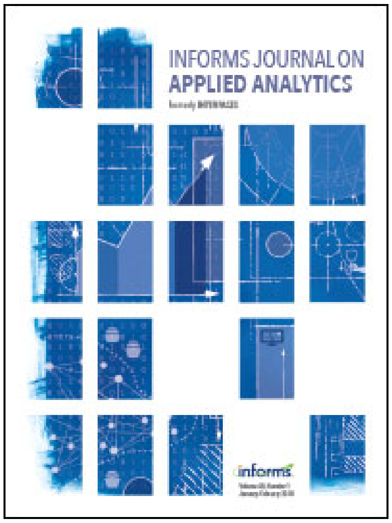


INFORMS Journal on Applied Analytics

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

# A Decision Support System for Attended Home Services

Bruno P. Bruck, Filippo Castegini, Jean-François Cordeau, Manuel Iori, Tommaso Poncemi, Dario Vezzali

informs
http://pubsonline.informs.org/journal/inte

INFORMS JOURNAL ON APPLIED ANALYTICS
Vol. 50, No. 2, March–April 2020, pp. 137–152
ISSN 0092-2102 (print), ISSN 1526-551X (online)
50 YEARS 1970–2020

# A Decision Support System for Attended Home Services

**Bruno P. Bruck,[a] Filippo Castegini,[b] Jean-François Cordeau,[c] Manuel Iori,[b] Tommaso Poncemi,[d] Dario Vezzali[e]**

[a] Centro de Informática, Universidade Federal da Paraíba, CEP 58058-600, João Pessoa, Brazil; [b] Department of Sciences and Methods for Engineering (DISMI), University of Modena and Reggio Emilia, 42122 Reggio Emilia, Italy; [c] HEC Montréal, Montréal, Québec H3T 2A7, Canada; [d] IRETI, External and Metering Operations, 42123 Reggio Emilia, Italy; [e] Marco Biagi Foundation (FMB), University of Modena and Reggio Emilia, 41121 Modena, Italy

**Contact:** bruno.bruck@ci.ufpb.br (BPB); filippocastegini@gmail.com (FC); jean-francois.cordeau@hec.ca (J-FC); manuel.iori@unimore.it, https://orcid.org/0000-0003-2097-6572 (MI); tommaso.poncemi@ireti.it (TP); dario.vezzali@unimore.it (DV)

**Received:** August 23, 2019
**Revised:** December 29, 2019
**Accepted:** December 31, 2019
**Published Online in Articles in Advance:** March 11, 2020

**https://doi.org/10.1287/inte.2020.1031**



**Abstract.** This paper describes a decision support system developed to solve a practical attended home services problem faced by Iren Group, an Italian multiutility company operating in the distribution of electricity, gas, and water. The company operates in several regions across Italy and aims to optimize the dispatching of technicians to customer locations where they perform installations, closures, or maintenance activities within time slots chosen by the customers. The system uses historical data and helps operations managers in performing a number of strategic decisions: grouping municipalities into clusters; designing sets of model-weeks for each cluster; evaluating the obtained solutions by means of a dynamic rolling horizon simulator; and providing as output several key performance indicators, as well as visual optimized technician routing plans to analyze different scenarios. The system uses mathematical models and heuristic algorithms that have been specifically developed to take into account different service levels. Computational experiments carried out on data provided by the company confirm the efficiency of the proposed methods. These methods also constitute a powerful tool that can be used by the company not only to reduce costs but also to help them in their strategic evaluation of existing and potential market opportunities.

**History:** This paper was refereed.
**Funding:** B. P. Bruck acknowledges financial support by Brazilian research agency Conselho Nacional de Desenvolvimento Científico e Tecnológico [Grant 432014/2018-7].

**Keywords:** attended home services • time slot management • routing • decision support system • simulation

## Introduction

*Attended home services* (AHS) are service delivery systems in which a supplying company and a customer agree on a time window during which the customer will be home and the service will be performed. AHS systems are common in many fields, such as the distribution of perishable goods, the delivery of furniture or kitchen appliances and pharmaceutical products, and the provision of repair or technical services of different types (Agatz et al. 2008a, b). Many AHS companies schedule the deliveries by hand and spend considerable resources on schedule evaluation. Their process becomes even more complex because of the dynamics of the activities performed and the presence of demand variability (Campbell and Savelsbergh 2005). Even when problems are solved with optimization tools, it is not obvious that the tradeoff between *quality of service* (QoS) level and cost is perceived as optimal by the customers. Indeed, to fulfill customer intervention requests, the development of an efficient system, providing a high QoS level while balancing the service cost, is essential.

Typically, the optimization of AHS requires solving a two-stage problem, combining appointment scheduling and vehicle routing (Han et al. 2017). In the specific case of IRETI, the optimization of appointment scheduling consists in designing a set of *timetables*, defined by five working days and eight daily time slots of one hour, that are associated to a specific group of municipalities, called *cluster*. Logically, the configuration of clusters is a decision that has to precede the creation of timetables.

Essentially, a timetable is a matrix of time slots in which customer intervention requests are booked by the consumption of a given capacity of allocated resources, corresponding to a certain number of working hours of technicians available to perform the services. The initial configuration of a timetable is called a *model-week* and may change on a seasonal basis according to the expected demand profile. The available time slots are gradually filled with services, and during the booking process, IRETI might dynamically change the model-weeks by moving or adding resources to meet a peak demand. Once the demand is known, the design of routing plans based on customer locations and selected time slots is performed.

As imposed by the authority that regulates the market, IRETI has to respect minimum QoS levels, which may concern the maximum lead time, in terms of working days, between the customer intervention request and the execution of the service, or the maximum delay from the assigned time slot, in terms of hours.

This paper describes a *decision support system* (DSS) designed and implemented to support IRETI in the setup and refinement of operations in a specific territory. In particular, the aim of the DSS is to support IRETI in (i) determining optimal cluster configurations for a given territory, (ii) designing an efficient set of model-weeks by determining the capacity allocated for each time slot, and (iii) simulating detailed routing plans for the technicians.

To solve this three-stage problem, we propose an integrated approach consisting of a series of optimization methods. In particular, the first stage is formulated as a simple *mixed integer linear programming* (MILP) model, based on the well-known *P-median facility location problem* (P-MFLP); the second stage builds on the heuristic algorithm proposed by Bruck et al. (2018); and the third stage is based on a model that dynamically simulates customer intervention requests and their fulfillment by using a rolling-horizon simulation approach in the creation of routing plans.

The introduction of the DSS has dramatically decreased the effort required by IRETI to identify the optimal cluster configuration and create an efficient set of model-weeks for a given territory. Furthermore, by using the DSS, IRETI can evaluate alternative scenarios in terms of strategic *key performance indicators* (KPI) over predetermined time periods (e.g., a week, a month, a year) and visualize the simulated routes of each technician on a real road network. Last, the DSS was initially conceived to optimize the operations of IRETI in a specific territory. However, it could be easily extended to other territories in which IRETI has no historical data or in other contexts, such as electricity and water distribution, thus representing a powerful strategic tool to set up the operations over any territory and context.

The computational results presented in Bruck et al. (2018), obtained without the clustering optimization, show that it is possible to save, on average, 10% in the routing costs considering the most realistic scenario. As expected, these results also showed that imposing QoS constraints has a negative impact on the solution cost, even though such constraints are very important to provide customers with a more substantial set of options and improve the overall satisfaction. This fact is particularly interesting for companies that must compete with each other to get public contracts, where every feature that improves the QoS levels counts. In the present paper, we build on this prior work and extend it to a more realistic setting that yields a powerful tool for strategic planning.

## Context Description

IRETI is a division of Iren Group, an industrial holding company operating in the Italian market of multiutilities, which distributes electricity, gas, and water in several Italian regions, such as Piemonte, Liguria, and Emilia-Romagna (Figure 1). The development of the DSS was justified by a public tender issued by the

**Figure 1.** (Color online) Regions Served and Municipalities in Which IRETI Operates as a Gas Supplier

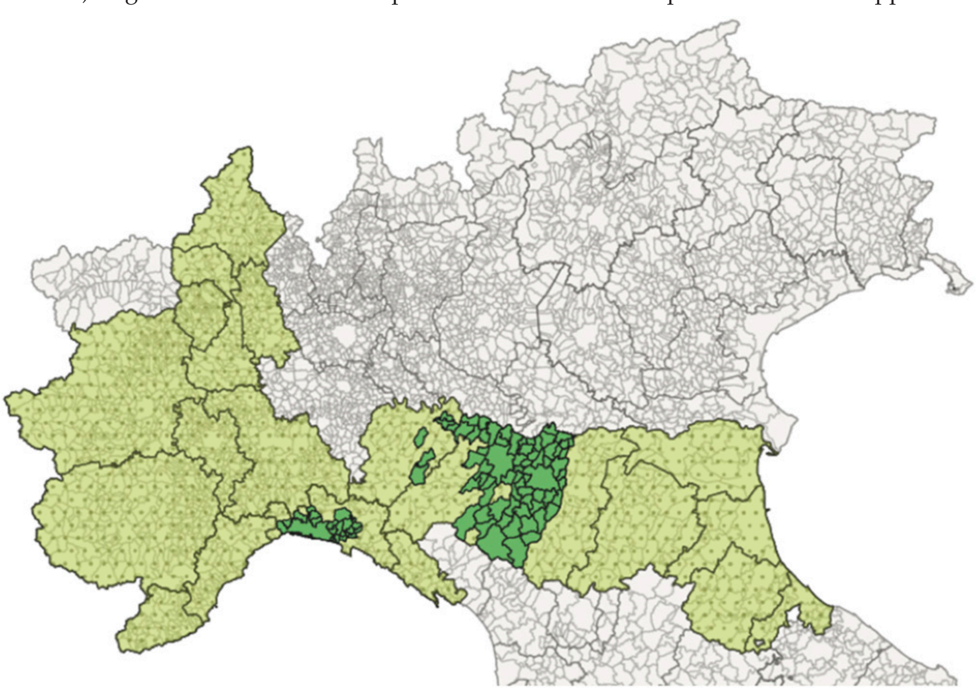

Italian government to renew the concession of gas distribution in an area currently served by IRETI. To clarify the procedure of public tenders in the market of multiutilities, a brief overview is provided.

## European Market of Multiutilities

The European Union is the institution responsible for the definition of directives to which every member state should conform concerning the regulation of the European gas market.

Every member state accepts the directives and adopts national laws to implement them. Local authorities use national laws and adopt specific deliberations to guarantee the observance of the regulations. The Italian market of gas distribution is regulated by the Authority for Regulation of Energy Networks and Environment (ARERA). Through the Accounting Unbundling Obligation (ARERA 2016), called TUIC, the authority defines two main actors: the independent network manager, also denoted as the *distribution company*, and the gas seller, also denoted as the *trading company*. The former operates in the territory assigned through a specific public tender and is responsible for managing and maintaining the gas distribution network. The latter stipulates commercial contracts and is responsible for managing the relationship with customers while competing with other trading companies in a free market. To sum up, several trading companies compete against each other to sell contracts to customers, whereas a single distribution company, selected by means of a public tender, is responsible for operating the gas distribution network.

Until 2012, customer appointments for AHS were defined by mutual agreements (e.g., phone calls) between the distribution company and the customers (Madsen et al. 1995). Now, as a consequence of changes in the institutional and regulatory framework (European Parliament 2009), the distribution company must provide an *online agenda* that consists of a portal containing the model-weeks for all clusters in a served territory. Trading companies book the appointments for their customers in the available time slots directly on the portal, and the distribution company (e.g., IRETI) receives the requests through the online agenda.

## Public Tenders

The public tenders published by ARERA concern the regulation of gas, water, and electricity distribution. The principle behind these tenders is to encourage improvements in the QoS levels by positively evaluating those proposals in which companies commit to reach advances in one or more of these levels. Usually, the QoS is evaluated on the basis of *service time*, defined as the number of working days between the customer intervention request and the execution of service, and *punctuality*, defined as the maximum time at which the technician can arrive at customer's location and start executing the service once a service is booked in a time slot.

Specific QoS levels are defined in the "Regimentation of Quality for Distribution and Measurement of Gas" (ARERA 2013), called TUDG, and are currently valid for the period from 2014 to 2019. The AHSs regulated by the TUDG include the following:

- installing a new meter;
- reopening a meter after closure due to a situation of potential risk;
- reopening a meter after closure due to being in arrears;
- closing a meter at the request of a customer;
- checking a meter at the request of a customer; and
- making available technical data from a meter.

A public tender refers to a so-called *minimum territorial area* (ATEM), a definition introduced by ARERA to represent a cluster of municipalities supplied by the same distribution company. Distribution companies applying to the public call for an ATEM submit their proposals, which are evaluated both in terms of financial evaluation and expected QoS levels compliance. Regarding the latter, the authority usually identifies one or more criteria for which improvements in relation to the current QoS levels are requested. For example, the authority might ask for a 50% reduction in the service time required for installing a new meter, from 10 to 5 days, or in the service time for reopening a meter after closure due to being in arrears, from 2 days to just 1 day. If an applicant distribution company declares to fulfill the tightened QoS levels, a higher technical evaluation is obtained, but the achievement of these performances should be guaranteed to avoid penalties in case of delays in the execution of services. In this sense, the tradeoff between cost and QoS level is crucial, and hence the optimization of operations is a lever to maintain low costs while increasing QoS levels.

The distribution company reaching the highest overall score undertakes the contract in the ATEM for a certain interval of years. This type of tender is very common in public procurement (Pasupathy and Medina-Borja 2008) and has a large number of real-world applications, not only regarding gas distribution, but also in subcontracting other commodities and services.

## Company Description

Iren Group is a large Italian corporate group established in 2010 through the merging of Enìa and Iride. Iren operates in the market of multiutilities with approximately 6,200 employees and achieved a revenue of 3.8 billion euros in 2018. The company is listed on the Milan Stock Exchange and is one of the leaders in its sector. The group consists of an industrial holding company, Iren S.p.A., and four fully controlled business units operating in their specific sectors either directly or through controlled companies in which they hold a share:

- Iren Energia operates in electricity and heat supply, managing some district heating networks and providing technological services;
- Iren Mercato is a trading company that stipulates commercial contracts with customers for the trade of commodities such as electricity, gas, water, and district heating;
- IRETI is a distribution company specialized in gas, electricity, and water distribution networks management; and
- Iren Ambiente operates in the field of waste collection, treatment, and disposal and in the design and management of renewable energy systems.

Based on the aforementioned Accounting Unbundling Obligation, IRETI is an independent network manager operating as a distribution company, whereas Iren Mercato is a trading company. The regulation settled by ARERA (ARERA 2016) imposes administrative and accounting separation for companies operating in gas or electricity distribution markets. In other words, even if distribution and trading companies can be part of the same group, they must be completely independent, both from the operational and the accounting point of view. To encourage fair competition, efficiency, and high QoS levels, they cannot share commercial or sensible information.

Iren Group is the first player in Italy in terms of volume heated in district heating systems, the third in terms of volume of water supplied and collected waste, and the fifth in terms of gas and electricity supplied to final customers.

This study was developed as part of a collaboration between the operations research group of the University of Modena and Reggio Emilia and the external and metering operations unit of IRETI, both based in Reggio Emilia. With almost 8,000 km of distribution network, IRETI provides gas supply to approximately 750,000 customers located in 95 municipalities in Emilia-Romagna, Piemonte, and Liguria, for a total volume of 1.2 billion cubic meters of gas per year. Such distribution consists in transporting gas from the pipelines of *Snam Rete GAS* (the company that manages the Italian national gas transportation system) to the local distribution networks. This includes a number of external operations such as filtering, preheating, and pressure regulation to provide safe and timely services to final users.

## Brief Literature Review

AHSs are commonly designed as the combination of two problems: (i) booking process and (ii) service execution (Han et al. 2017). During the booking process, the customers book a service (either directly or by means of their trading company) in one of the available time slots. Then, the distribution company must perform the services by sending technicians to customers' locations. In most AHS problems, the generation of the routing plans for the technicians can thus be represented as a variant of the vehicle routing problem with time window (VRPTW).

Among the works that focus on the booking process, we cite Punakivi and Saranen (2001), who study different success factors in attended home delivery of grocery, and Campbell and Savelsbergh (2005, 2006), who propose techniques to determine when to accept or reject requests and to influence customer behavior toward low-demand time slots. For the literature that concerns service execution, and thus the routing component of the problem, we refer the reader to the detailed reviews by Agatz et al. (2008b) and Ehmke and Mattfeld (2012) on the VRP for AHS, Baldacci et al. (2012) on the VRPTW, and Toth and Vigo (2014) on VRP variants in general. For the specific VRP faced in the routing phase of our DSS, we refer instead to Bruck et al. (2018), who proposed a heuristic optimization method based on a large neighborhood search (LNS).

A number of studies (see, e.g., Lin and Mahmassani 2002) analyze the correlation between decreasing time window size and increasing delivery costs. A loss in profit when offering shorter time windows is observed also by Campbell and Savelsbergh (2005), who highlight that a significant cost reduction can be reached with longer time windows but obviously at the expense of the QoS level provided. In contrast, the problem considered in this paper combines the necessity of increasing the QoS level and at the same time considering the possibility of decreasing the time window size (as this could be imposed by the authority). Because of these specific characteristics, a cost increase might be unavoidable.

It must be noted that many models in the literature are intended to support the decision makers during the booking period while the final vehicle routes are planned, for instance, by means of commercial routing software. In this sense, a valid contribution of our study is to provide an integrated tool capable of managing both these aspects, analyzing and simulating different scenarios with different assets to guarantee the maximum QoS level and minimize resource deployment.

A related interesting problem that combines booking and routing is the time window assignment vehicle routing problem (TWAVRP), originally proposed by Spliet and Gabor (2014) and later generalized by Spliet and Desaulniers (2015). In the TWAVRP, time windows must be assigned to a fixed set of customers before the demand is known. Once the demand is revealed, a routing plan is constructed with the objective of minimizing the travelling costs. Our problem is, however, different from the TWAVRP because the time windows are not assigned by the company but chosen by the customers.

We are not the first to propose an integrated approach for an AHS problem, as relevant methods have been presented in Madsen et al. (1995) and Agatz et al. (2011), among others. Other integrated approaches have been proposed for other multiutility activities, as the recent DSS developed by Fadda et al. (2018) for urban waste collection.

However, our research may provide a number of innovative and interesting contributions:

- In most of the literature on AHS problems, costs are related to routing (Subramanyam et al. 2018). IRETI, instead, puts a high emphasis on costs related to failures in reaching the required QoS, determined by compensations in favor of customers in the case precision range or service time constraints are not satisfied. This does not mean that IRETI ignores routing costs but that they put more effort in timetable creation to avoid exceeding service time and thus compensations imposed by the authority. In view of this, the timetable creation process takes a very strategic role for the company.
- Most of the literature has focused on heuristic algorithms (see, e.g., Hernandez et al. 2017) in an attempt to automate tasks normally performed by decision makers. Our DSS is equipped with a combination of simulation and MILP models and has been developed with the intent of finding an optimal configuration of assets for a given set of inputs and not just to replace operations. In this respect, the simulation aspect is particularly important, especially for making and evaluating predictions on new territories.
- Our simulation is based on real geographical and historical data, providing the decision makers with meaningful information. This is a relevant component of the first stage, that is, the division of an ATEM into clusters.
- Due to ARERA regulations, we cannot influence customer choices, as suggested, for example, in Campbell and Savelsbergh (2006) and Yang et al. (2014), or reject them because all customers must be served under the same conditions. Thus, demand uncertainty, a typical component of many real-life AHS applications (Cappanera et al., 2018), must be fully considered in the DSS.

To the best of our knowledge, no such integrated strategic tool with the same conditions exists in the AHS literature.

## Problem Description

The AHS problem managed by IRETI is a three-stage problem, consisting of (i) a clustering problem that aims at grouping the municipalities of a certain ATEM into clusters; (ii) a model-week design problem concerning the definition of a time table for each cluster; and (iii) a routing problem that aims at creating optimal routing plans for technicians appointed to perform the services. An in-depth description of the three stages is given in the following sections.

Solving this specific AHS problem means considering several interconnected decisions that increase the complexity of the problem. Therefore, designing an integrated approach capable of proposing cost-effective and balanced clusters, creating good-quality and efficient timetables, and simulating technicians routing plans are complex tasks.

To better understand the whole process, we start by describing the booking process (Figure 2). First, when

**Figure 2.** (Color online) Flowchart Representing the Booking Process for a Service Request

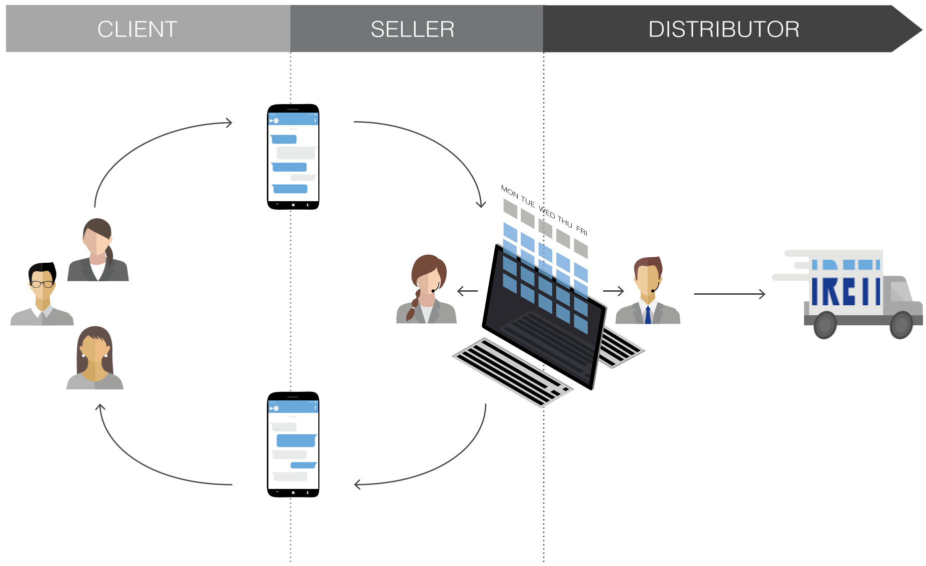

a customer needs to book a service, she must contact her trading company and agree on a date and time slot for the execution. Then, the service request is forwarded to the independent network manager (e.g., IRETI), through the online agenda. At this point, the independent network manager might contact the customer to anticipate the appointment. In case the anticipation is rejected, the date and time slot of the appointment are confirmed. Eventually, the customer might ask to postpone or cancel the appointment. In this case, she must contact her trading company to reschedule the appointment. If the request is made more than 24 hours before the previously arranged appointment, the entire booking process is repeated by searching for a new date and time slot. Only in the case where the request is made less than 24 hours before the previously arranged appointment can the independent network manager reject it. Once an appointment is confirmed, it is assigned to a technician and, if the independent network manager does not respect the required QoS level for the requested service, a compensation in favor of the customer is due. The only situation in which IRETI is relieved of this compensation is when the customer is absent on the execution date.

### Stage 1: Clustering

The clustering is the first stage of the AHS problem managed by IRETI and precedes the creation of the timetables. This stage is solved at a strategic level because the clustering of municipalities is seldom redefined. Until now, this stage has been performed manually by IRETI, based on experience and common sense. For instance, due to the weekly scheduling of outdoor markets or to changes in the road network, decision makers might decide not to group particular municipalities in the same cluster. The proposed method adopts and incorporates these best practices.

### Stage 2: Model-Week Design

The model-week design is the second stage of the problem and is another strategic activity that is performed only a few times over the year, usually on a seasonal basis depending on the demand profile of service requests in a specific ATEM. The definition of a model-week implies deciding which days and time slots to open for each cluster and how many resources to allocate for each time slot, given the availability of technicians. The objective of this stage is to minimize the unbalanced distribution of resources in time slots, between morning and afternoon and among the different days of the week, to provide a high QoS level to customers. IRETI is also responsible for managing the booking process and performing continuous adjustments to the timetables by adding or moving resources to fulfill peak demand.

At the beginning of each week, the timetable is reset to the original model-week for each cluster, except for the resources already booked from the previous weeks. Note that IRETI assumes that a resource corresponds to 30 working minutes of a technician. Given the fact that the length of a time slot is one hour, then two resources can be allocated per technician available. Thus, timetable adjustments refer to the additions or movements of technicians among clusters and time slots. Many types of services require only 30 minutes (i.e., one resource), whereas others require 60 minutes (i.e., two resources). Therefore, a technician could execute a single 60-minute service or two 30-minute services in a time slot. An explanation of how these aspects are managed during the simulation is given in the following sections.

From a practical point of view, the addition or movement of technicians from one cluster to another is performed to fulfill the required QoS level. Currently, the adjustment of timetables is performed manually, but handling a scheduling process in this way is a complex task (see, e.g., Kostuk and Willoughby 2012).

### Stage 3: Simulation of Detailed Routing Plans

The third stage of the problem concerns the building of routing plan for each technician, which must be done on a daily basis. In the proposed method, the routing is also used to evaluate and compare the cost of a given set of model-weeks. Technicians are routed from the depots to customers' locations, and if necessary, additional technicians might be employed from subcontractors that provide outsourcing services. Ideally, technicians are assigned to a determined depot and perform most of the services in the clusters served by that depot. In practice, technicians might be required to perform services outside their area of competence. This might happen in case the allocated capacity is not sufficient to completely fulfill the demand.

## Solution Method

We developed a suite composed by three modules, each solving one of the stages defined in the previous section.

The first module aims at supporting IRETI in the definition of the geographical clusters (stage 1). In this module, the evaluation of travel distances and times between municipalities is essential to determine the best cluster configuration. To that aim, we use the *open source routing machine* (OSRM), a routing engine from OpenStreetMap, which is able to efficiently evaluate the fastest path distance on the real road network between a pair of geographical coordinates received in input. These distances are then used for the generation of clusters, which are obtained by implementing a simple MILP model based on the P-MFLP.

**Figure 3.** (Color online) Three Main Components of the DSS and Their Inner Parts

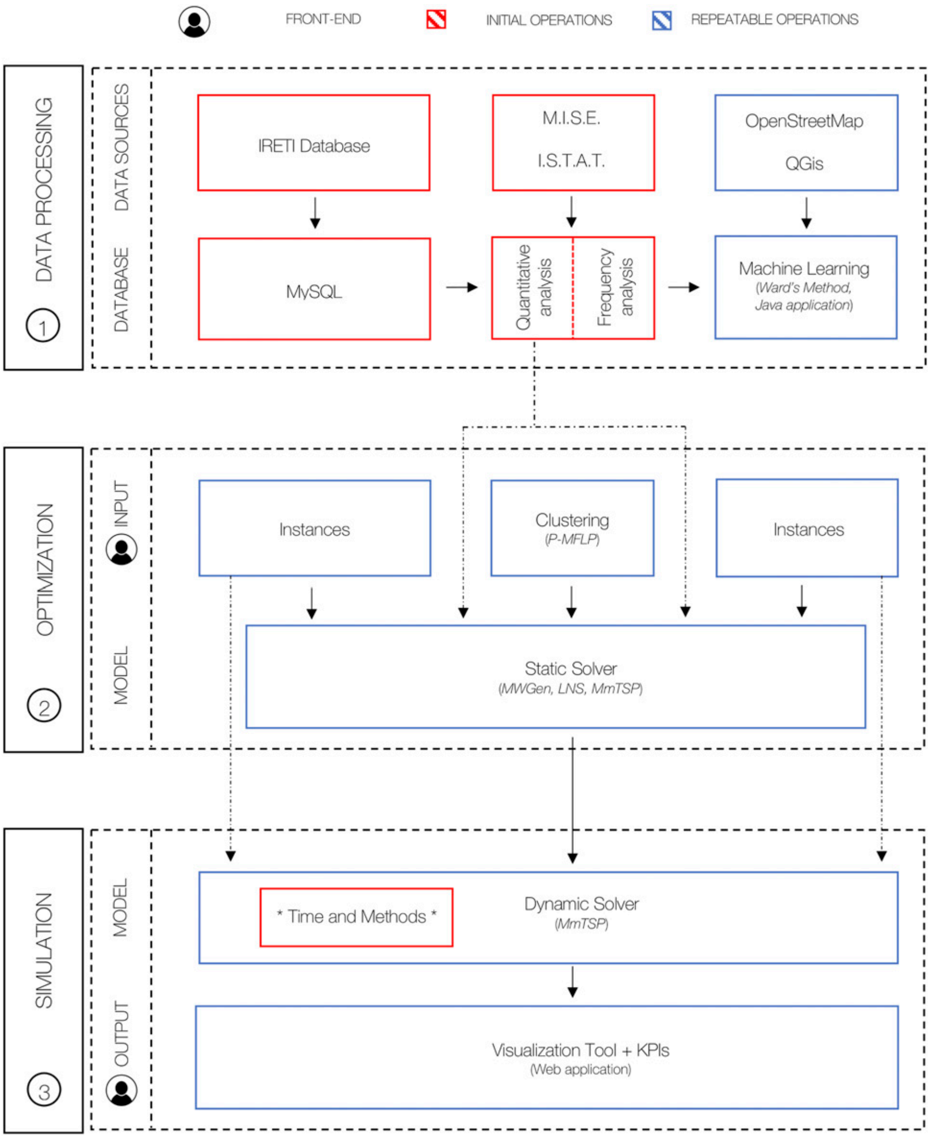


Basically, the clustering stage is characterized by several real constraints that could increase the complexity of the model. Nevertheless, we decided to implement a simple and flexible model, whose constraints could eventually embed some of the recommendations drawn from IRETI. The detailed P-MFLP formulation that we developed is reported in the Appendix.

The second module is the so-called *static solver*, which aims to create a model-week for each cluster, thus solving stage 2. Note that, to estimate the cost of a model week, the second module uses nontrivial algorithms, including a one-week simulation of the booking time slots and the consequent construction of routing plans for the technicians. We implemented this module building on the heuristic approach by Bruck et al. (2018). In particular, their formulation has been replaced by what we call a *model-week generator* (MWGen), which is an MILP model that introduces additional constraints that were not addressed in the original paper. A distinction between clusters containing large and small cities has been introduced, and, according to what the company suggested, different QoS constraints have been adopted. Indeed, as large cities (e.g., the regional county seats) give rise to strong imbalances in the demands, we decided to adopt specific constraints to better mitigate uneven distribution of resources among time slots. Furthermore, the resource constraints are now expressed in terms of number of technicians instead of number of resources. This not only reduces the problem size but also removes some symmetry from the formulation.

Based on the results of the computational experiments, the MWGen formulation has proved to be an effective tool to generate initial feasible model-weeks.

However, it uses a simplified objective function with respect to the real problem, so a further effort is required to understand what would be the resulting operational costs derived from the use of such model-weeks. Therefore, the set of model-weeks produced by the MWGen is given as an input to an LNS algorithm, which evaluates them and possibly modifies them by means of destroy and repair methods. Once a set of model-weeks has been generated, the actual demand registered in a one-week time horizon is revealed, the booking of time slots is simulated by assigning the customer intervention requests to the time slots, and then the routing plan for each technician is created. As in Bruck et al. (2018), the detailed technician routing plan creation is obtained by solving a variant of the *multidepot multiple traveling salesman problem* (MmTSP), in which the total traveled distance is minimized. The LNS continues for a certain number of iterations. At each iteration, the current set of model-weeks is randomly modified and then rebuilt by invoking once more model MWGen. To avoid cycling, we introduced a family of *no-good cuts* into the model. These impose that the configuration of open and closed time slots at a given iteration is different with respect to the configurations obtained in the previous iterations. All these changes helped improve the original model and fulfill the requirements of the new tenders. Model MWGen is described in detail in the appendix (objective function (A.9) and constraints (A.10)–(A.31) are used to generate the initial solution, whereas constraints (A.32)–(A.35) are added to avoid LNS cycling).

The static solver returns a feasible set of model-weeks, that is given to a third module, the *dynamic solver*, which dynamically simulates how customer intervention requests would be satisfied in practice for a long time horizon (typically a few months), thus solving stage 3. Similar to the second module, the creation of routing plans for each technician and each day is modeled as a variant of the MmTSP. Nevertheless, the swapping of resources through different model-weeks, in case of capacity excess, and the use of additional technicians, in case of demand peaks, are introduced. Again, as proposed by Bruck et al. (2018), a *time-extended network* is considered to ensure the respect of time window constraints and, consequently, define the specific schedule of each technician. Based on historical data, we identified three demand scenarios: low, medium, and high. By analyzing the output of this module, the specific KPIs logged during the simulation and the adjustments required in certain time horizons to fulfill the QoS levels imposed by the authority are evaluated. Furthermore, the tool also allows us to visualize and inspect the routes of each technician, providing information on average speed and route duration, among other KPIs.

## Implementation of the DSS

We embedded the proposed methods into a DSS that consists of three macro modules: data processing, optimization, and simulation (Figure 3). The data processing module represents the interface between the DSS and IRETI historical data. The optimization module is responsible for creating instances, generating cluster configurations by running the P-MFLP model, and designing sets of model-weeks by means of the static solver. The simulation module then simulates under different scenarios the proposed solution, given a set of previously designed model-weeks. The results of these simulations can be visualized in a web application, allowing decision makers to verify strategic KPIs. In the following, we analyze in detail each module.

### Data Processing

The data processing module consists of a set of methods, tools, and scripts that are used to preprocess and cleanse all of the data extracted from the IRETI database, which contains the list of intervention requests, performed services, and georeferences over an extended interval of years. Additional information might be integrated from the open data published by the National Institute of Statistics (ISTAT) and the Ministry of Economic Development (MISE). In particular, demographics and geographical data on administrative borders are extracted from the previous source, whereas technical data (e.g., users per municipality, volume of distributed gas, network distribution length, and category of supply) are obtained from the latter.

In case of a public call referred to a particular ATEM that is completely or partially not served by IRETI and for which historical data are not available or incomplete, we implemented simple machine learning methods that rebuild or predict missing information by processing determined features and historical

**Figure 4.** (Color online) Cluster Configuration for ATEM of Verona

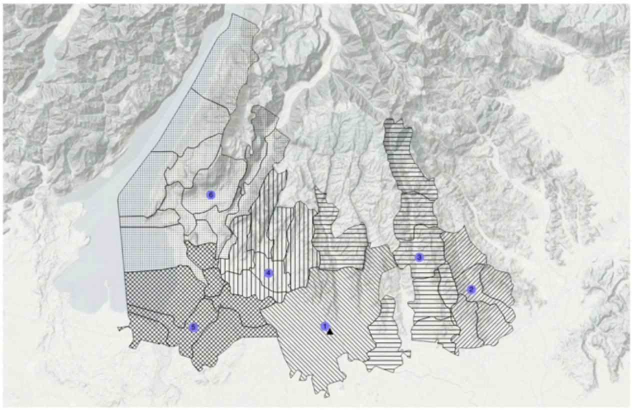

available data, both from other served ATEMs and from already served municipalities in the considered ATEM. The idea of this module is to create dummy intervention requests for all of the unknown municipalities, based on similarities with already served and known territories.

In the following, a brief description of the implementation of the aforementioned machine learning methods is provided.

Initially, the additional features obtained from the open data published by ISTAT and MISE are given as an input to a hierarchical clustering algorithm, known as Ward's method (WM) (Murtagh and Contreras 2012) and integrated in the software RStudio. That constitutes the first part of the machine learning algorithm, for which the output is a dendogram successively converted into a similarity matrix where the element $a_{ij}$ is a parameter that takes the value 1 if municipality $i$ is similar to municipality $j$ and 0 otherwise. After that conversion, the similarity matrix is given as an input to a Java application, along with a series of additional inputs such as the following:

- the calendar year in which the rebuilding is performed;
- the frequency analysis of historical customer intervention requests; and
- the set of real building coordinates in the ATEM aim of the public tender, previously extracted using a query of the QGis open-source software on an OpenStreetMap layer.

The Java application output is a set of dummy intervention requests that, together with historical

**Table 1.** Model-Week Tables for ATEM of Verona

| Cluster 1 | | | | | | Cluster 2 | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Time slots | Mon | Tue | Wed | Thu | Fri | Time slots | Mon | Tue | Wed | Thu | Fri |
| 08:30–09:30 | 10 | 10 | 12 | 12 | 12 | 08:30–09:30 | 0 | 2 | 0 | 0 | 0 |
| 09:30–10:30 | 10 | 10 | 12 | 12 | 12 | 09:30–10:30 | 0 | 2 | 0 | 0 | 0 |
| 10:30–11:30 | 10 | 10 | 12 | 12 | 12 | 10:30–11:30 | 0 | 2 | 0 | 0 | 0 |
| 11:30–12:30 | 10 | 10 | 12 | 12 | 12 | 11:30–12:30 | 0 | 2 | 0 | 0 | 0 |
| 12:30–13:30 | — | — | — | — | — | 12:30–13:30 | — | — | — | — | — |
| 13:30–14:30 | 10 | 10 | 8 | 8 | 8 | 13:30–14:30 | 0 | 0 | 0 | 0 | 2 |
| 14:30–15:30 | 10 | 10 | 8 | 8 | 8 | 14:30–15:30 | 0 | 0 | 0 | 0 | 2 |
| 15:30–16:30 | 10 | 10 | 8 | 8 | 8 | 15:30–16:30 | 0 | 0 | 0 | 0 | 2 |
| 16:30–17:30 | 10 | 10 | 8 | 8 | 8 | 16:30–17:30 | 0 | 0 | 0 | 0 | 2 |
| Cluster 3 | | | | | | Cluster 4 | | | | | |
| Time slots | Mon | Tue | Wed | Thu | Fri | Time slots | Mon | Tue | Wed | Thu | Fri |
| 08:30–09:30 | 0 | 2 | 2 | 0 | 2 | 08:30–09:30 | 2 | 2 | 2 | 2 | 0 |
| 09:30–10:30 | 0 | 2 | 2 | 0 | 2 | 09:30–10:30 | 2 | 2 | 2 | 2 | 0 |
| 10:30–11:30 | 0 | 2 | 2 | 0 | 2 | 10:30–11:30 | 2 | 2 | 2 | 2 | 0 |
| 11:30–12:30 | 0 | 2 | 2 | 0 | 0 | 11:30–12:30 | 2 | 0 | 2 | 2 | 0 |
| 12:30–13:30 | — | — | — | — | — | 12:30–13:30 | — | — | — | — | — |
| 13:30–14:30 | 0 | 2 | 2 | 0 | 0 | 13:30–14:30 | 0 | 2 | 2 | 2 | 0 |
| 14:30–15:30 | 0 | 2 | 2 | 0 | 2 | 14:30–15:30 | 2 | 2 | 2 | 2 | 0 |
| 15:30–16:30 | 0 | 2 | 2 | 0 | 2 | 15:30–16:30 | 2 | 2 | 2 | 2 | 0 |
| 16:30–17:30 | 0 | 2 | 2 | 0 | 2 | 16:30–17:30 | 2 | 2 | 2 | 2 | 0 |
| Cluster 5 | | | | | | Cluster 6 | | | | | |
| Time slots | Mon | Tue | Wed | Thu | Fri | Time slots | Mon | Tue | Wed | Thu | Fri |
| 08:30–09:30 | 2 | 2 | 4 | 2 | 2 | 08:30–09:30 | 0 | 2 | 0 | 4 | 2 |
| 09:30–10:30 | 2 | 2 | 4 | 2 | 2 | 09:30–10:30 | 0 | 2 | 0 | 4 | 2 |
| 10:30–11:30 | 2 | 2 | 2 | 2 | 4 | 10:30–11:30 | 0 | 4 | 0 | 2 | 2 |
| 11:30–12:30 | 2 | 2 | 2 | 2 | 4 | 11:30–12:30 | 0 | 4 | 0 | 2 | 2 |
| 12:30–13:30 | — | — | — | — | — | 12:30–13:30 | — | — | — | — | — |
| 13:30–14:30 | 2 | 2 | 2 | 2 | 2 | 13:30–14:30 | 0 | 2 | 0 | 2 | 2 |
| 14:30–15:30 | 2 | 2 | 2 | 2 | 2 | 14:30–15:30 | 0 | 2 | 0 | 2 | 2 |
| 15:30–16:30 | 2 | 2 | 2 | 2 | 2 | 15:30–16:30 | 0 | 2 | 0 | 2 | 2 |
| 16:30–17:30 | 2 | 2 | 2 | 2 | 2 | 16:30–17:30 | 0 | 2 | 0 | 2 | 2 |

**Table 2.** All Clusters Table: Sum of All Model-Week Tables

| | All clusters | | | | | |
|---|---|---|---|---|---|---|
| Time slots | Mon | Tue | Wed | Thu | Fri | Total |
| 08:30–09:30 | 14 | 20 | 20 | 20 | 18 | 92 |
| 09:30–10:30 | 14 | 20 | 20 | 20 | 18 | 92 |
| 10:30–11:30 | 14 | 22 | 18 | 18 | 20 | 92 |
| 11:30–12:30 | 14 | 20 | 18 | 18 | 18 | 88 |
| 12:30–13:30 | — | — | — | — | — | — |
| 13:30–14:30 | 12 | 18 | 14 | 14 | 14 | 72 |
| 14:30–15:30 | 14 | 18 | 14 | 14 | 16 | 76 |
| 15:30–16:30 | 14 | 18 | 14 | 14 | 16 | 76 |
| 16:30–17:30 | 14 | 18 | 14 | 14 | 16 | 76 |
| Total | 110 | 154 | 132 | 132 | 136 | 664 |

customer intervention requests, are added to a local database used to create sets of instances that are given as an input to the following modules.

## Optimization

The optimization module contains the tool responsible for generating optimized cluster configurations and the static solver. Both tools rely on the database for input data. However, they do not have direct access to it. Instead, we generate instance files that comprise data from specific periods of time that we consider in the simulations. These files are created by Python scripts that search through the database and select only the requested data. The database can be populated either with real data or with the machine learning algorithm.

The static solver requires additional input information, such as the cluster configuration, which can be either given by the company or generated using our clustering algorithm. Note that, if we choose to use the latter approach, it is important to consult the company experts to ensure that the proposed cluster configuration is indeed viable. In addition, it is necessary to input the demand scenario for the static solver, which basically defines the amount of demand that will be covered during the generation of the model-week tables. At this point, the decision makers are ready to run the static solver through a command line and select some customization elements to reach different QoS levels.

It is important to mention that the solver might not always be able to find a feasible solution. This might happen due to several reasons. For instance, the given number of technicians for each depot may be insufficient or the number of available time slots may be too low to reach the specified QoS levels. In these cases, the decision maker can analyze input and output data and correct any inconsistency or underestimation of the workload to overcome the problem and then run the solver again.

## Simulation

The simulation module refers to the dynamic solver and to the simulations that can be performed to fine tune the solutions found by the optimization module. Instead of considering a single week of data, the instance files for the dynamic solver may contain information from an arbitrary period of time. In our experiments, we usually run simulations for an entire year to test the efficiency of a given solution in the seasons of low, medium, and high demand.

In the dynamic solver, decision makers can customize specific parameters of the simulation such as the maximum execution time of each type of service and decide whether to allow the swap of resources between timetables to fulfill requests in peak demand scenarios. In particular, the former parameter is crucial to ensure that the company is able to simulate possible scenarios that might be presented in future tender roles. Furthermore, to perform more realistic simulations, the dynamic solver considers that the execution times of services from the same type may vary by a certain degree, which can be specified by the decision maker. Furthermore, during the simulation, a delay in the schedule of a technicians is tolerable and logged as a KPI. Consequently, an evaluation of delays is given as an output to the decision maker who can assess whether the total amount of delay is acceptable or not in the simulated time horizon.

## Use and Benefits

The DSS is a strategic tool for IRETI that substantially decreases the decision-making process, providing more efficient solutions and exploring different demand scenarios that would be hard and time consuming to

**Figure 5.** KPIs Provided in Output

| Simulation period | | | Demand | | Dynamic manipulation | | Lateness [One-hour precision range] | | | Driving distance [km] | | Driving time [hh:mm] |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Weeks | From | To | #appointments | #resources (30 min.) | #technicians movements | #technicians additions | #late appointments | % on total week appointments | Avg. Lateness [hh:mm] | Total driving km | Km per appointment | Driving time per appointment |
| 1 | 04/07 | 08/07 | 211 | 296 | 3 | 0 | 3 | 1% | 00:02 | 2954 | 14,0 | 00:18 |
| 2 | 11/07 | 15/07 | 360 | 532 | 7 | 0 | 11 | 3% | 00:20 | 4365 | 12,1 | 00:16 |
| 3 | 18/07 | 22/07 | 417 | 637 | 13 | 0 | 14 | 3% | 00:31 | 4809 | 11,5 | 00:15 |
| 4 | 25/07 | 29/07 | 418 | 632 | 9 | 0 | 14 | 3% | 00:23 | 4885 | 11,7 | 00:16 |
| 5 | 01/08 | 05/08 | 415 | 639 | 22 | 0 | 5 | 1% | 00:16 | 4981 | 12,0 | 00:16 |
| 6 | 08/08 | 12/08 | 376 | 578 | 12 | 0 | 5 | 1% | 00:26 | 4379 | 11,6 | 00:15 |
| 7 | 15/08 | 19/08 | 369 | 575 | 6 | 0 | 11 | 3% | 00:24 | 4424 | 12,0 | 00:16 |
| 8 | 22/08 | 26/08 | 240 | 401 | 5 | 0 | 2 | 1% | 00:10 | 2598 | 10,8 | 00:15 |
| **Avg Values** | | | **351** | **536** | **10** | **0** | **8** | **2%** | **00:19** | **4174** | **12** | **00:16** |

compute manually. The system also allows the team to consider different QoS levels in the simulations, which is important when trying to establish a reasonable tradeoff between the QoS offered to clients and the costs incurred by the company.

Prior to using our DSS, most of the decisions made by IRETI relied on analysis performed manually, with the aid of several spreadsheets. However, due to the sheer amount of data involved, the number of constraints, and possible scenarios, finding a good solution was a difficult task. Moreover, due to the complexity of the decisions that had to be made and considering that the process was mostly manual, the team was not able to reliably plan for a longer time horizon, having to limit their forecast to only a few weeks ahead. Besides being able to automate most of the process of generating efficient solutions, the DSS also provides easy access to the results of the simulations (i.e., graphs and tables) so that alternatives can be evaluated, and more informed decisions that do not rely solely on experience can be made.

Currently, IRETI is using the DSS in preparation for tenders that would give the opportunity to the company to operate on new territories or confirm its position in already served ones. The simulations performed using the DSS enable the decision makers to experiment and analyze how they can provide a high QoS level to customers while minimizing the cost. The DSS might also represent a tool to analyze past decisions and review them either to improve the current KPIs or to prepare for critical issues that might happen, being able to react quickly and accurately. For instance, in a currently served ATEM, the evaluation of alternative cluster configurations would be performed, aiming at reducing costs while maintaining or increasing the QoS level.

Another benefit of the system is that it is very flexible and can be adapted to changes in the regulations imposed by the market or to new strategies that the company might be interested to adopt. In a sector that is continuously changing, this sort of flexibility is crucial to the success of such an application. For example, in anticipation of possible changes in the regulations of water distribution, the company could already start up a tailored analysis by simply modifying the input data and performing some small adjustments to the constraints of the models.

## A Realistic Instance

Due to an agreement to confidentiality, we cannot present the results of a specific scenario that was optimized by the company. However, to illustrate and show the flexibility of the DSS, in the following, we present the results that were found using the system to analyze and find a solution for a possible tender in the ATEM of Verona, which might be of interest to the company. The data for the instance were generated using the machine learning approach described in the Data Processing section. To assess the efficiency of the proposed approach, we performed a range of experiments using both the static and dynamic solvers to simulate different demand scenarios and varying the execution times required by each type of service. In all cases, the DSS obtained efficient solutions even when considering tightened QoS constraints than what is required in real tenders.

**Figure 6.** (Color online) Visualization Interface: View of Routes for All Technicians during a Week

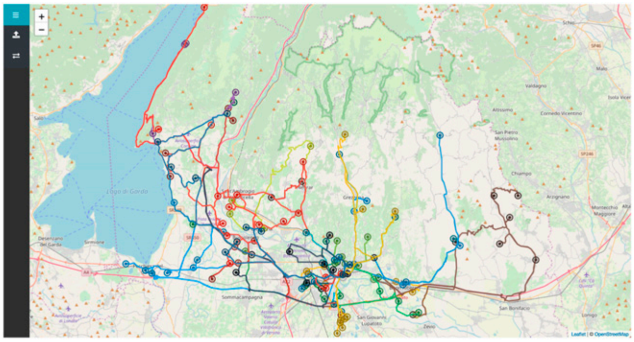

The cluster configuration created for this ATEM is depicted in Figure 4, where the depot is represented by a black triangle. The most efficient set of model-weeks is reported in Table 1. Note that, in this solution, cluster 1 corresponds to the county seat of Verona, and due to its predicted high demand, all time slots are opened, offering a high QoS level with an average number of technicians per time slot equal to five, whereas all other clusters have an average of one or two technicians. Two other interesting features with respect to QoS are that all clusters have time slots opened in at least two days of the week, and there is a balance between the distribution of time slots in the morning and afternoon. In Table 2, we show the sum

**Figure 7.** (Color online) Visualization Interface: Detailed Daily Route Plan for a Single Technician

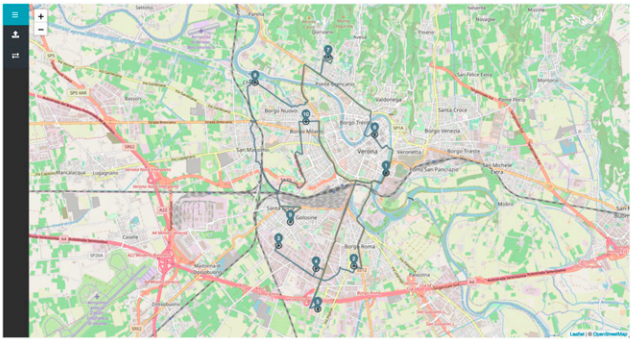

**Figure 8.** Service Time Output Chart

| Activity code | Avg simulated service time (working days) | Current QoS level | Hypothetic new QoS level |
|---|---|---|---|
| A01 | 4 | 10 | 5 |
| D01 | 2 | 5 | 2 |
| A40 | 4 | 10 | 5 |
| V01 | 10 | 20 | 10 |
| M01 | 4 | 10 | 5 |
| M02 | 3 | 15 | 8 |

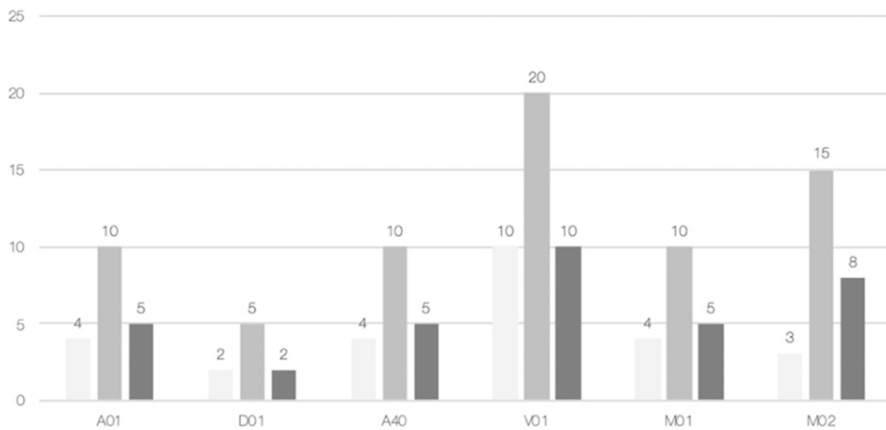


of all resources allocated to each time slot. This table is useful to analyze the distribution of resources per day and per time slot. Note that a total of 664 resources are provided, corresponding to 332 working hours per week, with well-balanced values per day (ranging from 110 to 154 resources) and per hour of the day (between 72 and 92).

In practice, the model-weeks are gradually populated as service requests are made by customers through the trading companies. Especially during high demand profile periods, the timetables might be altered to face demand peaks that would otherwise not be completely fulfilled with the resources that are normally allocated. In these scenarios, the decision makers usually have two options: (i) diverting technicians from one cluster to the other or (ii) employing third party technicians to compensate for the demand peak. Note that, although the former alternative does not lead to additional costs for the company, it might not always be feasible. Moreover, given that both alternatives lead to changes in the model-weeks, there is always the risk that poor decisions are made, leading to a higher increase in costs. Therefore, it is very important to simulate different scenarios in terms of demand profile when designing the configuration of model-weeks, not only to have a good initial solution but also to understand system behavior and to estimate which dynamic changes are expected to be performed. This is the purpose of the dynamic solver that is integrated in the DSS.

In the particular case of the ATEM of Verona, we present the results of a two-month (eight-week) simulation based on a realistic instance related to July and August. This is a typical medium demand profile scenario, and each intervention request is revealed to the dynamic solver in the day in which it was supposedly requested by the customer. The appointment is booked in the first available time slot of the cluster online agenda to which the customer belongs. In case no time slot allows the imposed QoS levels to be satisfied, the dynamic solver simulates the response from IRETI, providing another date, time slot, or both. The dynamic solver gives an evaluation of the performed simulation reporting different strategic KPIs such as (i) delays within one and two hours, (ii) driving times, and (iii) distance evaluations. An example of output provided by the DSS and containing the aforementioned KPIs is presented in Figure 5. Based on the obtained KPIs, a decision maker could conclude that no additional technicians are required during the simulation. This means that the number of technicians provided as an input is consistent. By contrast, comparing this scenario with an identical one except for the number of available technicians, reduced to 7 instead of 10, would return that three technicians should be added during weeks 2 and 6. Furthermore, in the current solution, week 6 could be characterized from a significant demand variation, due to the high number of technician swaps. In addition, the percentage of delays is always less than 3%, highlighting that delays are minimal with the reduction of the precision range. Note that a variable number of appointments (between 210 and 420) corresponds to a variable number of resources (between about 295 and 635).

Contextually, the system provides detailed information on the performance of technicians, analyzing their daily routes and giving evidence to the compliance with the QoS levels. The minimum routes on which the operators can move and the relative travel times are computed. The covered distance is about 4,100 km/wk on average, and it is consistent with the number of appointments, resulting in 12 km per appointment, with a minimum of 10.8 and a maximum of 12.1 km. Considering the dimensions of the ATEM, a KPI of about 12 km per appointment is plausible. Even the traveled time is compatible with the activities performed, as technicians spend an average of 16 minutes driving from one appointment to another. The respect of the imposed QoS levels has been achieved through continuous dynamic changes in the model-weeks, moving an average of 10 resources per week, which correspond to about two technicians per day. A visualization tool has been developed to test the set of routes with the company. Figures 6 and 7 show, respectively, a screenshot where all of the routes used by technicians during a week are depicted and another screenshot in which the detail of one of these routes in a particular day is displayed.

Another table is created for service times. Figure 8 reports the comparisons between different service times for each activity (e.g., average simulated service time, current QoS level, and hypothetical new QoS level). From this report, IRETI can verify if, on average, some activity is exceeding the hypothetical new QoS level. In our simulation, with the chosen inputs, all the activities respect the QoS levels. The simulated service time shows considerable improvements, especially concerning the activities D01 and M02.

## Conclusions

In this paper, we presented a DSS that was developed as a strategic tool to support IRETI in the achievement of new QoS levels for possible tenders. A major benefit of using the DSS is the possibility to simulate several different scenarios, trying to identify the best solutions for a given territory. The DSS is flexible, easy to replicate, and allows the company to simulate service performance in territories that are not currently served. As an example, IRETI is using the system to prepare technical documents for new tenders. The current version of the DSS is a prototype that requires specialized knowledge, given the multiple information technologies used. We do not, however, foresee any technical difficulty in implementing a version of the system comprising all its current features into a more user-friendly package (e.g., Hanne et al. 2009).

Furthermore, as suggested in Bent and Van Hentenryck (2004), another contribution for this study could be to develop new strategies to handle situations in which demand exceeds capacity. At the moment, in these scenarios, we consider that IRETI is capable of diverting resources from other sectors to make sure that all services are fulfilled in time, which obviously increases the cost for the company. We plan to modify the system to analyze the consequences of such planning strategy and compare with a scenario where it is possible to postpone services at the cost of a penalty per service that would be imposed to the company. The aim would be to find a tradeoff between diverting resources and accepting that paying the penalty sometimes might be less costly.

### Acknowledgments

The authors thank three anonymous referees for their valuable comments which helped improve the quality of the paper.

### Appendix. Mathematical Models

This appendix contains the details of the two mathematical formulations that are used for optimization purposes within the DSS.

#### Formulation P-MFLP

Formulation P-MFLP is the core of the clustering part of the solution approach described in the Problem Solution section and is used for dividing the territory managed by the company (ATEM) into clusters of municipalities that will share the same time table. Let $n$ be the number of municipalities in the ATEM and $c_{ij}$ be the traveling distance from the center of municipality $i$ to the center of municipality $j$, for $i, j = 1, \ldots, n$. Let also $q_i$ be the expected demand of $i$, and $F$ be a set of incompatible pairs of municipalities that cannot be part of the same cluster. Set $F$ is useful to impose additional constraints possibly required by the company based on practical experience and knowledge of the territory. The aim of model P-MFLP is to assign all municipalities to exactly $p$ clusters, in such a way that the sum of the distances between each municipality and the centroid of its cluster is minimized and a maximum unbalance among the total demands of the clusters is limited by an input ratio $\alpha$. In our model, cluster centroids must be chosen in the centers of the municipalities, and $p$ is an input parameter that can be varied to test different solutions. In other words, we need to select $p$ municipalities, among the $n$ available ones, whose centers will serve as cluster centroids.

Let $y_i$ be a binary variable that takes the value 1 in case cluster $i$ is selected and 0 otherwise for $i = 1, \ldots, n$, which means that a cluster with centroid located in the center of municipality $i$ is opened. In addition, let $x_{ij}$ be a binary variable that takes the value 1 if municipality $j$ is assigned to cluster $i$ and 0 otherwise for $i, j = 1, \ldots, n$. We obtain the following formulation.

$$\text{(P–MFLP)} \qquad \min z_{(\text{P–MFLP})} = \sum_{i=1}^{n} \sum_{j=1}^{n} c_{ij} x_{ij}; \tag{A.1}$$

s.t.

$$\sum_{i=1}^{n} y_i = p; \tag{A.2}$$

$$\sum_{i=1}^{n} x_{ij} = 1 \qquad j = 1, \ldots, n; \tag{A.3}$$

$$x_{ij} \leq y_i \qquad i, j = 1, \ldots, n; \tag{A.4}$$

$$\sum_{j=1}^{n} q_j x_{ij} \leq \left\lceil \sum_{j=1}^{n} q_j \Big/ p \right\rceil (1 + \alpha) \qquad i = 1, \ldots, n; \tag{A.5}$$

$$x_{ij} + x_{it} \leq 1 \qquad i = 1, \ldots, n; (j, t) \in F; \tag{A.6}$$

$$y_i \in \{0,1\} \qquad i = 1, \ldots, n; \tag{A.7}$$

$$x_{ij} \in \{0,1\} \qquad i, j = 1, \ldots, n. \tag{A.8}$$

The objective function (A.1) minimizes the total distance between each municipality and the centroid of its cluster. Constraints (A.2) ensure that exactly $p$ clusters are created, and constraints (A.3) specify that each city must be assigned to a cluster. Constraints (A.4) impose that municipality $j$ can be assigned to cluster $i$ only if $i$ has been selected (note that each municipality is a potential centroid of a cluster, so both indices $i$ and $j$ vary from 1 to $n$). Constraints (A.5) aim at balancing the distribution of resources among the clusters by ensuring that the demand of each cluster does not exceed the average demand per cluster by more than $\alpha$. Note that smaller values of $\alpha$ lead to clusters where the demand is well balanced, whereas larger ones neglect balancing but allow lower cost solutions to be obtained. In our experiments, we found good results with $\alpha \in [0.2, 0.6]$. Finally, constraints (A.6) specify that any pair of incompatible municipalities cannot be assigned to the same cluster, and (A.7) and (A.8) impose integrality on the variables.

## Formulation MWGen

The mathematical model described in this section is based on the formulation originally proposed by Bruck et al. (2018) but has been modified to better fit the needs of the project and generate more robust solutions. As described in the Problem Solution section, model MWGen is used for generating initial good-quality model-weeks for the clusters obtained by formulation P-MFLP that are then passed to the LNS for further optimization. Before presenting the model, we first need to introduce some notation.

We are given a set $M$ of depots, which are spread out in the territory managed by the company (ATEM). The territory is divided in a set $R$ of clusters. Each cluster is associated with a depot, in such a way that $R = \cup_{i \in M} R_i$, with $R_i$ being the subset of clusters associated with depot $i \in M$. Recall from the Introduction that we use the term *resource* to specify a time interval of half an hour, given that in our case study, services require either half an hour or one hour. Each cluster $r \in R$ has an expected demand of $q_r$ resources. Clusters that have significantly high demand values usually contain large cities, which require different QoS levels. These clusters are grouped in a subset $R_0 \subseteq R$. Furthermore, each depot $i$ has a certain number $Q_i$ of technicians that are available to perform services. Each technician is able to provide at most $\sigma$ resources per time slot, with $\sigma = 2$ in our case study, and its route must always start and end at the same depot.

The time horizon is divided into a set $D$ of days, each of which is further split into a set $T$ of nonoverlapping time slots. In our case study, $|D| = 5$ (corresponding to the working days in a week, from Monday to Friday) and $|T| = 9$ (corresponding to intervals of one hour each from 0830 to 1730 hours, where the interval 1230–1330 is reserved for the lunch break). In the following, a *time slot* is defined by a pair $(d, t)$, with $d \in D$ and $t \in T$, a *timetable* is an assignment of resources to the time slots, and a *solution* to the problem is a collection of time slot tables, one per cluster.

Let $u_{rdt}$ be an integer variable that specifies the number of technicians assigned to the time slot $(d, t)$ of cluster $r$. Note that, differently from Bruck et al. (2018), instead of resources, we assign technicians to time slots. This is a simple optimization based on the observation that each technician is always able to perform $\sigma$ resources per time slot, and thus any variable $u$ that is not a multiple of $\sigma$ could be rounded up to the next multiple. Furthermore, let $z_r$ be an integer variable specifying the number of working hours that could not be allocated to the available technicians in a certain cluster $r \in R$. These working hours will be assigned to a third-party logistics provider. In addition, let $v_{rdt}$ be an integer variable that determines the difference in the number of allocated technicians between two consecutive time slots. These variables are used to measure how well the technicians assigned to a certain cluster $r$ are distributed among consecutive time slots along the day. To penalize uneven distribution of technicians per day, we introduce variables $\rho_{\min}$ and $\rho_{\max}$, that evaluate, respectively, the minimum and maximum number of technicians assigned per day to large clusters, that is, to any cluster $r \in R_0$. We also define variables $\eta_{\min}$ and $\eta_{\max}$ for the same purpose, but this time we assign them to any cluster $r \in R \backslash R_0$. We are now ready to introduce the MWGen formulation.

$$\text{(MWGen)} \quad \min z_{(\text{MWGen})} = \sum_{d \in D} \pi (\rho_{\max} - \rho_{\min}) + \sum_{d \in D} \pi (\eta_{\max} - \eta_{\min}) + \sum_{r \in R} \sum_{d \in D} \sum_{t \in T} \gamma v_{rdt} + \sum_{r \in R} \Omega z_r \tag{A.9}$$

s.t.

$$\sum_{d \in D} \sum_{t \in T} u_{rdt} = \frac{q_r}{\sigma} - z_r \qquad \forall r \in R; \tag{A.10}$$

$$z_r \leq \lambda_1 \frac{q_r}{\sigma} \qquad \forall r \in R, d \in D, t \in T; \tag{A.11}$$

$$\sum_{r \in R_i} u_{rdt} \leq Q_i \qquad \forall i \in M, d \in D, t \in T; \tag{A.12}$$

$$\sum_{t \in T} u_{rdt} \geq \rho_{\min} \qquad \forall d \in D, r \in R_0; \tag{A.13}$$

$$\sum_{t \in T} u_{rdt} \leq \rho_{\max} \qquad \forall d \in D, r \in R_0; \tag{A.14}$$

$$\sum_{t \in T} u_{rdt} \geq \eta_{\min} \qquad \forall d \in D, r \in R \backslash R_0; \tag{A.15}$$

$$\sum_{t \in T} u_{rdt} \leq \eta_{\max} \qquad \forall d \in D, r \in R \backslash R_0; \tag{A.16}$$

$$v_{rdt} \geq u_{rdt} - u_{rd,t+1} \qquad \forall r \in R, d \in D, t \in T \backslash \{\ell - 1, \ell, |T|\}; \tag{A.17}$$

$$v_{rdt} \geq u_{rd,t+1} - u_{rdt} \quad \forall\, r \in R, d \in D, t \in T \setminus \{\ell - 1, \ell, |T|\}; \tag{A.18}$$
$$v_{rdt} \geq u_{rdt} - u_{rd,t-1} \quad \forall\, r \in R, d \in D, t \in \{\ell - 1, |T|\}; \tag{A.19}$$
$$v_{rdt} \geq u_{rd,t-1} - u_{rdt} \quad \forall\, r \in R, d \in D, t \in \{\ell - 1, |T|\}; \tag{A.20}$$
$$u_{rd\ell} = 0 \quad \forall\, r \in R, d \in D; \tag{A.21}$$
$$u_{rdt} \geq 0, \text{integer} \quad \forall\, r \in R, d \in D, t \in T; \tag{A.22}$$
$$z_r \geq 0, \text{integer} \quad \forall\, r \in R; \tag{A.23}$$
$$v_{rdt} \geq 0, \text{integer} \quad \forall\, r \in R, d \in D, t \in T; \tag{A.24}$$
$$\rho_{\min}, \rho_{\max}, \eta_{\min}, \eta_{\max} \geq 0, \text{integer}. \tag{A.25}$$

The objective function (A.9) minimizes the sum of four penalties. The first two penalize solutions have an uneven distribution of technicians among the time slots of each day for, respectively, large and small clusters. The third one penalizes uneven distribution of technicians among consecutive time slots, and the last one aims at minimizing the number of working hours that could not be assigned to the available technicians. Parameters $\pi$, $\gamma$, and $\Omega$ represent the weights associated with each penalty. Constraints (A.10) ensure that the number of resources assigned to each cluster does not exceed its demand, whereas (A.11) imposes an upper bound to the number of technicians that could not be assigned to fulfill the demand of a certain cluster. Thus, constant $\lambda_1$ specifies the minimum ratio of the demand that must be fulfilled. Constraints (A.12) ensure that the capacity of each depot is not exceeded, whereas (A.13)–(A.16) determine the minimum and maximum number of resources assigned per day. Constraints (A.17)–(A.20) evaluate how balanced is the distribution of technicians among the time slots of a day by connecting $u$ and $v$ variables. For an in-depth explanation of these constraints, we refer the reader to Bruck et al. (2018). Finally, constraints (A.21) impose that all time slots associated with the lunch break (time slot $l \in T$) are unavailable for services, and (A.22)–(A.25) require integrality of the variables.

Although the aforementioned model is able to generate complete solutions for the problem, it does not consider some key QoS elements that are important for the company. To improve the level of QoS offered in the solutions designed by the model, we introduce some additional variables and constraints. Let $y_{rd}$ be a binary variable that specifies whether there is at least one time slot opened in day $d \in D$ for cluster $r \in R$. Note that, a time slot $(d, t)$ of cluster $r$ is considered opened if $u_{rdt} > 0$. A higher QoS level is then imposed by including the following constraints.

$$\sum_{t \in T} u_{rdt} \geq y_{rd} \quad \forall\, r \in R, d \in D; \tag{A.26}$$
$$\sum_{l \in \phi(d,g)} y_{rl} \geq 1 \quad \forall\, r \in R, d \in D; \tag{A.27}$$
$$u_{rdt} \leq y_{rd}\, Q_i \quad \forall\, r \in R_i, d \in D, t \in T; \tag{A.28}$$
$$\left| \sum_{t=1}^{\ell-1} u_{rdt} - \sum_{t=\ell+1}^{|T|} u_{rdt} \right| \leq \lambda_2 \sum_{t \in T} u_{rdt} \quad \forall\, d \in D, r \in R_0; \tag{A.29}$$
$$\left| \sum_{d \in D} \sum_{t=1}^{\ell-1} u_{rdt} - \sum_{d \in D} \sum_{t=\ell+1}^{|T|} u_{rdt} \right| \leq \lambda_2 \sum_{d \in D} \sum_{t \in T} u_{rdt} \quad \forall\, r \in R \setminus R_0; \tag{A.30}$$
$$y_{rd} \in \{0,1\} \quad \forall\, r \in R, d \in D. \tag{A.31}$$

As in Bruck et al. (2018), constraints (A.26)–(A.28) impose a limit $g$ on the maximum number of consecutive days without any time slot opened in a cluster, which is particularly important for clusters with low demand. In (A.27), function $\phi(d, g) = \{(i \bmod |D|) + 1 : i = d - 1, d, \ldots, d + g - 1\}$ is used to switch indices from one week to the next one. Constraints (A.29) and (A.30) state that there must be a certain balance in the distribution of technicians between time slots in the morning and in the afternoon, by forcing the difference in number of assigned technicians to be smaller than or equal to $(\lambda_2 \times 100)\%$ of the total number of assigned technicians. In our case, the company specified that larger clusters (i.e., any cluster $r \in R_0$) should have even higher QoS levels, which explains why constraints (A.29) are defined for each day and cluster, whereas constraints (A.30), in contrast, are only imposed for each cluster. In our experiments, we observed that good results were obtained by setting $\lambda_2 = 0.15$.

One of the main goals of formulation MWGen is to repair solutions in the iterations of an LNS procedure (see Bruck et al. 2018). To this aim, it is important to avoid finding very similar solutions or even the same one over and over again. Let $s \in S$ represent a feasible solution found at a given iteration, with $S$ being the set of solutions explored. Let also $w_{rdt}$ be an additional binary variable that specifies whether time slot $(d, t)$ from cluster $r$ is opened. Let $\Upsilon = |R||D||T|$ be the total number of $w_{rdt}$ variables, and, in addition, let $W_s^0$ and $W_s^1$ denote the sets of $w_{rdt}$ variables that take value 0 and 1, respectively, in solution $s$, for $s \in S$. The following set of constraints can be used to remove solutions from the search space that are very similar to any solution $s \in S$.

$$u_{rdt} \leq w_{rdt}\, Q_i \quad \forall\, r \in R_i, d \in D, t \in T; \tag{A.32}$$
$$u_{rdt} \geq w_{rdt} \quad \forall\, r \in R, d \in D, t \in T; \tag{A.33}$$
$$\sum_{(r,d,t) \in W_s^0} (1 - w_{rdt}) + \sum_{(r,d,t) \in W_s^1} w_{rdt} \leq \Upsilon - 1 \quad \forall\, s \in S; \tag{A.34}$$
$$w_{rdt} \in \{0,1\} \quad \forall\, r \in R, d \in D, t \in T. \tag{A.35}$$

Constraints (A.32), (A.33), and (A.35) define the $w$ variables and link them with the $u$ variables, whereas constraints (A.34) are *no-good cuts* that ensure that the new solution found by formulation MWGen will have a different configuration of opened time slots with respect to the previous solutions found during the LNS search.

## References

Agatz N, Fleischmann M, van Nunen JA (2008a) E-fulfillment and multi-channel distribution—A review. *Eur. J. Oper. Res.* 187(2):339–356.

Agatz N, Campbell AM, Fleischmann M, Savelsbergh M (2008b) Challenges and opportunities in attended home delivery. Golden BL, Raghavan S, Wasil EA, eds. *The Vehicle Routing Problem: Latest Advances and New Challenges* (Springer US, Boston), 379–396.

Agatz N, Campbell AM, Fleischmann M, Savelsbergh M (2011) Time slot management in attended home delivery. *Transportation Sci.* 45(3):435–449.

ARERA (2013) Single text of the instructions of quality regulation and tariffs of gas distribution and measurement services for the 2014–2019 regulatory period. Accessed December 12, https://www.autorita.energia.it/allegati/docs/13/574-13all.pdf.

ARERA (2016) Integrated text of the authority regulations for accounting unbundling for companies operating in the same gas, water or electricity business market. Accessed March 24, https://www.arera.it/allegati/docs/16/137-16all.pdf.

Baldacci R, Mingozzi A, Roberti R (2012) Recent exact algorithms for solving the vehicle routing problem under capacity and time window constraints. *Eur. J. Oper. Res.* 218(1):1–6.
Bent RW, Van Hentenryck P (2004) Scenario-based planning for partially dynamic vehicle routing with stochastic customers. *Oper. Res.* 52(6):977–987.
Bruck BP, Cordeau JF, Iori M (2018) A practical time slot management and routing problem for attended home services. *Omega* 81:208–219.
Campbell AM, Savelsbergh M (2005) Decision support for consumer direct grocery initiatives. *Transportation Sci.* 39(3):313–327.
Campbell AM, Savelsbergh M (2006) Incentive schemes for attended home delivery services. *Transportation Sci.* 40(3):327–341.
Cappanera P, Scutellà MG, Nervi F, Galli L (2018) Demand uncertainty in robust home care optimization. *Omega* 80:95–110.
Ehmke JF, Mattfeld DC (2012) Vehicle routing for attended home delivery in city logistics. *Procedia Soc. Behav. Sci.* 39:622–632.
European Parliament (2009) Directive 2009/73/EC of the European parliament and of the council. Accessed July 13, http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2009:211:0094:0136:en:PDF.
Fadda E, Gobbato L, Perboli G, Rosano M, Tadei R (2018) Waste collection in urban areas: A case study. *Interfaces* 48(4):307–322.
Han S, Zhao L, Chen K, Luo Zw, Mishra D (2017) Appointment scheduling and routing optimization of attended home delivery system with random customer behavior. *Eur. J. Oper. Res.* 262(3):966–980.
Hanne T, Melo T, Nickel S (2009) Bringing robustness to patient flow management through optimized patient transports in hospitals. *Interfaces* 39(3):241–255.
Hernandez F, Gendreau M, Potvin JY (2017) Heuristics for tactical time slot management: A periodic vehicle routing problem view. *Internat. Transportation Oper. Res.* 24(6):1233–1252.
Kostuk KJ, Willoughby KA (2012) A decision support system for scheduling the Canadian Football League. *Interfaces* 42(3):286–295.
Lin II, Mahmassani HS (2002) Can online grocers deliver? Some logistics considerations. *Transportation Res. Record J. Transportation Res. Board* 1817(1):17–24.
Madsen OB, Tosti K, Vælds J (1995) A heuristic method for dispatching repair men. *Ann. Oper. Res.* 61(1):213–226.
Murtagh F, Contreras P (2012) Algorithms for hierarchical clustering: An overview. *Wires Data Mining Knowledge Discovery* 2(1):86–97.
Pasupathy KS, Medina-Borja A (2008) Integrating Excel, Access, and Visual Basic to deploy performance measurement and evaluation at the American Red Cross. *Interfaces* 38(4):324–337.
Punakivi M, Saranen J (2001) Identifying the success factors in e-grocery home delivery. *Internat. J. Retail Distribution Management* 29(4):156–163.
Spliet R, Desaulniers G (2015) The discrete time window assignment vehicle routing problem. *Eur. J. Oper. Res.* 244(2):379–391.
Spliet R, Gabor A (2014) The time window assignment vehicle routing problem. *Transportation Sci.* 49(4):721–731.
Subramanyam A, Wang A, Gounaris C (2018) A scenario decomposition algorithm for strategic time window assignment vehicle routing problems. *Transportation Res. Part B: Methodological* 117:296–317.
Toth P, Vigo D, eds. (2014) *Vehicle Routing: Problems, Methods, and Applications,* 2nd ed. (SIAM, Philadelphia).
Yang X, Strauss AK, Currie CS, Eglese R (2014) Choice-based demand management and vehicle routing in e-fulfillment. *Transportation Sci.* 50(2):473–488.

## Verification Letter

Rosario Di Bartolo, Metering and Distribution Services Director, IRETI, 16138 Genova, Italy, writes:

"IRETI S.p.A., in particular its Metering Operations branch, and the Operations Research group of the University of Modena and Reggio Emilia have jointly developed the Decision Support System (DSS) described in detail in the article 'A Decision Support System for Attended Home Services.'

"The DSS has been formulated and implemented to help IRETI operations managers and their staff in the decision-making process to be performed each time a specific public tender, issued by the Italian government to select a gas distribution company in a minimum territorial area, is published.

"With the present letter, we declare that IRETI has used and is using the DSS described in the article. In particular, the DSS is in use for a number of public tenders that are expected to be published soon by the Italian government to select gas distribution companies."

**Bruno P. Bruck** is an adjunct professor at the Federal University of Paraíba (Brazil). His main research area is combinatorial optimization and logistics, with a focus on traveling salesman, vehicle routing, and bike sharing problems. He also has experience in solving real-world optimization problems involving vehicle routing, intermodal railway transportation, and the development of decision support systems.

**Filippo Castegini** is currently employed as a software analyst for the IT automation division of an international Italian company. He obtained his BS in industrial engineering at the University of Padova and his MS in industrial engineering at the University of Modena and Reggio Emilia.

**Jean-François Cordeau** is a professor of operations management at HEC Montréal, where he also holds the Chair in logistics and transportation. He has authored or coauthored more than 140 scientific articles in combinatorial optimization and mathematical programming, mostly in the fields of vehicle routing and logistics network design. He has also acted as a consultant for several Canadian and European organizations in the private and public sectors.

**Manuel Iori** is an associate professor of Operations Research at the University of Modena and Reggio Emilia. He has published more than 70 papers in international journals in the areas of combinatorial optimization and logistics, with a focus on traveling salesman, vehicle routing, bin packing, and multidimensional cutting and packing problems. He collaborates with companies for the solutions of real-world optimization problems, by designing and developing decision support systems.

**Tommaso Poncemi** obtained his MS in mechanical engineering in 2000 at the University of Parma. Since 2004, he has been working at the IREN Group. He is now employed at IRETI, a subsidiary of the IREN Group, as head of the external and metering operations unit, where he directs activities related to gas and water distribution networks.

**Dario Vezzali** is a PhD candidate in labour, development, and innovation at the University of Modena and Reggio Emilia. Formerly a professional in the field of operations management, he is now approaching the study and development of decision support systems to solve real-world optimization problems. His focus is mainly on vehicle routing, time slot management, engineering economics, and multicriteria decision analysis.