22nd National Conference of the Australian Society for Operations Research, Adelaide, Australia, 1–6 December 2013
www.asor.org.au/conferences/asor2013

Constructing Water Tank Delivery Schedules through Combined Vehicle Routing and Packing Decisions

Jacob Stolk[a,c], **Isaac Mann**[a], **Arvind Mohais**[a,d], **Zbigniew Michalewicz**[b,a,e]

[a) SolveIT Software / Schneider Electric, Level 1, 99 Frome Street, Adelaide, SA 5000, Australia.
[b) School of Computer Science, University of Adelaide, Adelaide, SA 5005, Australia; also at the Institute of Computer Science, Polish Academy of Sciences, ul. Jana Kazimierza 5, 01-248 Warsaw, Poland. Polish-Japanese Institute of Information Technology, ul. Koszykowa 86, 02-008 Warsaw, Poland.
[c) Email: js@solveitsoftware.com (corresponding author).
[d) Email: am@solveitsoftware.com.
[e) Email: zbyszek@cs.adelaide.edu.au.

Abstract: This paper describes a decision-support system that was developed in 2011 and is currently in production use. The purpose of the system is to assist planners in constructing delivery schedules of water tanks to often remote areas in Australia. A delivery schedule consists of a number of delivery trips by trucks. An optimal delivery schedule minimises cost to deliver a given total sales value of delivered products. To construct an optimal delivery schedule, trucks need to be optimally packed with water tanks and accessories to be delivered to a set of delivery locations. This packing problem, which involves many packing and loading constraints, is intertwined with the transport problem of minimising distance travelled by road.

We have reviewed relevant literature on vehicle routing problems and on packing and cutting problems. Many of the problems studies in the literature have some similarities with the problem we had to solve, but none of these consider the same business rules and constraints we had to consider.

Therefore, we have combined several algorithms in a software system to assist planners of a water tank production company to construct good delivery schedules of ordered water tanks on an on-going basis. Algorithms used include: a clustering algorithm to cluster delivery locations and agent locations as a basis for constructing good delivery trips, i.e. trips with an acceptable transport cost / delivered value ratio; an adapted vehicle routing algorithm to calculate approximately optimal (shortest) delivery trips around a given intermediate storage/unbundling location; an algorithm to calculate shortest distances between geographical locations in a road network; a custom algorithm to load trucks with water tanks according to bundling possibilities; a 2D packing algorithm to load products that cannot be bundled on remaining open space on trucks and trailers; algorithms for constructing vertical and pyramid stacks of products, where products can be stacked. In addition, we have used a geolocation service to obtain geographical coordinates corresponding to delivery and agent addresses.

Each of the algorithms used solves a well-defined problem which is a part of the business problem to be solved. The main challenge was that the complete business problem is not well-defined, as business rules and constraints that have to be taken into account are often implicit; for example, a solution that might be optimal in a static sense, for delivering a given set of orders, is not necessarily optimal in a dynamic sense, for optimising business operations on an on-going basis. To obtain good solutions, not only does each of the algorithms used need to perform well, but their partial results need to be combined to produce acceptable overall solutions.

Keywords: *Optimisation, clustering, vehicle routing, packing, shortest path.*

1 INTRODUCTION

Large-scale business decision problems consist of interconnected components, but need to be solved to achieve optimal results for a business as a whole. Even if we know exact algorithms for solving sub-problems, it remains an open question how to integrate these partial solutions to obtain a global optimum for the whole problem. Businesses need global solutions for their operations, not partial solutions. This was recognised over 30 years ago in the Operations Research (OR) community, when Russell Ackoff wrote: "… *problems are abstracted from systems of problems, messes. Messes require holistic treatment. They cannot be treated effectively by decomposing them analytically into separate problems to which optimal solutions are sought*" (Ackoff, 1979). These remarks are directly applicable to planning of optimal delivery schedules, for which intertwined routing, packing and assignment problems need to be solved, that should not be considered in isolation.

Every time we solve a problem we must realise that we are in reality only finding the solution to a *model* of the problem. All models are a simplification of the real world – otherwise they would be as complex and unwieldy as the natural setting itself. Thus the process of problem solving consists of two steps: (1) creating a model of the problem, and (2) using that model to generate a solution (Michalewicz and Fogel, 2004):

$$\textit{Problem} \rightarrow \textit{Model} \rightarrow \textit{Solution}$$

The system described in this article embodies a model of a delivery schedule planning problem that is as accurate as possible. It combines various algorithms and heuristics to find approximate solutions for a complex real-world business problem and is currently in production use. The purpose of the system is to assist planners in constructing delivery schedules of water tanks to often remote areas in Australia. A delivery schedule consists of a number of delivery trips by trucks. An optimal delivery schedule minimises cost to deliver a given value. To construct an optimal delivery schedule, trucks need to be optimally packed with water tanks and accessories to be delivered to a set of delivery locations. This packing problem, which involves many packing and loading constraints, is intertwined with the transport problem of minimising distance travelled by road.

2 THE WATER TANK TRANSPORT AND PACKING PROBLEM

An Australian company produces, sells and delivers rain water tanks. These products with a large volume should be delivered to geographically dispersed customers, so transport cost is high and transport cost minimisation an important competitive advantage. The transport cost is considered with respect to the value of products delivered during the trips composing a delivery schedule, so the packing component is very important to maximise delivered value for a given transport cost. In practice the objective is to deliver all orders, so the value is given and the problem is reduced to finding a schedule with the lowest possible transport cost. Assuming a fixed transport cost per kilometre, the problem is further reduced to finding a schedule with minimum distance travelled.

The company's delivery schedule planners need to make decisions involving numerous variables related to production, intermediate storage, packing, transport and delivery of over 3000 products, divided in over 100 categories with different dimensions of rain water tanks and accessories, to fulfil on the order of 1000 to 2000 customer orders per month, while respecting numerous business rules and constraints that impact distribution decisions and are often implicit. Orders come in continuously, so today's optimal solution is not necessarily still optimal tomorrow and re-optimisation is frequently necessary. Decisions to be made in this complex logistics operation include:

- the choice of one of several plants to make each product to be delivered and dozens of storage facilities (agent sites, hubs, unbundling locations) from which products in stock can be shipped;
- selection of trucks for delivery trips from several types of trucks in a fleet of some 50 trucks, with or without trailers, to transport the goods, taking into account constraints such as truck availability at each plant, truck and trailer dimensions, etc.;
- packing of goods on trucks and trailers in an optimal way, taking into account numerous constraints related to truck, trailer and product dimensions, as well as several ways to bundle, pack and stack products according to product type;
- selection of drivers from over 100 possible drivers, taking into account their availability, including timetables, maximum hours of work, availability for making longer trips, qualifications (not all drivers can drive all trucks or transport all tanks);
- selection of optimal routes to travel in order to minimise distance travelled.

All these decisions need to be made to minimise the key performance indicator of transport cost as percentage of delivered sales value. Each of the above problems is hard to solve in isolation.

Further, the problems are connected, as decisions made for one problem may impact decisions for another problem. Thus, packing and routing problems are intertwined, as the destination locations of items packed on a truck/trailer for a trip determine the final delivery destinations to be visited. Water tanks can often be bundled inside each other, making possible efficient packing of a truck and/or trailer, but they can be unbundled only at specific agent locations of the company, which has to be done prior to final delivery to customer locations – so there is a trade-off between efficient packing and the need to travel a longer distance to visit unbundling sites in addition to final delivery sites.

Another important consideration is that the business problem to be solved is not limited to finding a delivery schedule for a static set of orders, but rather finding the best delivery schedules for orders that evolve over time. For example, rather than filling up a delivery trip with as many orders as possible from the current set of orders, causing it to visit a large geographical area, it is usually preferred to only partially load the truck with orders for a smaller area and delay its departure until new orders come in. Therefore, a heuristic solution for the current set of orders that takes into account such considerations may well be better in the long run than a strictly optimal solution for the current orders. Our scheduler is designed to achieve such a heuristic solution.

The overall water tank delivery problem has some similarities with other well-known optimisation problems described in the literature, reviewed in subsections 2.1 – 2.2, but also has peculiar features which are not found elsewhere. We have developed algorithms to solve the problem by adapting and combining selected optimisation algorithms for problems with similar characteristics.

2.1 Vehicle routing problems

A vehicle routing problem (VRP) consists of determining a set of vehicle trips from a depot to customers, of minimum total cost, such that each trip starts and ends at the depot, each client is visited exactly once, and the total demand handled by any vehicle does not exceed the vehicle capacity (Clarke and Wright, 1964; Cordeau et al, 2007a). Some variants of the VRP that have been studied include:

- the classical vehicle routing problem: a commodity is to be delivered at minimal cost from a depot to a set of customers, given a nonnegative demand for each customer, travel costs for each route, and a fixed capacity for each vehicle of a fleet of identical vehicles;
- VRPs with loading constraints, reviewed in (Iori and Martello, 2010): loading constraints that have been analysed include two and three-dimensional loading constraints, as well as some special variants such as multi-pile loading, taking into account order of loading, etc.;
- a multi-depot VRP: delivery trips take place to several delivery regions, each with a hub where unbundling is done; in the multi-depot VRP described in (Ho et al, 2008) it is assumed that each depot has enough stock to supply the demand of all customers;;
- pickup and delivery problems: VRPs where a set of transportation requests is satisfied by a given fleet of vehicles (Cordeau et al, 2007b); each request is characterised by its pickup location (origin), its delivery location (destination) and the size of the load that has to be transported from the origin to the destination.

For the water tank delivery problem an important limitation of the formulation of VRPs is the assumption that only one commodity is transported and both vehicle capacities and customer demand are expressed as quantities of this single commodity, while water tanks have many different sizes and shapes, trucks (with or without trailers) have different dimensions and drivers have different qualifications for driving trucks and transporting tanks depending on dimensions; thus the possibility of transporting a given load of tanks is determined by all these variables rather than a simple capacity expressed as a single number.

2.2 Cutting and packing problems

In cutting and packing problems a set of large objects and a set of small items are given. The problem is to select some or all small items, group them into one or more subsets and assign each of the resulting subsets to one of the large objects such that the small items of the subset lie entirely within the large object and do not overlap, and a given objective function is optimised (Wäscher et al, 2007).

The problem of packing a set of water tanks on available trucks is a kind of cutting and packing problem. However, the delivery cost depends on the route to be travelled by the truck on which the items are packed which in turn is not independent of items packed on and routes to be travelled by other trucks.

3 THE WATER TANK TRANSPORT AND PACKING SCHEDULER

We have combined several algorithms in a decision support system to assist planners to make numerous packing and routing decisions for the construction of good delivery schedules of water tanks and accessories on an on-going basis. We have made use of these algorithms to construct heuristics for finding delivery schedules with desirable characteristics, as detailed below.

The system generates a delivery schedule using the following steps:

1) Delivery destinations and intermediary agent locations:
    1.1) retrieve geolocations using a geolocation method to determine geographical coordinates from address data, and road network data from a dataset provided by Geoscience Australia (http://www.ga.gov.au)
    1.2) calculate distances between locations to be visited with Dijkstra shortest path algorithm (Dijkstra, 1959)
2) Initial clustering (an example of clusters is shown in Figure 1). Clustering of delivery locations is desirable, as delivery trips normally leave a plant to travel to a central location for a group of delivery locations, which should preferably be close together. We use a clustering algorithm that determines a clustering for any set of locations without assumptions about the number of locations, as the set of locations to be visited perpetually changes according to customer orders to be fulfilled and/or agent locations being activated or deactivated. Another desirable feature of a clustering algorithm is the possibility to fine-tune the algorithm by setting parameters to be able to control, for example, the maximum size of generated clusters, or the approximate number of locations in a cluster, so clusters can be adjusted depending on numbers of deliveries to be made in a region. A clustering algorithm with such features is the fuzzy clustering by Local Approximation of Membership (FLAME) algorithm (Fu and Medico, 2007), originally developed for bioinformatics applications:
    2.1) cluster destination locations with FLAME clustering algorithm:
        2.1.1) for each object (delivery location), find the k nearest neighbours, calculate their proximity, calculate object density and define the object type (CSO/outlier/other)
        2.1.2) assign fuzzy membership by local approximation
        2.1.3) construct clusters from the fuzzy memberships
    2.2) post-process clusters to avoid clusters that are too large: a maximum cluster size is specified - if the largest distance between two locations in a cluster exceeds the maximum size, the cluster is split into two clusters by allocating each location in the cluster to a new cluster according to distance to each of the two locations with maximum distance
    2.3) if desired, post-process clusters to avoid clusters that are too sparse: if the average distance between cluster locations is too large, split the cluster in two - users may prefer trips that are not completely loaded but are confined to a relatively small region, rather than better loaded trips covering a larger region, when they expect the incompletely loaded trips to be filled with new orders to come in later
3) Construct trips: for each cluster, select a truck and a driver, then attempt to add items to be delivered to a trip in order of decreasing distance of delivery locations to the plant, checking if items can be bundled, packed or stacked:
    3.1) select truck and driver, taking into account constraints such as possibility to transport certain types of tanks, driver availability for the duration of the trip, etc.
    3.2) bundling: a recursive method of cutting the top of the tank off, placing a smaller tank on the inside, and then repeating the process (Figure 2) - bundling can only be performed at plants, and can only be undone at specifically chosen sites that contain the equipment and expertise to unbundle
    3.3) packing and stacking: attempt to pack and/or stack an item, using custom packing algorithms: a 2D column packing algorithm, combined with “stacking” (placing items of the same product group on top of each other) and/or “pyramid stacking” (placing tanks of the same product group in a pyramid, e.g. 3 in the base layer, 2 in the next layer, 1 in the top layer) - items are checked for merging possibilities with stacking or pyramid stacking and placed into rectangular objects representing their final size; these rectangles can then be packed
4) Combine clusters of remaining destinations , if not all items have been assigned to trips in step 3) and if this is desired by the user; clusters are combined if they are next to each other and each have only a few

remaining delivery locations; as in step 2.3, users may decide to skip this step for similar reasons: not recombining clusters will tend to produce less well loaded trips confined to smaller delivery regions

5) Iterate 3), in combination with 4) if desired, until all ordered products are assigned to trips
6) For each trip, find the best unbundling location, i.e. the location that, if chosen, minimises trip distance travelled:
    6.1) calculate maximum and minimum latitude and longitude for all delivery locations of the trip
    6.2) for each possible unbundling location within the area contained in these bounds, construct trip routes with the algorithm described in step 7) – retain the unbundling location with the shortest trip distance
7) For each unbundling location, determine the best, i.e. shortest distance, routing of ‘hub run’ trips from unbundling location to customers with a modified Clark & Wright algorithm (Figure 3):
    7.1) initial solution: each vehicle serves exactly one customer
    7.2) for each pair two distinct routes, compute possible savings by merging them, for example: merging routes servicing customers i and j leads to savings $s_{i,j} = c_{i,0} + c_{0,j} - c_{i,j}$; if $s_{i,j} > 0$, the merging operation is convenient
    7.3) all saving values $s_{i,j}(i, j = 1, \ldots, n \; and \; j > i)$ are stored in a half-square matrix M
    7.4) matrix M is sorted in not-increasing order of the $s_{i,j}$ values to create a list L of saving objects composed by the triplets $(s_{i,j}, i, j)$ - the higher the saving value, the more appealing the associated merge operation
    7.5) the saving objects in list L are now sequentially considered: if the associated merge operations are feasible, they are carried out, where merge feasibility is determined as follows:
        7.5.1) overload of the vehicle: a merge operation is not feasible, if load to be transported violates vehicle capacity – in the original Clark & Wright algorithm vehicle capacity is a given quantity of goods; we have modifies this by a check if the goods can be packed on the truck, given the loading constraints
        7.5.2) internal customers: a customer which is neither the first nor the last at a route cannot be involved in merge operations
        7.5.3) customers both in the same route: if customers i and j suggested by saving $(s_{i,j}, i, j)$ are at extremes of the same route (first or last), the merge operation cannot be performed (no sub-tours are allowed)
    7.6) a solution is found, when no more merge operations are possible
8) The best set of delivery trips is retained as recommended delivery schedule.

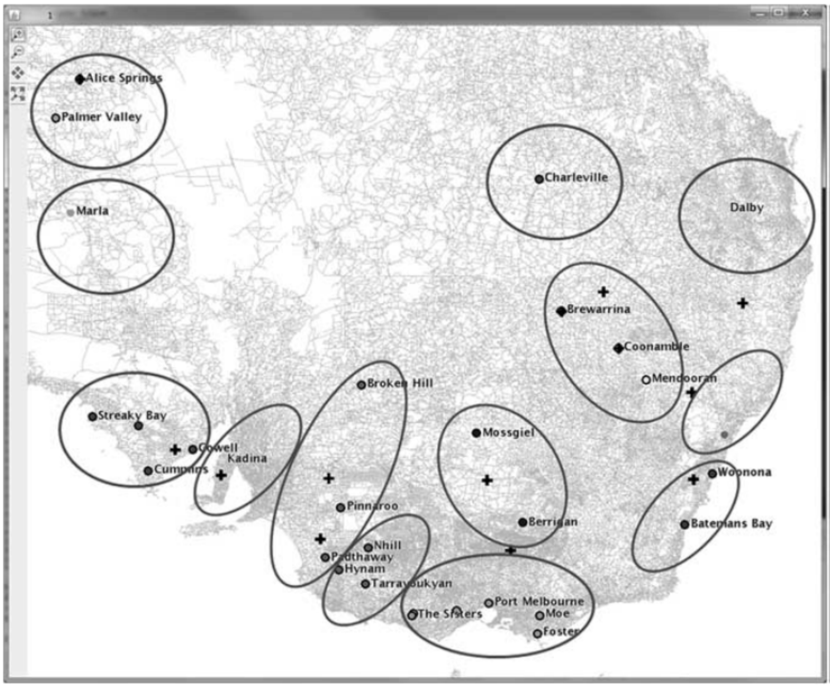

Figure 1. Clustered locations.

The coloured dots represent delivery destinations, with a different colour for each cluster (encircled) that has been found by the clustering algorithm. Crosses represent agent sites that can be used for unbundling.

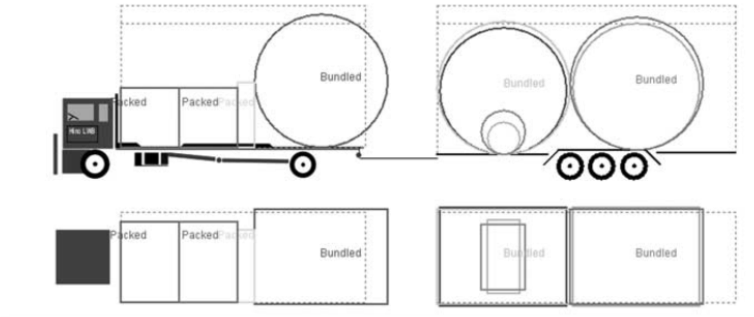

Truck: 1 bundles (1 bundled items) 3 packed items; Trailer: 2 bundles (6 bundled items)

Tanks

No.	Transport	Tank	Pyramid	Order No.	Product	Product Group	Product	Height	Width (m)	Depth (m)	Weight	Unbundling Location
1	Truck	Packed		480921	TP1112	3000L SEPT...	RECTA...	1.635	2.214	1.624	240	
2	Truck	Packed		480921	TP1112	3000L SEPT...	RECTA...	1.635	2.214	1.624	240	
3	Truck	Packed		480528	TP7497	3 MODULE S...	RECTA...	1.994	1.793	0.475	72	
4	Truck	Bundle		478508	TP11173	22500L SQU...	ROUND	2.631	3.7	3.7	378.05	
5	Trailer	Bundle		478508	TP11173	22500L SQU...	ROUND	2.631	3.7	3.7	378.05	
6	Trailer	Bundle		479239	TP3307	22500L TANK	ROUND	2.72	3.53	3.53	378.25	
7	Trailer	Bundle		479861	TP1952	1600L TANK	ROUND	1.809	1.23	1.23	50	
8	Trailer	Bundle		479861	TP2348	1000L TANK	ROUND	2.021	0.897	0.897	37	
9	Trailer	Bundle		478508	TP11173	22500L SQU...	ROUND	2.631	3.7	3.7	378.05	
10	Trailer	Bundle		479239	TP3307	22500L TANK	ROUND	2.72	3.53	3.53	378.25	

Figure 2. A loaded trip. The top part of this screen is a graphical view of a truck and trailer load with bundled and packed items, seen from the side and from above. The table at the bottom gives details on the items loaded: if they are loaded on the truck or the trailer, if they are packed or bundled, as well as product code, product group, shape, dimensions, weight and unbundling location.

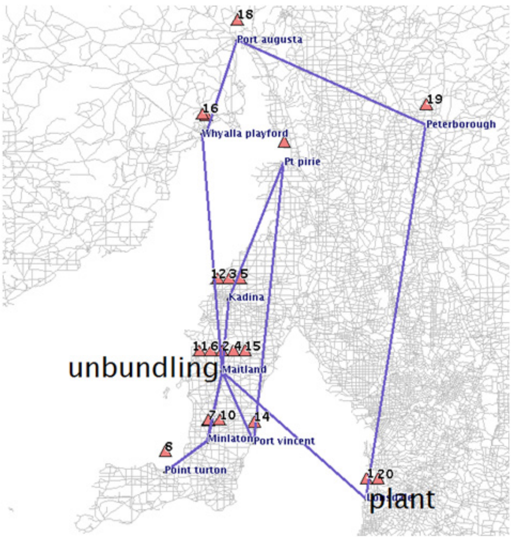

Figure 3. A trip routing. This trip leaves the central plant at Lonsdale and first goes to an agent site in Maitland, where unbundling is done. From the unbundling site hub run trips are made to the delivery locations. Hub run trips return to the unbundling location to pick up items for the next hub run, except the last hub run trip, which goes back to the plant. Locations of this trip are numbered from 1 to 20 in order of visit. Routes are shown as straight lines for rendering performance reasons only. Actual distances travelled and used by the algorithms are distances by road.

4 CONCLUSION

The water tank packing and transport scheduler produces delivery schedules with details of trip routes and packing. The scheduler is used in production by the delivery trip planners of the company it was built for. Key benefits obtained are: more profitable trips as the scheduler locates opportunities to maximise trip loads and minimise transport cost; fewer trips cancelled owing to trips not passing verifier criteria – fewer dissatisfied customers (end users and resellers) and improved service levels; more proactive order fulfilment, with reduced delivery times after order placement; reduced routine work for planning delivery schedules, so more time is available for customer service and sales work, resulting in improved customer relationship management; removal of duplication of information, voluminous paperwork, copying and storage of hard copies; reduction of carbon footprint across the business; greater visibility of the planning process and a single point of reference for information; more flexible automated reporting for management.

We have combined several algorithms in a software system to assist planners of a water tank production company to construct good delivery schedules of ordered water tanks on an on-going basis. Algorithms used include: a clustering algorithm to cluster delivery locations and agent locations as a basis for constructing good delivery trips, i.e. trips with an acceptable transport cost; an adapted vehicle routing algorithm to calculate approximately optimal (shortest) delivery trips around a given intermediate storage/unbundling location; an algorithm to calculate shortest distances between geographical locations in a road network; a custom algorithm to load trucks with water tanks according to bundling possibilities; a 2D packing algorithm to load products that cannot be bundled on remaining open space on trucks and trailers; algorithms for constructing vertical and pyramid stacks of products, where products can be stacked. In addition, we have used a geolocation service to obtain geographical coordinates corresponding to delivery and agent addresses.

Each of the algorithms used solves a well-defined problem which is a part of the business problem to be solved. The main challenge was that the complete business problem is not well-defined, as business rules and constraints to be taken into account are often implicit; for example a solution that might be optimal in a static sense, for delivering a given set of orders, is not necessarily optimal in a dynamic sense, for optimising business operations on an on-going basis. To obtain good solutions, not only does each of the algorithms used need to perform well, but their partial results need to be combined to produce acceptable overall solutions. For example, the initial clustering of delivery locations is not necessarily the best one, as it does not consider truck and driver selection or loading constraints. We have taken this into account using an iterative approach, re-clustering remaining delivery destinations when all initial destinations in a cluster cannot be served by a constructed trip. We have chosen a clustering algorithm which does not make assumptions about the number of clusters to be generated and is efficient. Similarly, once deliveries are assigned to a trip, the route of the trip is determined by evaluating routes for a set of possible unbundling locations and selecting the shortest one. This involves numerous executions of the shortest route algorithm. For this reason we have adapted the Clark & Wright algorithm, which gives only approximate solutions, but is very efficient.

ACKNOWLEDGEMENTS

The authors express their gratefulness to the following persons who have contributed to the development of the software described in this article: software developers Jingjing Dong, Alexis Pflaum and Ayaz Hassan; business analysts Josh Boston and Daniel Spitty; scientist Mohammadreza Bonyadi, various testers at SolveIT Software / Schneider Electric.

REFERENCES

Ackoff R (1979). The future of operational research is past. *Journal of the Operational Research Society* 30(2): 93-104.

Clarke G and Wright J W (1964). Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research* 12(4): 568-581.

Cordeau J-F, Laporte G, Potvin J-Y and Savelsbergh M W P (2007b). Transportation on demand. In: Barnhart C and Laporte G (eds). *Transportation, Handbooks in Operations Research and Management Science* 14: 429–466. Elsevier, Amsterdam.

Cordeau J-F, Laporte G, Savelsbergh M W P and Vigo D (2007a). Vehicle routing. In: Barnhart C and Laporte G (eds). *Transportation, Handbooks in Operations Research and Management Science* 14: 367–428. Elsevier, Amsterdam.

Dijkstra E W (1959). A note on two problems in connexion with graphs. *Numerische Mathematik* 1: 269–271

Fu L and Medico E (2007). FLAME, a novel fuzzy clustering method for the analysis of DNA microarray data. *BMC Bioinformatics* 8(3): http://www.biomedcentral.com/1471-2105/8/3.

Ho W, Ho G T S, Ji P and Lau H C W (2008). A hybrid genetic algorithm for the multi-depot vehicle routing problem. *Engineering Applications of Artificial Intelligence* 21(4): 548-557.

Iori M and Martello S (2010). Routing problems with loading constraints. *Top* 18: 4-27.

Michalewicz Z and Fogel D B (2004). *How to Solve It: Modern Heuristics*, 2nd edition. Springer-Verlag: Berlin, Heidelberg.

Wäscher G, Haußner H and Schumann H (2007). An improved typology of cutting and packing problems. *European Journal of Operational Research* 183: 1109–1130.