12	1.22	3,075,787	3,065,313	0.34
ID-200	PRH	3,020,727	11	0.12	3,101,099	3,081,151	0.64
IE-400	BTATVD-R	4,940,615	16	4.53	5,002,467		
IE-400	PRH	5,024,092	16	0.96	5,078,950		
IF-818	BTATVD-R	10,556,324	31	28.36	10,699,679		
IF-818	PRH	10,652,538	29	4.6	10,832,886		

This model has also been used to support multi-product optimization. Note that the capacity of each tank type is described independently of product, using water volume. When installed at a customer site, the mass capacity of a tank is determined based on product type stored and pressure of the customer site. If a customer requires three different product types, then this is represented as three different customers in the model. Likewise, trailer mass capacity is established by product type. If a single physical depot distributes two different product types, this is represented as two different depots for the model. Thus product type is essentially a depot attribute, and customers are assigned to depots which provide the required product. Routes are then generated for customers by depot. The resulting model provides an assessment of the preferred allocation of tanks across multiple products.

Future research will be devoted to exploring approaches to improve the scalability of the solution approach and quality of the obtained solutions. One possible improvement of the algorithm could involve using an alternative route generation method instead of the sweep heuristic which could produce higher quality routes depending on the geographical distribution of customers. We also plan to assess the benefits of tank allocations in addition to distribution cost savings by evaluating additional key performance measures.

ACKNOWLEDGMENTS

This research has been conducted as part of the Center for Excellence in Logistics and Distribution at Virginia Tech in collaboration with Air Liquide. The authors gratefully acknowledge the support received for this research from Air Liquide and the National Science Foundation under grant IIP-0801919. We also gratefully acknowledge Jared Schwalbe from Virginia Tech and Mukesh Rungta at Air Liquide for their contributions in implementing the solution approaches described in this manuscript.

REFERENCES

1. Aksoy, Y. and Derbez, A. (2003). "Software survey: supply chain management." OR MS TODAY, 30 (3), pp. 34–41.
2. Aleman, R.E. and Hill, R.R. (2010). "A tabu search with vocabulary building approach for the vehicle routing problem with split demands." International Journal of Metaheuristics, 1 (1), pp. 55–80.
3. Aleman, R.E., Zhang, X., and Hill, R.R. (2010). "An adaptive memory algorithm for the split delivery vehicle routing problem." Journal of Heuristics, 16 (3), pp. 441–473.
4. Andersson, H., Hoff, A., Christiansen, M., Hasle, G., and Løkketangen, A. (2010). "Industrial aspects and literature survey: Combined inventory management and routing." Computers and Operations Research, 37 (9), pp. 1515–1536.
5. Archetti, C. and Speranza, M. (2012). "Vehicle routing problems with split deliveries." International transactions in operational research, 19 (1-2), pp. 3–22.
6. Archetti, C., Bianchessi, N., and Speranza, M.G. (2011). "A column generation approach for the split delivery vehicle routing problem." Networks, 58 (4), pp. 241–254.
7. Archetti, C. and Speranza, M.G. (2008). "The split delivery vehicle routing problem: A survey." In The vehicle routing problem: latest advances and new challenges. Springer, pp. 103–122.
8. Baker, A.H. and Garvey, M.D. (2004). "Bulk Products Supply Chain: Complexity in Motion." CryoGas International, 42 (3), pp. 20–35.
9. Bard, J.F., Huang, L., Jaillet, P., and Dror, M. (1998). "A decomposition approach to the inventory routing problem with satellite facilities." Transportation science, 32 (2), pp. 189–203.
10. Bertazzi, L., Savelsbergh, M., and Speranza, M.G. (2008). "Inventory routing." In The vehicle routing problem: latest advances and new challenges. Springer, pp. 49–72.
11. Campbell, A., Clarke, L., Kleywegt, A., and Savelsbergh, M. (1998). "The inventory routing problem." In Fleet management and logistics. Springer, pp. 95–113.
12. Chowhan, P.K. (2013). "Supply Chain Planning Framework for Industrial Gas Producer." Available at SSRN 2385790.
13. Derigs, U., Li, B., and Vogel, U. (2010). "Local search-based metaheuristics for the split delivery vehicle routing problem." Journal of the Operational Research Society, 61 (9), pp. 1356–1364.
14. Dror, M. and Trudeau, P. (1989). "Savings by split delivery routing." Transportation Science, pp. 141–145.
15. Ellis, K.P., Sherali, H.D., Saunders, L., Shao, S., and Crawford, C. (2014). "Bulk tank allocation to improve distribution planning for the industrial gas industry." IIE Transactions, 46 (6), pp. 557–566.
16. Feillet, D., Dejax, P., Gendreau, M., and Gueguen, C. (2004). "An exact algorithm for the elementary shortest path problem with resource constraints: Application to some vehicle routing problems." Networks, 44 (3), pp. 216–229.
17. Gillett, B.E. and Miller, L.R. (1974). "A heuristic algorithm for the vehicle-dispatch problem." Operations research, 22 (2), pp. 340–349.
18. Industry Experts (2014). "Industrial, Medical and Specialty Gases - A Global Market Overview." Technical report, Market Publisher.

19. Jin, M., Liu, K., and Bowden, R.O. (2007). "A two-stage algorithm with valid inequalities for the split delivery vehicle routing problem." International Journal of Production Economics, 105 (1), pp. 228–242.
20. Kleywegt, A.J., Nori, V.S., and Savelsbergh, M.W. (2002). "The stochastic inventory routing problem with direct deliveries." Transportation Science, 36 (1), pp. 94–118.
21. Kleywegt, A.J., Nori, V.S., and Savelsbergh, M.W. (2004). "Dynamic programming approximations for a stochastic inventory routing problem." Transportation Science, 38 (1), pp. 42–70.
22. Lee, C.G., Epelman, M.A., White III, C.C., and Bozer, Y.A. (2006). "A shortest path approach to the multiple-vehicle routing problem with split pick-ups." Transportation research part B: Methodological, 40 (4), pp. 265–284.
23. Moin, N.H. and Salhi, S. (2006). "Inventory routing problems: a logistical overview." Journal of the Operational Research Society, 58 (9), pp. 1185–1194.
24. Ryan, D.M., Hjorring, C., and Glover, F. (1993). "Extensions of the petal method for vehicle routeing." Journal of the Operational Research Society, pp. 289–296.
25. You, F., Pinto, J.M., Cap´on, E., Grossmann, I.E., Arora, N., and Megan, L. (2011a). "Optimal distribution-inventory planning of industrial gases. I. Fast computational strategies for large-scale problems." Industrial & Engineering Chemistry Research, 50 (5), pp. 2910–2927.
26. You, F., Pinto, J.M., Grossmann, I.E., and Megan, L. (2011b). "Optimal distribution-inventory planning of industrial gases. II. MINLP models and algorithms for stochastic cases." Industrial & Engineering Chemistry Research, 50 (5), pp. 2928–2945.

APPENDIX A

The routes used by the model in the gas distribution system are divided into two categories: single- customer routes and multiple-customer routes. On a single customer route, gases are delivered on direct routes starting at a depot, visiting a source, then replenishing a customer, and then returning to the depot. The costs and durations of these routes are easily estimated using the locations of depots, sources, and customers. For multiple customer routes where more than one customer is visited on a route, the costs and durations are more difficult to obtain, because both the customer clusters, and the sequence that the customers are visited affect the routes costs and durations. A customer's demand may be satisfied by a combination of single customer or multiple customer routes over the time horizon. Therefore, we need to generate a sufficient number of each type to input to the model. We initially normalize all demands δ_i using full load drops (fld_i) and partial load drops (pld_i) where:

$$fld_i = \left\lfloor \frac{\delta_i}{g} \right\rfloor \tag{22}$$

$$pld_i = \frac{\delta_i}{g} - fld_i \tag{23}$$

The full load drops (fld_i) represent the part of the customer demand that could potentially be delivered with single customer routes (depending on the tank capacity). For the sweep heuristic, the partial load drops (pld_i) are used to generate routes on which multiple customer are visited.

Single-Customer Routes
Single-customer routes are direct delivery routes to one customer site and back to the depot (single customer on each route). To ensure feasibility, a sufficient number of single customer routes are input to the model. The number of single-customer routes needed for each customer is obtained by taking the maximum of the number of full load drops (fld_i) or the customer demand divided by the allocated tank capacity minus the safety stock requirement for that customer using the following expression:

$$\max\left\{fld_i, \left\lceil \frac{\delta_i}{a_{it}\rho_i v_t - \sigma_i} \right\rceil\right\} \tag{24}$$

Multiple-Customer Routes
Multiple-customer routes are routes on which more than one customer is visited. To generate multiple-customer routes, first, clusters of customers, depots, and gas sources are defined and then, a sweep-based heuristics is used to generate potential delivery routes within each cluster.

Before routes are constructed, customers are grouped based on preassigned depots and sources, referred to as a cluster. The sweep-based heuristic generates individual routes using the depot, source, and set of customers from a single cluster. On each route, a trailer leaves the cluster depot, obtains product from the cluster source, then delivers product to cluster customers. The trailers may then either return to the depot and conclude the route or may obtain additional product during delivery by refilling at a source. When trailers visit multiple sources during routes, this is referred to as a *continuous route*. The inclusion of continuous movement adds additional decisions to the route construction process: when to refill, where to refill, and which customers to visit following a refill.

A modified version of the sweep algorithm has been adapted for continuous movement and addresses the decision of when to visit a source to refill, which source to visit, and which customers to serve following a refill. The algorithm begins by representing each customer i in polar coordinates (ρ_i, θ_i) with the source as the pole or origin and a polar axis (1, 0). Let N be the number of customers assigned to this source and R be the set of routes generated. The customers are then sorted in non-decreasing order of θ_i, to form set Ip,γ, with the first customer indexed as 1 and the last as N. In the following algorithm, the route construction heuristic starts from one source. Each customer has the opportunity to lead the sweep procedure, thus, we repeat the algorithm N times (iterations) for every sorted cluster, Ip,γ, each time starting from a unique customer. We define two bookkeeping variables, i_{start} to track the starting customer ($i_{start} \in Ip,\gamma$) for that iteration, and i_{next} to identify the next sorted customer ($i_{next} \in Ip,\gamma$) to be visited.

Additional Notation

Γ	set of sources indexed by γ.
Γp	subset of sources assigned to serve customers in Ip.
Ip,γ	sorted set (cluster) of customers assigned to depot p and source γ.
I'	sorted set (cluster) of customers under consideration.
maximum-source-visits	limit on the number of source visits per route.
maximum-trip-size	maximum number of customers added to a trip between source visits.
trailer-refill-threshold	amount of product left on trailer before refill is allowed.

Route Construction Method

Initialization Set $I' \leftarrow Ip\gamma$, $N = |I'|$, $istart = 1$, and $R = \emptyset$.

Step 1 (Route Start) Set $i_{next} \leftarrow i_{start} + 1$. Create a new route r, which begins at the cluster depot, visits the cluster source, and start to consider customers in the clockwise direction from the i_{start} position on the ordered list.

Step 2 (Sweep) If the number of customers served on r has reached *maximum-trip-size*, go to Step 3 to consider a refill. Otherwise, consider adding customer i_{next} to the current route. If the remaining capacity cannot serve customer i_{next} but is above the *trailer-refill-threshold*, set $i_{next} \leftarrow i_{next} + 1$ and repeat Step 2 to consider the next customer. If the trailer capacity can serve i_{next} then add i_{next} to r, set $i_{next} \leftarrow i_{next} + 1$, and repeat Step 2. If the remaining trailer capacity cannot serve i_{next} and is below *trailer-refill-threshold*, go to Step 3 to refill.

Step 3 (Consider Refilling) If *maximum-source-visits* has been reached, go to Step 4 and end the route. Otherwise, search the set of sources serving customers from this depot, Γp, for the source nearest the last customer served. If the nearest source is the original cluster source, visit to refill on product and go to Step 2 to continue serving customers. Otherwise, visit the nearest source, γ', to refill on product, replace the current cluster customers with customer assigned to that source (*hopping*). Set the current cluster to the sorted customers of the new cluster, $I' \leftarrow Ip\gamma'$, $N = |I'|$, $inext = 1$, and go to Step 2.

Step 4 (End Route) Add the route to the set of completed routes $R \leftarrow R \cup r$. Set $I' \leftarrow Ip\gamma$ and $N = |I'|$. If $istart < N$, then set $istart \leftarrow istart + 1$ and go to Step 1.

The algorithm is applied to each source and the associated customers. In the procedure shown above, customers are considered in the clockwise direction. The heuristic is also run in the counterclockwise direction to create additional candidate routes.

During the process of route construction, it is possible for a trailer to obtain product from multiple sources on a single route. After a trailer obtains product from a source, it begins a polar sweep to add additional customers to the trip. Eventually, the trailer may consider refilling on product. The algorithm then searches the set of sources allocated to serve customers assigned to the current depot for the source located nearest the trailers current location. If this source is different from the last one, the trailer refills on product and proceeds to serve customers assigned to this new source. This process is denoted as *hopping* because the trailer is moving to serve a new set of customers. When a route hops from one source to another, a new set of customers must be considered. After refilling on product at a new source, the next customer to visit is determined for that source. The same route construction method may be used to select the next customer. In our implementation, we assume that, after hopping, the trailer is most likely approaching maximum allowable route time. Therefore, we adjust the orientation of polar coordinates such that customers are selected in the direction of the depot.