




INFORMS Journal on Applied Analytics

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Saving Time and Money and Reducing Carbon Dioxide Emissions by Efficiently Allocating Customers

Christian Haket, Bo van der Rhee, Jacques de Swart

To cite this article:

Christian Haket, Bo van der Rhee, Jacques de Swart (2020) Saving Time and Money and Reducing Carbon Dioxide Emissions by Efficiently Allocating Customers. INFORMS Journal on Applied Analytics 50(3):153-165. <https://doi.org/10.1287/infe.2020.1028>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2020, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Saving Time and Money and Reducing Carbon Dioxide Emissions by Efficiently Allocating Customers

Christian Haket,^a Bo van der Rhee,^b Jacques de Swart^c

^a Anders Invest, Maarsbergen 3953 BA, Netherlands; ^b Center for Marketing and Supply Chain Management, Nyenrode Business University, 3621 BG Breukelen, Netherlands; ^c Center for Accounting, Auditing and Control, Nyenrode Business University, 3621 BG Breukelen, Netherlands

Contact: christian_haket@hotmail.com (CH); b.vdrhee@nyenrode.nl,  <https://orcid.org/0000-0003-4217-7966> (BvdR); j.dswart@nyenrode.nl (JdS)

Received: September 26, 2018

Revised: April 15, 2019; September 7, 2019; November 16, 2019

Accepted: December 13, 2019

Published Online in Articles in Advance: April 29, 2020

<https://doi.org/10.1287/inte.2020.1028>

Copyright: © 2020 INFORMS

Abstract. In many industries, multifacility service providers can save time and money and reduce carbon dioxide (CO₂) emissions by more efficiently allocating customers to their facilities. However, firms incur a reallocation cost when reassigning a customer to a different facility, something that has not received much attention in the literature. Software packages such as CPLEX can find the optimal solution for this type of problem, but managers rarely use them because they lack the specific knowledge, overestimate the cost, and/or underestimate the benefits. Including the reallocation costs, we modeled several common heuristics in Excel's Visual Basic and compared the results with the optimal solution found by CPLEX. We collaborated with Van Dorp, a large service provider in the Netherlands, and found that (1) substantial savings can be achieved, (2) reallocation costs play a major role, and (3) the best heuristic achieves near-optimal results. Specifically, reallocating Van Dorp's 20 "worst allocated" customers realizes a savings of more than 4,000 driving hours and €360,000 in cost and a reduction of 41 tons of CO₂ emissions.

History: This paper was refereed.

Keywords: heuristics • assignment problem • reallocation cost

Introduction

For this paper, we intensively worked together with Van Dorp, one of the largest service providers in the technical installation industry in the Netherlands. Van Dorp specializes in service and maintenance of technical installations in utility buildings. The company provides nationwide coverage to 4,888 customers from 14 regional facilities, with approximately 1,000 employees, of whom approximately 300 work in the main office and 700 are on-site service providers. Typically, a service provider visits one customer per day and must drive to this customer from the facility at which he or she is stationed, even if another facility is closer. The reason for this apparent suboptimality is that separate facilities are largely independent, and the directors of the facilities are responsible for their own profit and loss. The facility directors are hired as, and told to be, local entrepreneurs. As such, central coordination on customer allocation interferes with local independence and is therefore rarely discussed. Occasionally, customer allocation inefficiencies are visible enough to provoke awareness and induce actions by general management, but these reallocation decisions are based on qualitative arguments without an in-depth knowledge of the quantitative effects of a certain (re)allocation decision.

To illustrate the potential savings, the following calculation applies to Van Dorp: in 2016, the total distance traveled by the company's service providers was approximately 6 million kilometers (approximately 3.75 million miles). A 1% improvement in efficiency already results in a reduction of 60,000 kilometers per year. With an average speed of 60 km/h, this means saving 1,000 productive hours, which translates into approximately €90,000 by considering the savings in cost of labor, vehicle depreciation, maintenance, and fuel (i.e., collectively approximately €90 per hour), but ignoring the additional reallocation costs. The reduction in carbon dioxide (CO₂) emissions, using an average theoretical value of 125 g/km (Kadijk et al. 2015), would be approximately 7.5 tons annually.

In general, service providers who must visit their customers are increasingly interested in reducing their employees' time and distance spent on the road (Buijs et al. 2016), making it remarkable that most service providers pay little attention to operations research. For example, our recent discussions with six directors of large (>1,000 employees) service providers revealed that for multifacility service providers, the allocation of customers among different facilities is known to be particularly inefficient, but little to no effort is made to address this. In fact, when

we started working with Van Dorp, the company acknowledged that the problem existed but was reluctant to work on it because, according to an unidentified manager, “the effort of reallocating would not be worth the benefits.”

Of course, reallocating a customer to another facility incurs some cost, both tangible (i.e., operational and administrative changes in the organization’s systems) and intangible (i.e., idiosyncratic knowledge of the customer). We combine these in a single cost measured in terms of the time required to reallocate a customer, something not investigated in-depth in the literature. Because the intangible costs are hard to estimate, we offset this by only considering a reallocation if the time it takes can be regained by the reduced travel time for that customer within one year. This is a conservative approach because a reallocation incurs a one-time cost but continues to save time and money and reduce CO₂ emissions over the duration of the multiyear contract. In summary, the reallocation costs are a fixed value measured in terms of time, which acts as a threshold for reallocating a specific customer.

The complexity and size of many real-life problems require the construction of a simplified representation of reality in a model. Modeling customer allocation problems has received plenty of attention in the field of facility location research (see, e.g., Brandeau and Chiu 1989, ReVelle and Eiselt 2005, Daskin 2013), and we study a subclass of facility location problems where the locations are fixed, with the focus on the (re)allocation of customers, the so-called assignment problem. As such, rather than focusing on determining where to locate the service provider’s facilities (see, e.g., Allen et al. 2017), we focus on which customers should be assigned to which facilities.

Problems such as this can be solved through the application of CPLEX, and indeed, we have an optimal solution for the problem we investigated, as described in more detail later. However, this is rarely applied because managers do not possess the necessary skills to use such sophisticated and expensive software. Therefore, we modeled several heuristics in Excel to develop insights into Van Dorp’s ability to improve efficiency. We also consider the effect of the facilities’ ability to increase their workload (i.e., *capacity relaxation*) and run sensitivity analyses on the reallocation cost.

Although theoretical research on these kinds of assignment problems abounds, there is a deficit in the application of theory to a real-life case study (ReVelle and Eiselt 2005). The literature review shows that most research is purely theoretical or uses standard databases such as those developed by Beasley (1990). Our aims for this research therefore were threefold: (1) to determine the extent of the savings (time,

money, and reduction of CO₂ emissions) for Van Dorp, (2) to determine the role that reallocation costs plays in the optimal allocation of customers, and (3) to determine the performance of easy-to-implement heuristics relative to the optimal solution in a real-life business setting.

Related Literature

Type of Problem: Assignment Problem

The general class of facility location problems concerns the location of facilities and the allocation of customers to those facilities, and if the facility’s locations are fixed, it becomes an assignment problem. We refer the interested reader to Current et al. (1990) for multiple-objective models, although our single objective is to minimize the sum of the distances (Brandeau and Chiu 1989) between the fixed service-providing facilities and service-demanding customers. We express the distances in time because we also include a reallocation cost, which is measured in hours. Therefore, our objective is to minimize the total service and reallocation time, which impacts profits because the overall cost is reduced, people because they spend less time traveling, and the planet because CO₂ emissions are reduced (Roobeek et al. 2018).

Because more than 70% of Van Dorp’s mechanics only visit one customer per day, we do not focus on routing between customers (see, e.g., Salhi and Rand 1989). Instead, we focus on allocating customers in such a way that the time spent on round trips is minimized. Because the locations of both the customers and the facilities are known and generally quite stable for the duration of one year, this makes our problem discrete and static (ReVelle and Eiselt 2005). Furthermore, although travel time is stochastic by nature (e.g., unpredictable traffic jams), aggregated data can be used to state travel times in term of averages, making this input deterministic (Klose and Drexl 2005).

The type of customers considered in our research are served by one facility, and the capacities of the facilities are considered limited to some extent in the short term, which makes our problem one with capacitated facilities (Tragantalerngsak et al. 2000) that provide a single type of service. Given all of the preceding, our work is a version of the general assignment problem (GAP), studied at least as early as in Barr et al. (1977). However, we need to incorporate the costs associated with reallocating a customer in the model, which, to the best of our knowledge, has not been done before in the literature. For example, in the survey paper by Öncan (2007), 11 extensions to the GAP were discussed, but none of these dealt with reassigning jobs to resources.

Type of Solution: Exact and Heuristic Algorithms

Small-integer linear problems can be solved through brute-force search techniques (Geoffrion and Van

Roy 1979), but larger problems can quickly become intractable for solving by enumeration. The preferred way of solving location problems is by applying exact solution techniques such as the branch-and-bound and cutting-planes algorithms (Beasley 1998, Daskin 2013). Although exact algorithms guarantee that the global optimal solution is found, they may take an exponential number of iterations to find that solution (Mladenović et al. 2007). Given the increase in computer speed, this is not really a computational issue anymore, and we applied CPLEX to our case study to find the optimal solution as well (refer to Appendix A for the implementation). However, as mentioned in the Introduction, managers often do not favor this approach, which is why we also turn to heuristics.

Existing heuristic methods to solve location problems can also be applied to our capacitated assignment problem (Ross and Soland 1977). Different algorithms show different results for different problems, which implies that no single best algorithm has been identified (see, e.g., Reese 2006, Gendreau and Potvin 2010). Heuristic algorithms are often classified as construction, improvement, or metaheuristic methods. Construction algorithms start with an empty initial solution and iteratively add a new element to the solution until a complete solution is found (Daskin 2013). Improvement algorithms start with a nonempty solution and iteratively try to improve the solution until no longer possible, until the last improvement is negligible, or until the maximum number of iterations is reached (Korupolu et al. 2000). Combinations of construction and improvement heuristic algorithms often form the basis for metaheuristic algorithms (Mladenović et al. 2007), such as the greedy randomized adaptive search procedure (GRASP). For more information on metaheuristic algorithms, the reader is referred to Amrani et al. (2011), Gendreau and Potvin (2010), Klibi et al. (2010), and Nagy and Salhi (2007).

Method

Mathematical Model

Let $I = \{1, \dots, i, \dots, 4,888\}$ be the set of demand nodes (i.e., customers of Van Dorp), $J = \{1, \dots, j, \dots, 14\}$ the set of facility locations, each with capacity c_j in terms of the number of customer visits that the facility can facilitate, and D the $4,888 \times 14$ matrix with distances d_{ij} representing the driving time between customer i and facility location j . Furthermore, let h_i be customer i 's demand in terms of the number of required visits per year, T the reallocation cost (in terms of time), a_{ij} the current allocations, and the decision variables X_{ij} the new allocations (which may equal a_{ij} , indicating that

there was no reallocation). We now specify the linear model as follows:

$$\text{Min } \sum_{i \in I} \sum_{j \in J} h_i \cdot d_{ij} \cdot X_{ij} \quad (1)$$

subject to

$$\sum_{j \in J} X_{ij} = 1, \quad \forall i \in I, \quad (2)$$

$$\sum_{i \in I} h_i \cdot X_{ij} \leq c_j, \quad \forall j \in J, \quad (3)$$

$$X_{ij} \in \{0,1\}, \quad \forall i \in I, \forall j \in J, \quad (4)$$

$$\sum_{j \in J} h_i \cdot d_{ij} \cdot (a_{ij} - X_{ij}) \geq R_i \cdot T, \quad \forall i \in I, \quad (5)$$

where

$$R_i = 1 - \sum_j a_{ij} X_{ij} \forall i, \quad \forall i \in I. \quad (6)$$

This formulation is very similar to a GAP (see, e.g., Öncan 2007), where the typical Constraints (2)–(4) ensure that all customers' demand must be served, that a facility's capacity is not exceeded, and that single sourcing is guaranteed. We extend this with Constraint (5) that ensures that a customer is reallocated only when the reallocation savings for that customer exceed the reallocation costs, where (6) defines the binary reallocation variable that indicates whether customer i is reallocated (i.e., $X_{ij} \neq a_{ij}$; thus $R_i = 1$) or not ($X_{ij} = a_{ij}$; thus $R_i = 0$). Although Constraint (5) is enforced for each possible reallocation at the customer level, we acknowledge that removing it would enable a system-wide optimization that could capture trade-offs between reallocation cost and assignment/travel cost. In such an alternative formulation, the solution could increase the cost to serve a particular customer in order to reduce cost for serving others as long as the system-wide cost is minimized. However, we include (5) on Van Dorp's request because the company believes that customers are unwilling to change to a new service provider that must travel further. Alternatively, one could include a large penalty to the objective function for reallocations that would lead to increased cost for an individual customer, but we believe that Constraint (5) is more transparent. We discuss the impact of this additional constraint later.

Heuristic Algorithms

We programmed each algorithm in Visual Basic for Applications (VBA) in Excel, which is not the most elegant programming language but does come with the benefit that managers understand it because no

specialized software is required. As discussed later, there are three basic types of heuristics: construction, improvement, and metaheuristics.

We first developed four variations of construction algorithms that start with an empty solution and then add the customers one by one: (1) the sequential, (2) the random, (3) the savings regret, and (4) the randomized adaptive greedy adding algorithm. Each selects the next customer to assign in a slightly different way, leading to different performances. The sequential version goes through customers in sequential order from 1 to 4,888, adding each to the best facility with available capacity. The advantage of this heuristic is that it is very quick, giving a runtime of less than 90 seconds in Van Dorp's case. However, it gets stuck in a local optimum and therefore is the worst-performing adding algorithm. Because the runtimes of the other heuristics are not orders of magnitude higher (see Table B.1 in Appendix B), we recommend against using the greedy sequential heuristic in this setting.

The random version goes through the customers in random order. This requires a runtime of approximately 320 seconds but has the advantage of not always getting stuck in the same local optimum. In fact, our analysis revealed that this version outperformed the other variations for most of the reasonable capacity relaxation levels. The pseudocode for this version is shown in Figure 1, whereas the others are available from the authors upon request. Figure 1 shows a single loop in which each customer is assigned to the nearest facility with available capacity.

The savings regret version first assigns the customers with the highest demand to available facilities, similar to packing a car with the larger bags first and the smaller ones last. However, whereas this version's runtimes are also less than 90 seconds and it performs better than the sequential version, it also gets stuck in a local optimum. Therefore, we do not recommend this heuristic in this setting.

Figure 1. Pseudocode of the Randomized Greedy Adding Algorithm

```

Sub GreedyAddingRandom ()
  1 initiate facilities  $k$ , capacities, demands, costs, empty solution and seed
  2 calculate weighted costs including reallocation cost
  3 generate quasi-random distribution of customers  $c$ 
  4 for  $c = 1$  to  $i$  do
  5   find customer with value  $c$ 
  6   for  $j = 1$  to  $k$  from lowest to highest cost of allocation do
  7     if facility  $j$  has enough spare capacity then
  8       allocate customer  $c$  to facility  $j$ 
  9   else goto next facility
  10  update solution
  11 next  $c$ 
  12 return solution
End sub GreedyAddingSequential

```

Finally, in the randomized adaptive version, the heuristic first creates a shortlist of a random number of the largest remaining customers and then randomly selects a customer from this list to be assigned to a facility. This version performs in-between the random and the savings regret versions in terms of setup, runtime (~170 seconds), and performance. In conclusion, although the random version takes the longest, the runtime is still reasonable for Van Dorp's problem size at slightly more than five minutes, and it performs best. Hence, we recommend the random greedy adding heuristic as the best greedy adding heuristic for a problem with Van Dorp's characteristics.

We then developed three versions of the improvement heuristics that each start with the current allocation at Van Dorp: (1) the "first" improvement, (2) the random "best" improvement, and (3) the random "one-opt" local search algorithm. Each version tries to iteratively improve on the current allocation in different ways. The first version reallocates a customer to the first facility it finds that can realize an improvement. This means that it goes through the list of customers several times because it is possible to reallocate a customer multiple times, each time slightly improving the solution. The runtime for all improvement heuristics is substantially higher than for the greedy adding heuristics, with the first version clocking in at around 15 minutes (refer to Table B.2 in Appendix B). In terms of the runtime, the first version is the second fastest of the improvement heuristics and requires the most iterations for reasons explained earlier. However, it is also the improvement heuristic that finds the lowest value of the objective function, which was surprising, as we discuss in more detail later.

The best version *only* reallocates a customer when the *best* facility for that customer has available capacity. This means that this heuristic does *not* reallocate a customer to a closer facility that has available capacity if there is an even closer facility without available capacity. The advantage of this heuristic is that it is the fastest of the improvement heuristics (runtime of approximately seven minutes) with the fewest iterations. Also, following this heuristic, no customer would say, "If you reallocate me anyway, then why not to the closest facility?" However, the disadvantage is that this heuristic also gets stuck in a local optimum and "refuses" to reallocate customers to better facilities, even if those facilities have available capacity. Because the differences are relatively small (<1% worse compared with the first version) and the runtimes are about half, we would recommend this heuristic in situations in which it is important to *only* reallocate customers if the closest facility has available capacity. However, because that

Table 1. Solution Quality of the Three Heuristic Algorithms

Capacity relaxation	Greedy random			First improvement			GRASP		
	Value in minutes	Number of reallocations	Gain over base case (%)	Value in minutes	Number of reallocations	Gain over base case (%)	Value in minutes	Number of reallocations	Gain over base case (%)
+0%	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
+1%	4,583,129	392	7%	4,448,055	416	10%	4,666,931	379	5%
+2%	4,529,284	399	8%	4,407,311	419	11%	4,599,924	387	7%
+3%	4,493,259	404	9%	4,388,854	422	11%	4,538,091	395	8%
+4%	4,473,956	410	9%	4,366,015	424	12%	4,485,025	407	9%
+5%	4,433,059	417	10%	4,348,904	429	12%	4,449,587	415	10%
+10%	4,325,658	432	12%	4,268,466	468	14%	4,335,614	436	12%
+15%	4,247,001	471	14%	4,215,021	499	15%	4,222,682	489	14%
+20%	4,208,792	505	15%	4,198,832	510	15%	4,178,853	511	15%
+25%	4,180,935	514	15%	4,172,590	515	15%	4,172,590	515	15%
+50%	4,170,608	518	15%	4,170,608	518	15%	4,170,608	518	15%
Uncapacitated	4,170,608	518	15%	4,170,608	518	15%	4,170,608	518	15%

Notes. Value in minutes includes reallocation costs. GRASP, greedy randomized adaptive search procedure; N/A, not applicable.

was not a constraint for Van Dorp, we used the first version because its solutions were slightly better.

Finally, the one-opt version is very similar to the best version but *will* assign a customer to a better facility even if it is not the best. This makes the performance a bit worse than the best version because facilities “fill up” with suboptimal reallocations and can no longer take on optimal reallocations. More important, this heuristic works differently by assigning a customer to all facilities with enough available capacity and then compares the values of the objective function to determine where to (re)allocate the customer. This is done to also allow for a two-opt version (or higher-order opts) in which two (or more) customers are swapped between facilities. However, this different setup is very time consuming, with runtimes well above an hour. Although still reasonable for a problem of this size, especially because Van Dorp would not run the model

daily, the fact that it does not outperform the faster and more easily understood first version leads us to recommend against using the one-opt version in this setting.

Figure 2 shows the pseudocode of the first improvement local search algorithm, in which there are two loops, one for the customers where the minimum allocation cost is identified and one where the customer is assigned to the facility with the lowest possible cost with still-available capacity. The improvement heuristics are generally randomized because sequential local search methods are quickly mired in a local optimum, especially when facilities operate close to capacity. A local search algorithm also requires a parameter that states when the algorithm should stop such that it stops after the last improvement was smaller than x . In Van Dorp’s case, we simply used the last iteration that did not improve the solution (i.e., $x = 0$).

Table 2. Solution Quality of the “First” Improvement Heuristic Compared with CPLEX

Capacity relaxation	CPLEX					First improvement
	Total travel cost	Number of reallocations	First-year total cost	Gain over base case (%)	Second-year gain (%)	Loss over CPLEX (%)
+0%	4,241,968	421	4,393,528	11%	14%	N/A
+1%	4,206,375	425	4,359,375	12%	15%	2.0%
+2%	4,181,637	434	4,337,877	12%	15%	1.6%
+3%	4,163,097	430	4,317,897	13%	16%	1.6%
+4%	4,145,039	443	4,304,519	13%	16%	1.4%
+5%	4,127,377	456	4,291,537	13%	16%	1.3%
+10%	4,057,457	480	4,230,257	14%	18%	0.9%
+15%	4,016,334	509	4,199,574	15%	19%	0.4%
+20%	3,996,588	513	4,181,268	15%	19%	0.4%
+25%	3,986,473	516	4,172,233	15%	19%	0.0%
+50%	3,984,128	518	4,170,608	15%	19%	0.0%
Uncapacitated	3,984,128	518	4,170,608	15%	19%	0.0%

Note. The second-year % gain (over the base case) is excluding the reallocation cost because that is a one-time investment.

Figure 2. Pseudocode of the Random First Improvement Local Search Algorithm

```

Sub LocalSearchFirstIterativeImprovementRandom
  1 initiate capacities, demands, costs, reallocation cost, stopping criterion
    and non-empty solution
  2 calculate weighted costs including reallocation cost
  3 while local search stopping criterion is not met do
  4   randomly select allocated customer
  5   set cost as current allocation cost
  6   if current allocation cost > minimum allocation cost then
  7     for facility 1 to  $j$  do
  8       if cost < current allocation cost then
  9         if facility  $j$  has enough spare capacity then
 10           allocate customer to facility  $j$ 
 11           goto nextcustomer
 12         else: end if
 13       else: end if
 14     next facility
 15   else: end if
 16 nextcustomer:
 17 loop
 18 return solution
End sub

```

Finally, the GRASP consists of two phases. The construction phase starts with an empty solution and quickly constructs a feasible solution using the randomized adaptive greedy algorithm from which it takes its name. This solution is then optimized in the second phase using the first improvement local search algorithm. We refrain from showing the pseudocode here and refer to Table B.3 in Appendix B for the performance and runtimes.

CPLEX Solution

To compare the heuristic algorithms with CPLEX, we programed the model specified by Equations (1)–(6) as a mixed-integer programming (MIP) mathematical program within AIMMS (Roelofs and Bisschop 2017). Please refer to Appendix A for implementation details.

Measurement and Analysis Procedure

The current allocation at Van Dorp resulted in a total travel time of 4,935,181 minutes in 2016, and we determined the algorithms' performances by comparing their solutions to this base case. Randomized algorithms may return different solutions per run, which is why we ran each heuristic 10 times and report the (local) minimum. Also, we ran the heuristics for different facility capacity relaxation levels, as shown in the results section in Tables 1 and 2. Finally, we translated the improvement into monetary savings and reduction in CO₂ emissions.

Data

We retrieved the input data of Van Dorp customers as of July 1, 2017, from the enterprise resource planning

system. For those who were also customers during the 2016 calendar year (91%), we used the number of service visits of 2016 to represent their total demand for 2017. Year-over-year changes in demand are less than 5%, so this is not a strong assumption. For customers who joined January 1, 2017 (5%), we doubled the number of visits in the first half of 2017 to estimate the total number of required visits in 2017. Finally, we estimated the total number of visits to short-term customers (4%) based on similar contracts.

We defined the capacities of the different facilities as the current number of visits to customers without taking routing into consideration. In reality, capacity is stated in terms of the number of available hours to work on-site rather than the number of visits that can be made. However, when aggregating total demand over the course of a year for all facilities, the total number of visits to a customer represents the facility capacity. However, in collaboration with Van Dorp's management, we assumed that up to 10% additional capacity over the current capacity is feasible for all facilities. For all heuristics, we compare the results for different capacity relaxations to determine the influence of this assumption on the outcome.

The cost matrix required for the model is a $4,888 \times 14$ matrix with round-trip travel times between the facilities and the customers multiplied by the total number of visits to each customer. We sent the empty distance matrix to an external company that calculated all the distances between the facilities and the customers in both kilometers and travel time. The company based these on the real travel time including average traffic delays on a typical Tuesday morning, which is representative of the mechanics' average travel times. For our heuristics, we used travel time, with the corresponding distances in kilometers used to calculate the reduction in CO₂ emissions. We then started with the current allocation of customers (a_{ij}) in the same matrix, with zero representing that a customer is not serviced by a facility and one that it is. As a side note, the algorithms' runtimes were reduced by more than 75% by leaving a cell blank rather than explicitly adding a zero.

Finally, we modeled the reallocation costs as an exogenous parameter, determined through internal research at Van Dorp. For most customers' buildings, a similar amount of time is required to make the proper administrative adjustments in the financial and operational systems. Introducing the customer to the new point of contact also generally requires a similar amount of time. Intangible costs, such as idiosyncratic knowledge of customers' assets, are very hard to determine exactly and differ widely between facilities and employees. Therefore, we omitted these costs from the scope but posit that these are at least partially and potentially more than completely offset

by the one-year horizon of our model. Van Dorp estimated that two hours of administrative changes and four hours for the introduction of a new service provider are required, for six hours (360 minutes) in total. A possible extension of our work could be to make the reallocation cost depend on the customer's demand.

Results

Performances of the Heuristic Algorithms

Interestingly, none of the heuristic algorithms could find an improvement on the current allocations for a 0% relaxation of the facility capacities, even though CPLEX could. Table 1 shows the solution in terms of total time in minutes, which is the combination of the travel time (i.e., the objective function) and the reallocation time (i.e., the number of reallocations multiplied by the reallocation time). It then shows the number of required reallocations to achieve this solution, as well as a relative gain over the base case. Please note that this is a comparison for the first year only because the reallocation continues to benefit Van Dorp, whereas the reallocation costs are only incurred once. This is made more explicit in Table 2, which shows the second-year gains as well.

Contrary to our expectations based on the literature, the GRASP algorithm was not always able to find a better solution than the construction phase alone. We tried the other three construction algorithms as well, but the sequential and savings regret greedy algorithms did not result in improvements in the local search phase, and the randomized greedy algorithm gave some variations in the outcome but no more than the adaptive version of the algorithm. Another possible option is to use a more sophisticated local search algorithm in future research.

The fact that the GRASP algorithm is sometimes unable to improve the intermediate solution after the construction phase is known and addressed in the literature (see, e.g., Feo and Resende 1995). Our goal of including the GRASP algorithm was to determine whether it would be as beneficial in a real-life business setting as touted in the literature rather than a methodological goal of developing a better metaheuristic. We found that the theoretical benefits cannot be achieved in Van Dorp's case, so we now focus on the first improvement heuristic to compare its performance with the optimal solution.

Comparison with CPLEX


Because the first improvement's heuristic achieves the best results up to and including a 15% capacity relaxation (refer to Table B.4 in Appendix B), we compare that heuristic with the optimal solution found by CPLEX. Table 2 first shows detailed solutions found by CPLEX: the value of the objective

function in terms of total travel cost, followed by the number of reallocations for the different capacity relaxation levels. These two add up to the first-year total cost, after which we show the gain compared with the base case in the first year as well as the second year, where reallocation costs are no longer incurred. Please note that in the noncapacitated case (listed as "Uncapacitated" in Tables 1 and 2), every customer is simply assigned to its closest facility, where *closest* is defined as the travel time plus the possible reallocation time. Of the 1,300 customers who are currently not allocated to their nearest facility, 782 of them have reallocation gains lower than the reallocation costs and are therefore not reallocated even in the noncapacitated case.

The last column in Table 2 shows the first improvement local search heuristic's performance compared with the optimal solution. As the capacity relaxation increases, so does the performance of the heuristic, which finds the optimal solution at a capacity relaxation of 50%. The focus for our case study was on a capacity relaxation of 10%, where the heuristic result is only 0.9% worse than optimal. In this solution, the first improvement local search heuristic would reallocate 468 customers.

We also let CPLEX determine the globally optimal solution at different capacity relaxations without including Constraint (5). This resulted in a further cost reduction in the range of 0%–1.50%, with higher additional gains for lower-capacity relaxations. However, these solutions require many customer reallocations to facilities that are farther away: 31 at a 0% capacity relaxation, and this number increases for higher-capacity relaxations. At the 10% capacity relaxation, the additional cost reduction is only 0.54%, and because Van Dorp believes that it is impossible to convince a customer to be served by a facility that is farther away, we recommend keeping Constraint (5) in the model.

Finally, because reallocating 480 customers is not a simple task in practice, Van Dorp also requested that we determine the 20 most inefficiently allocated customers with relatively high demand—that is, the "worst allocated" customers. Of course, any specific number may be chosen, granted that no feasibility problems occur, such as reallocating customers to a location without enough available capacity. An elegant way to avoid this is to increase the reallocation time T such that only a certain number of reallocations remain (refer also to Figure 3). In our case, the 20 worst-allocated customers were the same for both CPLEX and the first improvement local search heuristic, and they could be reallocated without any feasibility issues. The reallocation of these worst-allocated customers resulted in total savings of more than €360,000 with more than 4,000 travel hours saved, leading to a 41-ton reduction in CO₂ emissions. Note that the

Figure 3. Number of Customer Reallocations and Total Savings for Different Reallocation Costs

average speed here is 80 km/h, rather than the 60 km/h from the Introduction, because these 20 customers are currently allocated to facilities for which the majority of travel is highway driving.

Takeaways and Sensitivity Analysis

We now explore the similarities and discrepancies between the expectations based on the literature review and our findings, and we finish with a sensitivity analysis with respect to the reallocation cost.

Algorithms' Takeaways

Some of our findings conflict with the expectations based on the literature. For example, Feo and Resende (1995) stated that in situations in which the costs of allocation vary widely between customers, it can be beneficial to apply a savings regret rule. However, the randomized greedy algorithm, which does not apply the savings regret rule, was the best-performing construction algorithm, except for high-capacity relaxations of 20% or more. Therefore, prioritizing customers with the savings regret rule is only beneficial if there are hardly any capacity constraints, which was not the case for Van Dorp.

The improvement algorithms generally returned better solutions. This indicates that starting with the current allocation situation and improving this situation iteratively is preferable over starting with an empty solution. However, that the first version gave the best results was unexpected. Further analysis of the algorithms' steps taken revealed that the best version made allocation decisions in an earlier stage, which appeared to be detrimental to the outcome of the model. The best version returned nearly the same results as the first version and comes with the additional benefit that a customer would only be reallocated to the nearest facility. If this is important, then the best version should be adopted rather than the first version.

The GRASP metaheuristic algorithm seemed suited to our problem because Feo and Resende (1995) stated that it is relatively easy to implement and has a proven track record on assignment problems. Although it was indeed relatively easy to implement after we programmed the separate parts of the algorithm, the results show no real advantage over the other algorithms. It even returned lower-quality solutions than the first improvement heuristic because the improvement phase was unable to improve on the construction phase.

Sensitivity Analysis

In addition to determining the impact of the capacity relaxation as shown in Tables 1 and 2, we ran a sensitivity analysis on the impact of the reallocation costs, currently estimated to be six hours (360 minutes) per customer reallocation. We varied this reallocation cost using the first improvement heuristic algorithm at a capacity relaxation of 10%. Clearly, the savings compared with the current situation are highest when the reallocation costs are the lowest. For example, at no reallocation cost, the potential savings in Van Dorp's case are more than €1.4 million; at 6 hours, they are more than €1 million; and at 30 hours, they are just more than €0.6 million.

Figure 3 shows that the number of customer reallocations quickly decreases for increasing reallocation costs. The model suggests approximately 1,250 reallocations when reallocation costs are not taken into consideration, which decreases to fewer than 200 reallocations when these costs are 30 hours. Because all our data are deterministic, these 200 reallocations show up in all solutions, whereas additional customers are reallocated as the reallocation cost decreases. Figure 3 also can be interpreted by looking at the number of customers that offer potential savings above the reallocation threshold. For example, Figure 3 shows that there are approximately 280 customers

that offer savings greater than 1,000 minutes, adding up to total savings of more than €0.8 million.

Conclusions, Limitations, and Implementation Suggestions

Recall that our aim was threefold: (1) to determine the extent of the savings (time, money, and reduction of CO₂ emissions) for Van Dorp, (2) to determine the role that reallocation costs play in the optimal allocation of customers, and (3) to determine the performance of easy-to-implement heuristics relative to the optimal solution in a real-life business setting. Our approach shows that the potential savings in Van Dorp's case are substantial and also leads to a considerable reduction in CO₂ emissions. The reallocation costs proved to be a very important parameter in the model because the total number of suggested reallocations and the total savings are quite sensitive to the reallocation costs. Finally, although the heuristics were only able to find the optimal solution as determined by CPLEX for higher-capacity relaxations, they perform extremely well on intermediate relaxations and are perfectly suitable to identifying the quick wins.

The complexity and size of real-life problems require the construction of a simplified representation of reality in a model, which presents some limitations. For example, routing affects the total time and distance driven by service providers. Therefore, routing is another factor that could influence the allocation of customers. We did not take routing into consideration because more than 70% of the mechanics' trips are roundtrips to a single customer, but this might not be the case for all multifacility service providers. The nature of our model rules out a sensitivity analysis on this assumption, and we therefore leave this interesting avenue for future research.

We also assumed that we were dealing with a fixed number of facilities, although mergers and acquisitions or closures of facilities could occur. It could even be the case that opening a new facility in a location close to many customers would reap additional savings, but we did not investigate this option. We also did not capture the nature of the relationship between a customer and a facility. A customer could potentially refuse to be reallocated, which is a major challenge for the practical implementation of the suggested reallocations. However, including Constraint (5) at least ensures that customers will only be reallocated if the new service provider is closer.

Another limitation is that the reallocation costs in this study are assumed to be tangible and independent of the size of the customer. Because we have shown that the model is sensitive to the reallocation cost parameter, the costs associated with the intangible knowledge of a customer's assets could be substantial.

Therefore, another extension of the model is to make these intangible costs depend on the size and complexity of a customer's assets.

Finally, our first implementation suggestion is to start with the quick-win reallocations and develop a measurement procedure to check whether the theoretical savings are achieved. If possible, it is also beneficial to pair quick-win reallocations of two facilities (i.e., one customer from *A* to *B* and one from *B* to *A*) to dampen the changes in required capacities. Coordination of customer reallocations over multiple facilities is inherently a centralized decision-making process, but coordinated quick-win reallocations should also bring local managers on board. Finally, to convince the customers, we recommend focusing on the CO₂ reduction rather than on the financial impact for Van Dorp.

This is exactly what Van Dorp did: the company started with reallocations of the worst-allocated customers after discussing the model and results with facility managers. More than 20 reallocations between facilities have so far shown that customers accept the reallocation provided that the reasons are clearly explained and the emphasis is placed on the reduction in CO₂ emissions. The long-term impact of these reallocations is not available yet, but anecdotal evidence gives reason for optimism. Consider, for example, the reallocation of a university to a nearer facility. The quarterly performances review shows that Van Dorp's financial performance has improved since the reallocation, and the customer has also reacted positively.

Acknowledgments

The authors are grateful for the detailed feedback provided by the two reviewers and in particular the associate editor, who allowed the authors to improve the paper in numerous ways. In addition, the authors acknowledge the content editor, who improved the exposition of the paper significantly. Finally, the authors thank Van Dorp for its willingness to cooperate in this research.

Appendix A. CPLEX Implementation

AIMMS can be seen as a shell through which a MIP problem can be offered to CPLEX. The CPLEX version we used was 12.5 (IBM 2017). In the AIMMS code of this implementation, there are five input parameters, one calculated parameter, one mathematical program, one variable determined by the mathematical program, two calculated variables, and three constraints. The five input parameters are (more on the capacity relaxation level *r* in line 2 later)

1. Demand(*c*), the demand of customer *c*, an index running through the set of Customers;
2. Supply(*r*,*f*), the supply at capacity relaxation level *r* for facility *f*, an index running through the set of Facilities;
3. Distance(*c*,*f*), the distance between customer *c* and facility *f*;

4. $\text{Current}(c,f)$, a binary matrix in which one indicates that customer c is served by facility f ; and

5. T , representing the scalar reallocation cost and therefore also threshold.

The calculated parameter is CurrentCosts , representing the costs induced by the current allocation, defined as $\text{CurrentCosts} = \sum((c,f), \text{Demand}(c) \times \text{Distance}(c,f) \times \text{Current}(c,f))$. Finally, the variable is $\text{New}(c,f)$, a binary matrix, in which one means that customer c should be served by facility f .

The first calculated variable is the objective function $\text{NewCosts} = \sum((c,f), \text{Demand}(c) \times \text{Distance}(c,f) \times \text{New}(c,f))$. The second calculated variable is $\text{Reallocate}(c)$, which contains a one when customer c should be served by a facility other than the facility that is currently serving it, and is defined as $\text{Reallocate}(c) = 1 - \sum(f, \text{Current}(c,f) \times \text{New}(c,f))$.

The constraints are modeled in AIMMS as follows (i.e., Constraints (2) and (4) are covered together in $\text{SingleSourcing}(c)$, because we already defined $\text{New}(c,f)$ as a binary matrix):

1. $\text{SingleSourcing}(c)$: $\sum(f, \text{New}(c,f)) = 1$;
2. $\text{Capacity}(f)$: $\sum(c, \text{Demand}(c) \times \text{New}(c,f)) \leq \text{Supply}(f)$; and
3. $\text{Threshold}(c)$: $\text{Demand}(c) \times (\sum(f, \text{Distance}(c,f) \times (\text{Current}(c,f) - \text{New}(c,f))) \geq \text{Reallocate}(c) \times T$.

Finally, the mathematical program is simply defined as a MIP, with NewCosts as the minimizing objective and including all aforementioned variables and constraints. As such, this is a straightforward implementation of the standard GAP, although the definition of $\text{Reallocate}(c)$ and the third constraint may look less trivial. As mentioned earlier in the mathematical model, the $\text{Reallocate}(c)$ constraint ensures that a customer is only reallocated if the gain of less time on the road is larger than the reallocation cost. Implementation of the alternative model without the constraint on the individual customers follows:

Determine the binary variable $\text{NewAllocation}(c,f)$ to minimize

$\text{FirstYearTravelAndAllocationCosts}$

subject to

- $\text{FirstYearTravelAndAllocationCosts} := \text{NewTravelCosts} + \text{ReallocationCosts}$
- $\text{NewTravelCosts} := \sum((c,f), \text{Demand}(c) \times \text{Distance}(c,f) \times \text{NewAllocation}(c,f))$
- $\text{Reallocated}(c) := 1 - \sum(f, \text{CurrentAllocation}(c,f) \times \text{NewAllocation}(c,f))$
- $\text{ReallocationCosts} := \sum(c, \text{Reallocated}(c)) \times \text{Threshold}$
- $\sum(f, \text{NewAllocation}(c,f)) = 1$ {Single Sourcing constraint}
- $\sum(c, \text{Demand}(c) \times \text{NewAllocation}(c,f)) \geq \text{Supply}(\text{RelaxationLevel}, f)$ {Capacity constraint}
- Given
 - o $\text{Supply}(f)$ per RelaxationLevel
 - o $\text{Demand}(c)$
 - o Threshold
 - o $\text{Distance}(c,f)$
 - o $\text{CurrentAllocation}(c,f)$

Appendix B. Results of the Heuristics

The tables in this appendix show the quality of the different heuristics' solutions. In Tables B.1 (construction) and B.2 (improvement), the best solution shows the numeric value of the objective function (i.e., the total travel and reallocation time), whereas the other solutions show how much worse they are relative to the best solution. The rank and the runtimes in seconds are also shown. Table B.3 shows the results for the GRASP heuristic, including the relative comparison with the greedy adaptive (GA) and first improvement (FI) heuristics. Finally, Table B.4 shows all heuristic performances in relative terms compared with the best solution. That table clearly shows the superiority of the FI heuristic, except for capacity relaxations of 20% and 25%, where it is outperformed slightly.

Table B.1. Solution Quality and Average Runtime of Four Construction Heuristic Algorithms

Capacity relaxation	Greedy sequential (GS)				Greedy random (GR)				Greedy savings regret (GSR)				Greedy adaptive (GA)			
	Value (min)/deviation (%)		Average runtime Rank (s)		Value (min)/deviation (%)		Average runtime Rank (s)		Value (min)/deviation (%)		Average runtime Rank (s)		Value (min)/deviation (%)		Average runtime Rank (s)	
+0%	Infeasible				Infeasible				Infeasible				Infeasible			
+1%	3.5%	4	86		4,583,129	1	320		2.0%	3	89		1.9%	2	170	
+2%	2.5%	4	89		4,529,284	1	323		1.7%	3	89		1.6%	2	171	
+3%	1.9%	4	89		4,493,259	1	322		1.1%	3	89		1.1%	2	171	
+4%	1.3%	4	87		4,473,956	1	316		0.3%	3	88		0.3%	2	170	
+5%	1.4%	4	89		4,433,059	1	318		0.5%	3	89		0.4%	2	170	
+10%	1.3%	4	87		4,325,658	1	316		0.2%	3	89		0.2%	2	170	
+15%	3.0%	4	87		0.6%	3	318		0.1%	2	89		4,221,671	1	171	
+20%	1.2%	4	89		0.5%	2	318		4,187,853	1	88		4,187,853	1	169	
+25%	0.3%	4	87		0.2%	2	321		4,172,590	1	89		4,172,590	1	169	
+50%	4,170,480	1	88		4,170,480	1	322		4,170,480	1	88		4,170,480	1	169	
Uncapacitated	4,170,480	1	87		4,170,480	1	319		4,170,480	1	88		4,170,480	1	170	

Table B.2. Solution Quality and Average Runtime of Three Improvement Heuristic Algorithms

Capacity relaxation	First improvement (FI)				Best improvement (BI)				One-opt			
	Value (min)/deviation (%)	Rank	Number of iterations (average)	Average runtime (s)	Value (min)/deviation (%)	Rank	Number of iterations (average)	Average runtime (s)	Value (min)/deviation (%)	Rank	Number of iterations	Runtime (s)
+0%			Infeasible				Infeasible				Infeasible	
+1%	4,448,055	1	6.5	702	2.5%	3	4.8	551	1.8%	2	4	8,971
+2%	4,407,311	1	6.6	801	0.6%	2	4.0	449	0.6%	3	4	8,134
+3%	4,388,854	1	6.7	813	0.2%	2	3.5	406	1.8%	3	3	7,993
+4%	4,366,015	1	7.5	901	0.3%	2	3.4	413	0.8%	3	3	7,567
+5%	4,348,904	1	7.6	922	0.1%	2	3.5	401	0.7%	3	3	6,896
+10%	4,268,466	1	7.3	908	0.3%	2	3.1	407	0.5%	3	3	6,729
+15%	4,215,021	1	7.2	842	0.4%	2	3.0	394	0.6%	3	3	5,969
+20%	4,198,832	1	7.0	829	0.1%	2	2.7	371	0.2%	3	3	5,709
+25%	0.0%	2	7.0	801	0.0%	3	2.4	269	4,172,590	1	3	5,614
+50%	4,170,480	1	7.0	786	4,170,480	1	2.0	207	4,170,480	1	2	4,975
Uncapacitated	4,170,480	1	7.0	772	4,170,480	1	2.0	206	4,170,480	1	2	4,973

Table B.3. Solution Quality and Average Runtime of the GRASP Metaheuristic Algorithm

Capacity relaxation	GRASP-FI			
	Value (min)	Deviation from GA (%)	Deviation from FI (%)	Average runtime (s)
+0%			Infeasible	
+1%	4,666,931	−0.1%	4.7%	330
+2%	4,599,924	0.0%	4.2%	359
+3%	4,538,091	−0.1%	3.3%	344
+4%	4,485,025	−0.1%	2.7%	361
+5%	4,449,587	0.0%	2.3%	356
+10%	4,335,614	0.0%	1.5%	294
+15%	4,222,682	0.0%	0.2%	276
+20%	4,187,853	0.0%	−0.3%	284
+25%	4,172,590	0.0%	0.0%	279
+50%	4,170,480	0.0%	0.0%	265
Uncapacitated	4,170,480	0.0%	0.0%	256

Note. GRASP-FI, greedy randomized adaptive search procedure–first improvement, FI, first improvement; GA, greedy adaptive.

Table B.4. Comparison of Solution Qualities of the Heuristic Algorithms

Capacity relaxation	Construction				Improvement			Meta
	GS	GR	GSR	GA	FI	BI	One-opt	GRASP
	Deviation (%)							
+0%					N/A			
+1%	6.6%	3.0%	5.1%	5.0%	Best	2.5%	1.8%	4.9%
+2%	5.3%	2.8%	4.5%	4.4%	Best	0.6%	0.6%	4.4%
+3%	4.3%	2.4%	3.5%	3.5%	Best	0.2%	1.8%	3.4%
+4%	3.8%	2.5%	2.8%	2.8%	Best	0.3%	0.8%	2.7%
+5%	3.4%	1.9%	2.4%	2.3%	Best	0.1%	0.7%	2.3%
+10%	2.7%	1.3%	1.5%	1.5%	Best	0.3%	0.5%	1.6%
+15%	3.2%	0.8%	0.3%	0.2%	Best	0.4%	0.6%	0.2%
+20%	0.9%	0.2%	Best	Best	0.3%	0.4%	0.5%	Best
+25%	0.3%	0.2%	Best	Best	0.0%	0.0%	Best	Best
+50%	Best	Best	Best	Best	Best	Best	Best	Best
Uncapacitated	Best	Best	Best	Best	Best	Best	Best	Best

Note. GS, greedy sequential; GR, greedy random; GSR, greedy savings regret; GA, greedy adaptive; FI, first improvement; BI, best improvement; Meta, metaheuristic; GRASP, greedy randomized adaptive search procedure; N/A, not applicable.

References

- Allen DJ, Tobin RL, Calderan A (2017) Verizon optimizes work center locations to reduce installation and repair operations costs. *Interfaces* 47(2):111–121.
- Amrani H, Martel A, Zufferey N, Makeeva P (2011) A variable neighborhood search heuristic for the design of multicommodity production–distribution networks with alternative facility configurations. *OR Spectrum* 33(4):989–1007.
- Barr R, Glover F, Klingman D (1977) The alternating basis algorithm for the assignment problem. *Math. Programming* 13(1):1–13.
- Beasley JE (1990) OR-Library: Distributing test problems by electronic mail. *J. Oper. Res. Soc.* 41(11):1069–1072.
- Beasley JE (1998) Heuristic algorithms for the unconstrained binary quadratic programming problem. Working paper, Management School, Imperial College, London.
- Brandeau ML, Chiu SS (1989) An overview of representative problems in location research. *Management Sci.* 35(6):645–674.
- Buijs P, Alvarez JAL, Veenstra M, Roodbergen KJ (2016) Improved collaborative transport planning at Dutch logistics service provider Fritom. *Interfaces* 46(2):119–132.
- Current J, Min H, Schilling D (1990) Multiobjective analysis of facility location decisions. *Eur. J. Oper. Res.* 49(3):295–307.
- Daskin MS (2013) *Network and Discrete Location: Models, Algorithms, and Applications*, 2nd ed. (John Wiley & Sons, Hoboken, NJ).
- Feo TA, Resende MG (1995) Greedy randomized adaptive search procedures. *J. Global Optim.* 6(2):109–133.
- Gendreau M, Potvin JY, eds. (2010) *Handbook of Metaheuristics*, 2nd ed., vol. 146 (Springer, New York).
- Geoffrion AM, Van Roy TJ (1979) Caution: Common sense planning methods can be hazardous to your corporate health. *Sloan Management Rev.* 20(4):31–42.
- IBM (2017) IBM ILOG CPLEX Optimization Studio—CPLEX user's manual, v12.5. Retrieved January 29, 2020, <https://www.ibm.com/support/knowledgecenter/>.
- Kadijk G, Ligterink N, Spreen J (2015) TNO report: On-road NO_x and CO₂ investigations of Euro 5 light commercial vehicles. Retrieved January 29, 2020, https://www.tno.nl/media/4980/nox_and_co2_investigations_of_commercial_vehicles.pdf.
- Klibi W, Martel A, Guitouni A (2010) The design of robust value-creating supply chain networks: A critical review. *Eur. J. Oper. Res.* 203(2):283–293.
- Klose A, Drexl A (2005) Facility location models for distribution system design. *Eur. J. Oper. Res.* 162(1):4–29.
- Korupolu MR, Plaxton CG, Rajaraman R (2000) Analysis of a local search heuristic for facility location problems. *J. Algorithms* 37(1):146–188.
- Mladenović N, Brimberg J, Hansen P, Moreno-Pérez JA (2007) The *p*-median problem: A survey of metaheuristic approaches. *Eur. J. Oper. Res.* 179(3):927–939.
- Nagy G, Salhi S (2007) Location-routing: Issues, models and methods. *Eur. J. Oper. Res.* 177(2):649–672.
- Öncan T (2007) A survey of the generalized assignment problem and its applications. *Inform. Systems Oper. Res.* 45(3):123–141.
- Reese J (2006) Solution methods for the *p*-median problem: An annotated bibliography. *Networks* 48(3):125–142.
- ReVelle CS, Eiselt HA (2005) Location analysis: A synthesis and survey. *Eur. J. Oper. Res.* 165(1):1–19.
- Roelofs M, Bisschop J (2017) AIMMS—The Language Reference. Retrieved January 29, 2020, https://manual.aimms.com/_downloads/AIMMS_ref.pdf.
- Roobeek A, de Swart JJB, van der Plas M (2018) *Sustainable Business—Making Strategic Decisions to Benefit People, the Planet and Profits*, 1st ed. (KoganPage, London).
- Ross GT, Soland RM (1977) Modeling facility location problems as generalized assignment problems. *Management Sci.* 24(3):345–357.
- Salhi S, Rand GK (1989) The effect of ignoring routes when locating depots. *Eur. J. Oper. Res.* 39(2):150–156.
- Tragantalerngsak S, Holt J, Rönnqvist M (2000) An exact method for the two-echelon, single-source, capacitated facility location problem. *Eur. J. Oper. Res.* 123(3):473–489.

Verification Letter

J. P. M. Remmerswaal, CEO, Van Dorp, Koraalrood 161, 2718 SB Zoetermeer, Netherlands, writes,

“With this writing I wish to inform you of the practical implications of the academic research regarding efficiency improvements in customer allocation for which our company has provided a case study. The outcomes of the research have provided our company both with practical guidelines regarding allocations of customers and with a model that we can further integrate in our business analytics program.

“The company I am responsible for, Van Dorp, is one of the largest service providers in the technical installations industry in the Netherlands. We specialize in service and maintenance of technical installations in utility buildings and have nationwide coverage with 14 regional facilities and approximately 1,000 employees, roughly 300 working in the office and 700 mechanics working on-site.

“Like most of our competitors in the industry, we face profitability challenges even in current economic heydays. We also face challenges regarding scarcity of technical personnel and the reduction of our environmental footprint. And I am happy to report that the implementation of the research's recommendations has resulted in higher profit margins, higher utilization of our mechanics' working hours, and also a reduction in CO₂ emissions as a result of reduced transport movements.

“As was suggested by the author of the research, we started with the quick-win reallocations of customers that were the ‘worst allocated’ with respect to their optimum. The managerial challenges of reallocating customers within our company have been good lessons in general. Though it is difficult to give an exact figure at this point in time, we estimate that long-term savings of those quick wins add up to about €300,000 annually. The majority of these savings are due to productivity increases in our mechanics' working hours, but they are also due to lower operational costs such as mobility costs. The impact on CO₂ emissions is likewise difficult to assess, but we estimate the figure at around 20 tons annually.

“The research on customer reallocation has helped us in assessing our current commercial processes and in particular the division of work among our facilities. As a result, we have redefined our regional facility's commercial areas and improved cooperation between facility directors. While I am writing you this letter, we are implementing the method and algorithms of the research within the framework of our own software tools. The model constructed in the research was new for our business, but it has proven its value and we are grateful for that.”

Christian Haket was the deputy director and sales and marketing manager at Van Dorp at the time of the research. Currently, he is an investment manager of the Anders Invest Solar fund at the private equity firm Anders Invest in the Netherlands. He specializes in strategic and operational management of long-term growth investments in the Dutch

manufacturing industry and is focused on investments in renewable energy generation such as large-scale wind and solar farms.

Bo van der Rhee is a professor of operations management at the Nyenrode Business University in the Netherlands. He studies new product development and supply chain optimization issues and teaches courses on business statistics, management science, operations management, and

research methodology. He is currently the research director for the university.

Jacques de Swart is a professor of applied mathematics at the Nyenrode Business University in the Netherlands. He has an interest in applying Bayesian statistics, simulation, and optimization for strategic and operational decision making. He is also partner at PwC, where he is responsible for the data analytics team within the Dutch consulting practice.