Journal of the Operational Research Society (2003) 54, 1209–1221 0160-5682/03 $25.00

www.palgrave-journals.com/jors

District design for arc-routing applications

L Muyldermans*, D Cattrysse and D Van Oudheusden

Centre for Industrial Management, Katholieke Universiteit Leuven, Celestijnenlaan 300A, B-3001 Heverlee, Belgium

In this paper we address the problem of district design for the organisation of arc-routing activities. In particular, the focus is on operations like winter gritting and road maintenance. The problem involves how to allocate the road network edges to a set of depots with given locations. The collection of edges assigned to a facility forms a district in which routes have to be designed that start and end at the facility. Apart from the ability to support good arc routing, well-designed districts for road-maintenance operations should have the road network to be serviced connected and should define clear geographical boundaries. We present three districting heuristics and evaluate the quality of the partitions by solving capacitated arc routing problems in the districts, and by comparing the solution values with a multi-depot CARP cutting plane lower bound. Our experiments reveal that based on global information about the distribution system (ie the number of facilities or districts, the average edge demand and the vehicle capacity) and by using simple guidelines, an adequate districting policy may be selected.

Journal of the Operational Research Society (2003) **54,** 1209–1221. doi:10.1057/palgrave.jors.2601626

Keywords: districting; capacitated arc routing; heuristics; winter gritting

Introduction

Districting involves the partitioning of a large geographical region (or network) into smaller subareas (subnetworks) for organisational and/or administrative purposes. The applications of district design are manifold and include political districting,[1] the design of territories for salesmen,[2] health-care districting,[3] school district design,[4] police districting,[5] etc. Nearly all districting approaches in the literature are of the agglomeration type.[6,7] This means that the region to be subdivided is pre-partitioned into a large number of small basic areas (or units) that are aggregated into districts afterwards. Often, the aggregation phase consists of specifying (sometimes fictitious) district centres, followed by allocating the units to the centres while an objective is optimised and subject to side constraints. Usually, these constraints ensure the contiguity, the 'compactness' and the workload balance of the districts, and prevent a unit from splitting among several territories.

In this paper, we address the problem of district design for the organisation of arc-routing activities; in particular, for operations like salt spreading and road maintenance. Most naturally, the routing of gritters[8,9] is modelled as a capacitated arc routing problem (CARP)[10,11] — a routing problem where the demand for service occurs along the edges of a network and where the vehicles have a finite capacity. The districting problem involves how to allocate the network edges to a set of facilities with given locations. The collection of edges assigned to a centre is called a district, and each facility is independently responsible for the organisation of the routing within its district borders.

The primary goal in this article is to identify a number of district design guidelines that have an important influence on the routing efficiency — the total deadheading distance and the number of vehicles used — of the CARPs to be organised within the district borders. These guidelines relate to the definition of suitable units to build districts as well as to the selection of the appropriate objective(s) to guide the aggregation process. Three districting heuristics, either making use in the aggregation phase of individual edges or of groupings of edges into small cycles, are presented and evaluated. The units are assigned to the nearest facility in a rather greedy manner, or through the solution of an integer linear programming model (minimising a lower bound on the number of vehicles to be used). In both cases, we also ensure that the road network to be serviced in each district is connected. We test the procedures on large graphs constructed from the road network in Flanders (Belgium) and evaluate the quality of the partitions by solving CARPs in the districts and by comparing the solution values with a multi-depot CARP cutting plane lower bound. Depending on the size of the vehicle capacity, different districting policies are recommended.

In the sections that follow, we first describe the role of districting in the planning of (winter) road maintenance activities. We then derive several guidelines for district design — among these an approach for pre-clustering edges into small cycles, which intend to account for the local

Correspondence: L Muyldermans, Centre for Industrial Management, Katholieke Universiteit Leuven, Celestijnenlaan 300A, B-3001 Heverlee, Belgium.
E-mail: Luc.Muyldermans@cib.kuleuven.ac.be

routing cost. Next, we present the three heuristics and report the results of our computational experiments. Some recommendations for improving the gritting operations in Flanders are given and conclusions are drawn.

District design for road-maintenance activities

The partitioning of the road network into districts is a real, distinct stage in the planning and the organisation of (winter) road-maintenance activities.[12] Typically, districting is performed after the location of the facilities has been determined and before the routes are fixed. From a planning point of view, the difference between district design and routing is fundamental: whereas routing is performed at the operational level, districting (like location) is of a non-operational nature and more related to the managerial level. Apart from being a frame for routing, districts also serve administrative purposes. It is at the depot of the district that, over many years, useful, inter-related data are collected and that the road surface condition is monitored so as to make pertinent decisions about where, when and how much salt to be spread. Since modifying the district boundaries usually results in a lot of operational and administrative adjustments, the borders should not be changed too frequently, but only when major improvements in carrying out the operations can be reached or when important changes in the distribution system and/or activities take place, such as closing a facility, the introduction of a new task or the construction of an important new road. Furthermore, possibly, the same districts should support different types of arc routing or facilitate the organisation of other operations. For instance, in Flanders, both preventive and curative interventions (requiring different amounts of salt) are carried out within the same districts. Hence, whereas routes are more sensitive to specific constraints (capacity, time, distance, etc), districts should be more robust and not influenced by minor changes in the operational characteristics. Therefore, different guidelines should be used for districting and routing, while keeping the interactions between the two levels in mind.

In the literature, the problem of district design for (arc) routing applications receives very limited attention or is assumed to be solved *a priori*. Often, however, long-term savings can be achieved if more careful attention is devoted to the drawing of the district boundaries.[13] One approach in the literature is to assume that the districts are mainly determined by the location of the facilities and that the partition is obtained by assigning the customers to the nearest depot. In a probabilistic node routing context with split deliveries in the Euclidean plane,[14] it is shown that the average distribution cost behaves asymptotically as the cost of the classical p-median problem, so that the nearest depot services each customer. Another approach is to perform multi-depot routing and define the districts by the routes emanating from each depot. However, procedures based on the first approach may over-emphasise the importance of the radial travel cost (ie the distance from the facility to a demand entity and back); while districts based on the second approach may be too sensitive to specific routing constraints and lack a more global view.

In contrast to the clear objectives in pure location or routing problems, it appears to be more difficult to define exact criteria for designing good districts for (arc) routing. Instead, a number of more intuitive characteristics of a good district partition are formulated. It is clear that districts should not be too small in order to keep the fixed costs reasonable (depot infrastructure, administrative staff, etc) and on the other hand, they should not be too large since otherwise the organisation of the services becomes too complex. Since the routes in a district all start and end at the distribution centre, odd-shaped districts with the depot located near a boundary result probably in long, inefficient tours. A better partition would have the demand entities within each district near to each other and near the service centre. Thus, good districts will very likely be compact in shape and have centrally located facilities. In order to facilitate the planning and to delineate the responsibilities between neighbouring districts, districts for salt-spreading operations should define clear geographical borders (no overlap), have the road network to be serviced connected and, preferably, be somewhat balanced in workload. In order to reach partitions with these 'desirable' characteristics, a districting procedure for salt-spreading operations was proposed[15] earlier. This procedure applied similar units (ie. based on cycles) as the $C_{\text{min\_ratio}}$ and C_{ILP} procedures in this paper, but a two-phase allocation heuristic was used for the aggregation. In every phase one iteration, the least-workload district was expanded with the largest-weight adjacent unit. In phase two, the remaining units were assigned by a multi-criteria approach, taking into account a measure for compactness, an estimate on the number of vehicles to be used and an indication for the imbalance in workload. The eligible units in both phases were determined by a threshold value for a proximity measure based on the radial distances for reaching a unit's edges from each of the facilities. On a small, real-life sample network, the districts obtained by this procedure were shown (for a curative intervention) to incur about 14% less deadheading compared with an improved routing within the actual district borders. Nevertheless, no strong argument was provided as to why and when a district design based on cycles might lead to an improved routing. Especially, this last consideration is studied in this article. We will derive and focus on a number of guidelines for district design that substantially influence the routing efficiency, so that based on global information about the distribution system only (the number of facilities, the vehicle capacity, the average customer demand), a more adequate districting policy can be selected.

Guidelines for district design

We assume a connected, undirected, planar graph $G(V, E)$ with vertex set V and edge set E. All the edges $e_{rs} = (v_r, v_s)$ have a positive length c_{rs} and a positive demand for service q_{rs}. Additionally, a set of facilities $X = \{v_1, \ldots, v_i, \ldots, v_p\} \subset V$ is given, each facility housing identical vehicles with a finite capacity Q. In our districting problem, we are interested in a partition of the edges of G such that each district $G_i(E_i, V_i)$ makes up a connected subgraph in G with $v_i \in V_i$ and so that the edges can be serviced efficiently. The edges in E_i may be serviced by trucks from depot v_i only, and two aspects related to efficient routing will be considered: the total deadheading distance (ie the distance driven by the vehicles while they are not servicing) and the number of vehicles to be scheduled.

As the procedures in this paper are of the agglomeration type, the first step deals with the definition of suitable basic units in G (or groups of edges) to be aggregated into districts subsequently. In the following subsection, we analyse two extreme situations where 'optimal' district design turns out to be easy if the only concern is to minimise the total deadheading. In general, by transformation from the NP-complete 'partition problem', it can be shown that finding a district partition that incurs minimal deadheading is NP-hard.[6] The two extreme situations offer, however, useful information concerning the definition of the units (individual edges or units based on cycles) as well as regarding to some objective(s) to guide the aggregation process. We elaborate the approach for pre-clustering of the edges into cycles in the next subsection and explain afterwards when it might be meaningful to address the minimisation of the number of vehicles as an objective in district design for capacitated routing.

Two extreme situations

Let us first consider the situation where a vehicle can service in a tour one edge only. In this case, an optimal partition (minimum deadheading) is found by allocating individual edges to the nearest facility, the distance from facility v_i to an edge $e_{rs} = (v_r, v_s)$ being measured as the sum of the shortest path distances to reach v_r and v_s, from v_i. Hence, when the vehicle capacity Q is (very) small compared with the average edge demand, the radial distances to every edge are important and it is probably a good idea in district design to focus on the centrality of the facilities and on district compactness.

Another extreme occurs when a vehicle can service in a single tour all the edges of E. The minimum deadheading in this giant tour—basically an undirected Chinese postman tour[16] or a minimum-length tour traversing the edges of E at least once — is found by solving a minimum-cost perfect matching problem[16] between the vertices of odd degree in G. This involves duplicating some edges in the original graph G in order to create a Eulerian graph G' (ie a graph whose vertices all have an even degree and that can be covered in a closed walk by traversing every edge exactly once) at minimum cost. Consequently, any division of the augmented graph G' into connected Eulerian subgraphs, each subgraph containing one facility, creates a district partition incurring minimum deadheading (Figure 1).

It is clear from Figure 1 that — for large Q — the local routing cost, that is grouping the edges in a way so that low-cost Eulerian districts can be formed, is more important than assigning individual edges to the nearest facility than focussing on the centrality of the depots. In order to facilitate the construction of low-cost Eulerian districts, we propose decomposing the augmented graph G' into small, non-overlapping cycles, from which the basic units in G are subsequently derived. A potential for good arc routing is then easily maintained by ensuring that each district is composed of a collection of connected cycles.

A cycle decomposition of a planar Eulerian graph

Any Eulerian graph can be partitioned into edge disjoint cycles. It is important, however, to note that we work in the specific, planar representation of a connected, Eulerian

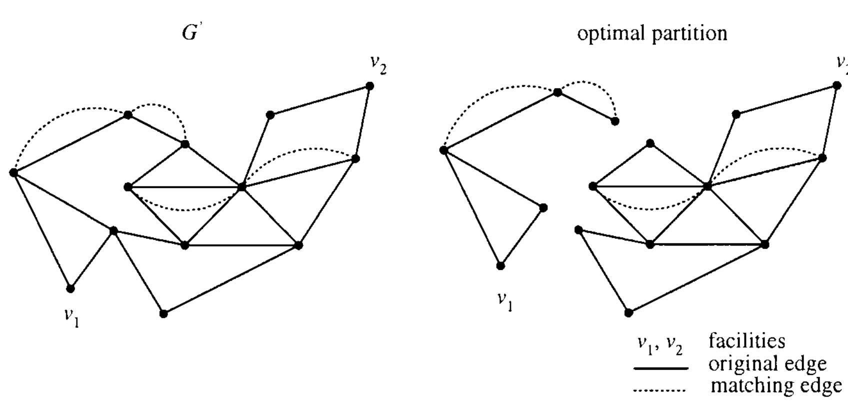

Figure 1 G' and an optimal district partition for $Q = \infty$.

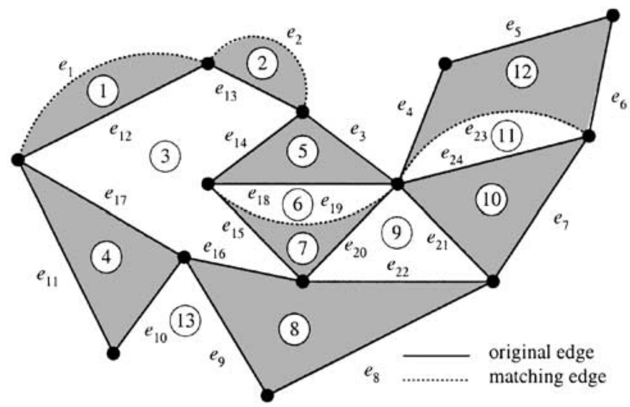

Figure 2 A cycle decomposition of a planar Eulerian graph.

graph that corresponds to the geographical layout of the road network G. This is accounted for in the vertex-edge incidence list of G, where the edges incident to a vertex appear in a cyclic order (eg counterclockwise ordering), agreeable with the planar representation in mind. The graph G is made Eulerian at minimum cost and the matching edges are inserted at proper positions in the vertex-edge incidence lists to preserve the structure of the planar representation. The matching edges have zero demand. The cycle decomposition of the augmented graph G' relies on the property that a planar graph is 2-face-colourable (ie its faces can be coloured in two colours without any two adjacent faces having the same colour) if and only if it is Eulerian (Theorem III. 68)[17] and this property has been exploited earlier for constructing vehicle tours in an arc-routing context.[8,18] In the example in Figure 2, the faces one to 12 are the bounded faces and face 13, with the contour e_1–e_2–e_3–e_4–e_5–e_6–e_7–e_8–e_9–e_{10}–e_{11} and coloured in white, is the unbounded face (the entire area outside the graph). It is readily observed that the contours of the faces that do not belong to the colour class of the unbounded face, partition G' into spatially non-overlapping cycles with every edge of G' contained in exactly one of them.

An initial partition of the edges of the original graph G into basic units is given by the collections of non-matching edges on each of the cycles selected in G'. The augmented graph G' contains, however, pairs of parallel edges; in each pair, one edge being the original (required) edge in G, the other one its duplicate due to the matching. When the edges in a pair belong to the contours of two different selected faces of G', interchanging their role produces other units in G. Therefore, we call the edges in such a pair exchangeable. In Figure 2, for example, the faces five and seven (with the contours e_{14}–e_3–e_{18} and e_{15}–e_{19}–e_{20}) currently define the units e_{14}–e_3–e_{18} and e_{15}–e_{20}, but if e_{18} refers to the matching edge and e_{19} to the required edge, the units e_{14}–e_3 and e_{15}–e_{19}–e_{20} are obtained. The possibility to exchange some of the edges in G' is exploited in a heuristic[6] that we have not described in detail here. Its main purpose is the creation of units u_j, whose edges induce connected subgraphs in G and thus facilitate the construction of contiguous districts. Additionally, we prefer the creation of (many) units with a small demand and (few) units with a large demand above the averaging of the loads, since small weight units are likely to be more useful for adjusting a district partition when the objective of minimising the number of vehicles is considered or when load-balancing constraints are to be taken into account.

In order to express the adjacency relations between the units and the facilities in G, we define a unit-adjacency graph $H(W, F)$. The vertex set W has a vertex w_j for each unit and for each facility in G. The edge set F contains an edge f_{ij} between w_i and w_j if the corresponding units u_i and u_j in G have a vertex in common or, if the facility associated with w_i, is located on the unit represented by w_j. The vertex subsets corresponding to units and facilities are W_{U} and W_{F}, and with every vertex $w_j \in W$, we associate a weight $q(w_j)$, equal to the total demand on unit u_j if $w_j \in W_{\mathrm{U}}$ and zero otherwise.

Finally, we introduce a distance indicator R_{ij} — the ratio — that measures the proximity between the facility $v_i \in X$ ($w_i \in W_{\mathrm{F}}$) and the edges of unit u_j in G ($w_j \in W_{\mathrm{U}}$). Let $D(v_i, u_j)$ denote the sum of shortest deadheading distances in G for servicing each edge e_{rs} of unit u_j separately from facility v_i, that is,

$$D(v_i, u_j) = \sum_{e_{rs} \in u_j} (d(v_i, v_r) + d(v_i, v_s)) \qquad (1)$$

with $d(v_i, v_r)$ being the shortest path distance between v_i and v_r. The ratio R_{ij} compares the radial distances to the edges of unit u_j from the facility v_i, and from the facility v_k ($\neq v_i$) being closest to u_j:

$$R_{ij} = D(v_i, u_j) / \min\{D(v_k, u_j) | v_k \in X, v_k \neq v_i\} \qquad (2)$$

It follows immediately from the definition that there is for each unit u_j at most one facility v_i with $R_{ij} < 1$.

Minimising the number of vehicles as an objective

Apart from a small portion of deadheading, efficient routing for salt spreading also means that the number of vehicles to be scheduled should remain limited. This is because the gritting in Flanders is mostly performed by private subcontractors, whose remuneration typically includes a fixed fee for being on standby in wintertime (with their own truck) and a variable part depending on the number of interventions and the total mileage. While the precise number of vehicles K (or subcontractors) to be used is known only after the routes have been determined, it is easily observed that an uncareful partitioning of the road network (or the total workload) — due to rounding — can lead to an increase in the number of vehicles required and possibly, also to an increase in the mileage. Indeed, the continuous lower bounds on the number of vehicles to be used prior to and after partitioning the network into p districts are, respectively, $K_{lb} = \lceil q(E)/Q \rceil$ and $K_{\text{part},lb} = \sum_i \lceil q(E_i)/Q \rceil$, with $q(E)$ ($q(E_i)$) being the total demand in G (serviced by facility v_i). It can be easily shown[6] that by partitioning G into p districts, in the worst case $K_{\text{part},lb}$ can grow to $K_{\text{part-wc},lb} = K_{lb} + p - 1$.

The ratio $(p-1)/K_{lb}$ can thus be used as an indicator for the importance of minimising the number of vehicles as an objective in districting: the larger $(p-1)/K_{lb}$ or the lower the average number of vehicles needed in each district, the more attention to be paid to the possible increase in the number of vehicles required caused by the partitioning. Owing to the bin-packing aspect in capacitated routing problems, there is, however, no guarantee that district partitions with a low $K_{\text{part},lb}$ value allow the construction of good routings using only $K_{\text{part},lb}$ (or slightly more) vehicles. Apart from the condition of $(p-1)/K_{lb}$ being large, the objective of minimising the number of vehicles might therefore turn out to be effective only in case of the vehicle capacity Q being rather large compared with the average demand on the edges.

Three heuristics

The first two procedures differ only in the definition of the units, while the way of aggregation is similar. Either individual edges (procedure $E_{\text{min\_ratio}}$) or units defined by the cycle approach (procedure $C_{\text{min\_ratio}}$) are allocated to the nearest facility. In the third procedure, (C_{ILP}) units obtained by the cycle decomposition are allocated to the facilities through the solution of an integer linear programming (ILP) model, minimising a lower bound ($K_{\text{part},lb}$) on the number of vehicles required. In each of the three heuristics, we ensure that the network to be serviced within each district is connected and we measure the proximity between a unit u_j and facility v_i by the ratio R_{ij}. The unit-adjacency graph $H(W, F)$ is used instead of the original graph G, although this is not necessary for the $E_{\text{min\_ratio}}$ procedure. Recall that the vertex set W is partitioned into W_{U} and W_{F} corresponding, respectively, to the units and the facilities.

The C_{min_ratio} and E_{min_ratio} procedure

In the following algorithmic description, Step 1′ replaces Step 1 in order to obtain the $E_{\text{min\_ratio}}$ version from the $C_{\text{min\_ratio}}$ procedure.

Step 1 ($C_{\text{min\_ratio}}$): Define the units u_j in G by the cycle decomposition approach (see Figure 2) and the edge exchange heuristic.[6]

Step 1′ ($E_{\text{min\_ratio}}$): Define the units u_j in G by taking every edge $e_{rs} \in E$ as a unit.

Step 2 Construct the unit-adjacency graph H and calculate the ratios R_{ij}. Initially all units $w_j \in W_{\text{U}}$ are unassigned.

Step 3 If all units $w_j \in W_{\text{U}}$ are allocated, go to Step 5.

Step 4 For each facility $w_i \in W_{\text{F}}$, select the unassigned unit $w_j \in W_{\text{U}}$ with the lowest ratio value R_{ij}, adjacent to w_i or to a unit w_k already assigned to facility w_i. Among these candidate facility-unit assignments (w_i, w_j), allocate the unit to the facility in the pair with the lowest R_{ij} value. Go to Step 3.

Step 5 Translate the unit allocations in H into a district partition in G.

$E_{\text{min\_ratio}}$ focuses only on the radial distances to reach every edge; $C_{\text{min\_ratio}}$ tries to take into account local routing aspects as well. Both procedures operate independent of the vehicle capacity Q, so that the generated partitions can be termed robust. The two heuristics run in polynomial time and consume, even for relative large graphs, at most a few seconds of CPU-time on a modest PC, the majority of the computation time being spent on solving the matching problem ($C_{\text{min\_ratio}}$) or on calculating shortest path distances for the ratio computations ($E_{\text{min\_ratio}}$). As an example, we show in Figure 3 the regional road network in the province of Antwerp and the $C_{\text{min\_ratio}}$ partition. The road network edges (342 in total) are pre-clustered into 143 units. Each unit is then allocated to one of the six depots involved.

An integer linear programming approach

Prior to solving the ILP model, the unit-adjacency graph H is reduced in size. The idea is to allocate immediately the units that are considered as very near the depots and further to merge some units by exploiting structural properties in H. In order to accomplish the first goal, we introduce a parameter R_{lim} ($\geqslant 1$) — a limit value for the ratios — by which the units are partitioned into three classes with respect to each facility $w_i \in W_{\text{F}}$: the units w_j that are prohibited from allocation to facility w_i ($R_{ij} > R_{lim}$), the units w_j that are immediately assigned to w_i ($R_{ij} < 1/R_{lim}$) and the units w_j that will be possibly assigned to w_i through the ILP model ($1/R_{lim} \leqslant R_{ij} \leqslant R_{lim}$). The effect of the R_{lim} parameter is illustrated in Figure 4 for a bounded region in the Euclidean plane Π. Figure 4(a) shows the zones defined around three facilities for $R_{lim} = 1$: the points in the region are partitioned according to the classical Voronoi diagram,[19] each point x being assigned to the nearest facility. In Figure 4(b) ($R_{lim} = 2$), each straight line bisector of Figure 4(a) is replaced by two circles (Apollonius circles), indicating the dominance regions with respect to the two facilities; for example, circle $C_{2\cdot 1} = \{x \in \Pi | d_{\text{Eucl}}(2, x)/d_{\text{Eucl}}(1, x) \leqslant R_{lim}\}$. In each zone defined by the circles, the facilities to which the points can be allocated are indicated.

Similarly in a network, a low R_{lim} value limits the number of units and the possible allocations in the ILP model and forces the solution towards a compact district partition. Setting R_{lim} too low, however (eg < 1.25), may prevent solutions from being found with a low $K_{\text{part},lb}$ value or

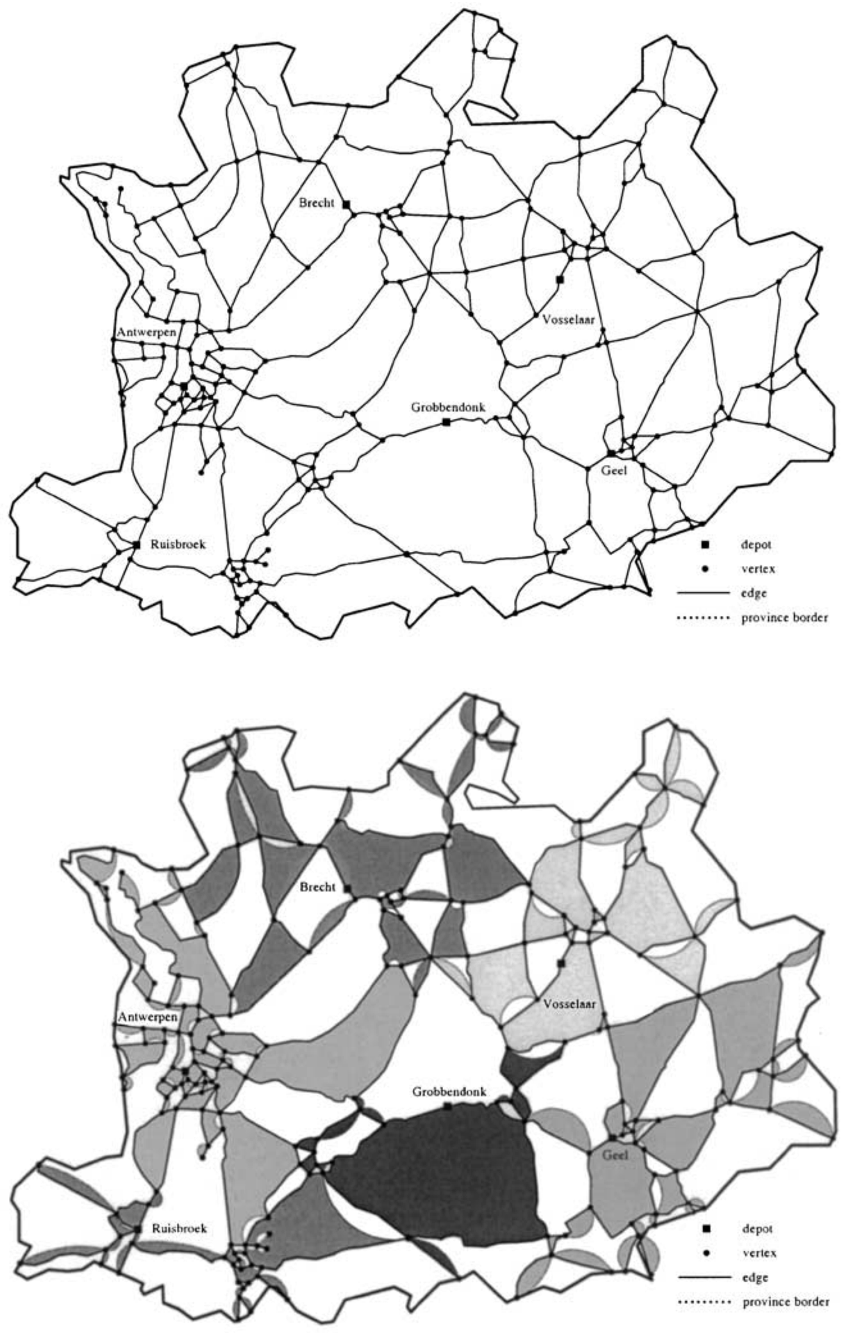

Figure 3 Road network and the $C_{\mathrm{min\_ratio}}$ partition in the province of Antwerp.

prevent the road network within each district from being connected. The latter is usually a local effect and is by-passed in the graph reduction procedure by the possibility to make R_{lim} facility-dependent (ie defining $R_{lim,i}$ values instead of R_{lim}) and by adapting some $R_{lim,i}$ parameters whenever necessary.

The first step in reducing H involves successively merging a (vertex) unit w_j and a facility w_i if the unit can be immediately assigned to w_i ($R_{ij}<1/R_{lim,i}$) and provided that w_j and w_i are adjacent. Next, we investigate H for cut-vertices. The removal of a cut-vertex from H disconnects H into several components (Figure 5). If such a component contains no vertex that corresponds to a facility, all the vertices on that component are merged with the cut-vertex (due to the connectivity). Finally, the reduced graph H is checked whether it still allows a feasible district partition to be constructed: for each facility w_i, the subgraph H_i induced by w_i and the units w_j that can be assigned to w_i should be connected. If this is not the case, some units are further disregarded for assignment to w_i. When it then becomes

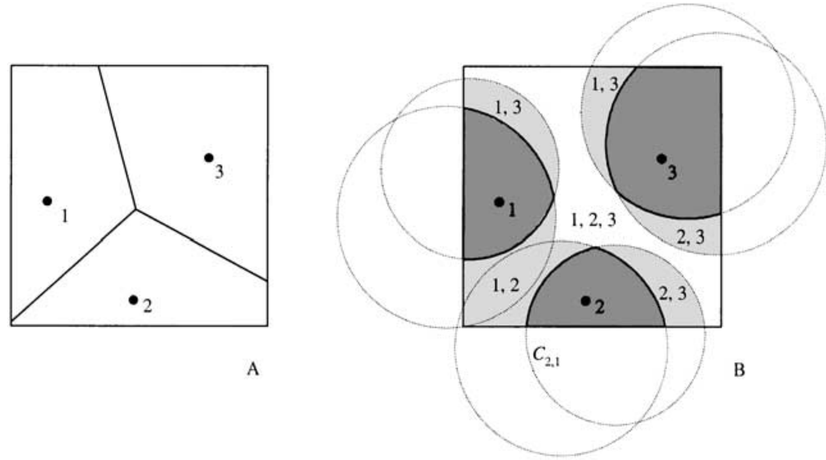

Figure 4 Regions defined by R_{lim} in the Euclidean plane. A: the nearest point Voronoi diagram ($R_{lim}=1$). B: the zones defined by $R_{lim}=2$.

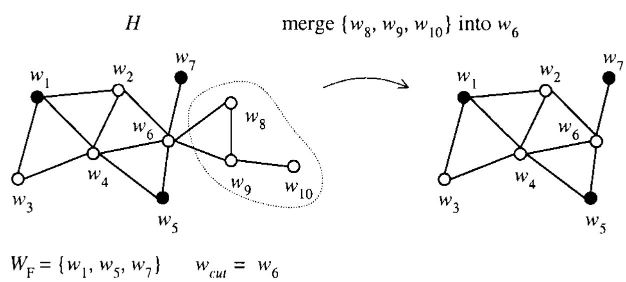

Figure 5 Investigation of a cut-vertex during the reduction of H.

impossible to allocate a unit w_j to any facility, we increase the $R_{lim,i}$ value of the two facilities nearest to w_j and restart the reduction procedure from scratch.

The adjacency graph reduction can be summarised as follows.

Step 1 Set $R_{lim,i} := R_{lim}\ \forall w_i \in W_F$ (eg $R_{lim}=1.5$) and mark all facilities as *unexamined*.

Step 2 If all the facilities are investigated, go to Step 4; else select an unexamined facility w_i.

Step 3 If there exists a unit w_j in H, adjacent to w_i and such that $R_{ij}<1/R_{lim,i}$, merge w_j and w_i into a single vertex (denoted by w_i), set $q(w_i) := q(w_i)+q(w_j)$ and go to Step 3; else, mark w_i *examined* and go to Step 2. (Merging w_j into w_i requires that new edges are added from w_i to the vertices w_k ($\neq w_i$) adjacent to w_j and not already adjacent to w_i, followed by deleting w_j from H.)

Step 4 Make a list L_{cut} that contains the cut-vertices in H.

Step 5 If L_{cut} is empty, go to Step 6; else select the first vertex w_{cut} from L_{cut} and determine the components in $H\backslash\{w_{\text{cut}}\}$. Select the components $(H\backslash\{w_{\text{cut}}\})_k$ that do not contain a facility $w_i \in W_F$. Merge all the vertices w_j on these components into w_{cut} and adapt the weight of w_{cut} by adding the total weight on the vertices in these components. Remove from L_{cut} the vertices that were merged with w_{cut} and also appeared in the list L_{cut}. Remove w_{cut} from L_{cut}. If $w_{\text{cut}} \notin W_F$, recalculate the ratios R_{ij} for w_{cut} with respect to every facility as if all the edges of G that are now represented by w_{cut} in H, would define a single unit. Go to Step 5.

Step 6 For each facility w_i, check if the subgraph H_i induced by w_i and the units w_j that can be assigned to w_i ($R_{ij} \leqslant R_{lim,i}$) is connected. If H_i is disconnected, ignore the vertices w_j on the components different from the one containing w_i, for possible allocation to w_i, (eg set $R_{ij} > R_{lim,i}$ for these units).

Step 7 Check if there are units $w_j \in W_U$ that cannot be allocated. These units have, with respect to every facility, a ratio $R_{ij} > R_{lim,i}$. If none are found, stop; else, for each of these units identify the two nearest facilities. For the selected facilities w_i, set $R_{lim,i} := R_{lim,i}+0.1$. Label all facilities *unexamined* and restart the reduction procedure (with the original graph H) from Step 2.

With the reduced unit-adjacency graph H as the input, we formulate the following ILP model. Three types of variables are defined: a general integer variable y_i for each facility $w_i \in W_F$, corresponding to the continuous lower bound on the number of vehicles scheduled from w_i; binary x_{ij} allocation variables denoting whether unit $w_i \in W_U$ is assigned to facility $w_i \in W_F$ or not, and a continuous $R_{\max,i}$ variable for each facility $w_i \in W_F$ denoting the maximum R_{ij} value among the units assigned to w_i. We further apply the following notations: W_i, for the set of units that can be allocated to facility w_i (ie the vertices of H_i, w_i excluded); W_j, to denote the facilities w_i that can receive unit w_j ($R_{ij} \leqslant R_{lim,i}$) and, W_{ij}, denoting a subset of the vertices in H_i to express the district connectivity when unit $w_j \in W_i$ is assigned to w_i. The formulation is then:

$$\text{Minimise} \sum_{w_i \in W_F} y_i + \alpha \sum_{w_i \in W_F} R_{\max,i} \tag{3}$$

$$\text{Subject to} \sum_{w_i \in W_j} x_{ij} = 1 \quad \forall w_j \in W_U \tag{4}$$

$$\sum_{w_j \in W_i} q(w_j)x_{ij} + q(w_i) \leqslant Qy_i \quad \forall w_i \in W_F \tag{5}$$

$$x_{ij} \leqslant \sum_{w_k \in W_{ij}} x_{ik} \quad \forall W_{ij} \neq \varnothing, w_i \in W_F, w_j \in W_i \tag{6}$$

$$R_{ij}x_{ij} \leqslant R_{\max,i} \quad \forall w_i \in W_F, w_j \in W_i \tag{7}$$

$$x_{ij} \in \{0,1\} \quad \forall w_i \in W_F, w_j \in W_i \tag{8}$$

$$y_i \text{ integer} \quad \forall w_i \in W_F \tag{9}$$

$$R_{\max,i} \geqslant 0 \quad \forall w_i \in W_F \tag{10}$$

The first summation in objective (3) is the continuous lower bound $K_{\mathrm{part},lb}$, on the number of vehicles to be used. Additionally, we penalise the 'non-compactness' of a district partition by taking into account the ratio of the most distant unit that is assigned to each facility. Formulation (3)–(10) may thus be considered as a multiple objective problem, yet the scale factor α is chosen suitably small, for example, $\alpha < 1/(p \times \max\{R_{lim,i} | w_i \in W_F\})$, in order to make the contribution of the second part in (3) less than one so that it does not affect the minimum value of the first term. Constraints (4) express that every unit must be assigned to exactly one facility. Constraints (5) express that the total capacity of the vehicles needed in a facility should cover at least the total demand assigned to it. Constraints (6) ensure the network connectivity within a district. They are explained as follows. For each facility $w_i \in W_F$, the units $w_j \in W_i$ are partially ordered according to the minimum number of edges $n_i(w_j)$ needed for reaching w_j from w_i in $H_i(W_i \cup \{w_i\}, F_i)$. We require that at least one unit of the set $W_{ij} = \{w_k \in W_i | n_i(w_j) = n_i(w_k) + 1 \text{ and } (w_k, w_j) \in F_i\}$ is allocated to w_i, before w_j can be assigned to w_i. Several other ways exist to express district connectivity. Set (6) has the advantage of being polynomial in size and using only the x_{ij} variables, while the drawback is that possibly some feasible solutions are not considered. The number of excluded solutions however, remains limited in the case of $R_{lim,i}$ being rather low. Constraints (7) select the maximum ratio value within each district and constraints (8)–(10) express integrality and non-negativity conditions.

Apart from formulation (3)–(10) ($C_{\mathrm{ILP\_3}}$), two other models are explored. In the first version, only the first summation in (3) is taken into account and constraints (7) and (10) are removed ($C_{\mathrm{ILP\_1}}$). In the second version ($C_{\mathrm{ILP\_2}}$), a single variable $R_{\max}$ is used (instead of a $R_{\max,i}$ variable for each facility) and we penalise for 'non-compactness' by the maximum R_{ij} value over all unit-allocations (with $\alpha < 1/\max\{R_{lim,i} | w_i \in W_F\}$).

Computational experiments

Test instances

In order to explore the quality of the district partitions, we constructed five, planar, undirected graphs (A, B, L, O, W), each of them corresponding to the (simplified) inter-city road network in a province of Flanders (Antwerpen, Vlaams-Brabant, Limburg, Oost-Vlaanderen and West-Vlaanderen). The networks in adjacent provinces were then combined into larger instances (AB,…, WOABL), giving 22 graphs in total. These networks are also relevant for the organisation of salt-spreading activities on regional roads by the Flemish Administration. The average edge length equals 32.8 (or about 3.3 km), and the demand on an edge was set equal to its length. The location of the facilities in the graphs matches the location of the depots from where the gritting is organised in reality. Simplifying the road networks involved that we considered only the secondary regional roads, while the highways were not taken into account as these are serviced in practice by separate tours. The selected roads make up nearly planar graphs. The (very) rare occasions where bridges actually take one road over another were replaced by road intersections. In this study, every road is modelled as an undirected edge, which is a suitable representation for the majority of them ((1 × 1) roads that can be gritted in one pass, servicing both directions together). Although few in number, there are in reality some larger (2 × 2) roads for which the two directions have to be gritted separately. It would be more appropriate to model these by pairs of arcs with opposite direction, so that the road network graph would be mixed. The cycle decomposition approach however, remains applicable for these networks as the resulting mixed graphs are of a special type.[15]

The characteristics of the test problems are summarised in Table 1. The column headings show the problem, the number of facilities p, the number of (required) edges $|E|$, the total demand, the continuous lower bound K_{lb} on the number of vehicles for $Q = 225$, the number of vertices $|W|$ in the unit-adjacency graph H defined by the cycle decomposition and the number of remaining units $|W|_{1,5}$ after reducing H with $R_{lim,i} = 1.5$.

The average reduction in problem size by the cycle decomposition ($|E|/|W|$) is 2.55. Using the H-reduction with $R_{lim,i} = 1.5$ as well, the average reduction factor ($|E|/|W|_{1.5}$) increases to 8.04.

Experiments

Three types of experiments are performed. First, we compare the quality of the partitions generated by $C_{\mathrm{min\_ratio}}$ and

Table 1 Characteristics of the test problems

Problem	p	|E|	Total demand	K_{lb} (Q = 225)	|W|	$\|W\|_{1.5}$
A	6	342	11 340	51	143	29
B	4	259	8170	37	122	35
L	6	369	10 185	46	148	51
O	5	344	12 565	56	152	39
W	6	378	13 180	59	142	48
AB	10	601	19 510	87	244	68
AL	12	711	21 525	96	284	85
BL	10	628	18 355	82	256	96
OA	11	686	23 905	107	280	74
OB	9	603	20 735	93	247	70
WO	11	722	25 745	115	271	97
ABL	16	970	29 695	132	371	130
OAB	15	945	32 075	143	361	121
OAL	17	1055	34 090	152	422	130
OBL	15	972	30 920	138	381	131
WOA	17	1064	37 085	165	399	132
WOB	15	981	33 915	151	365	127
OABL	21	1314	42 260	188	483	165
WOAB	21	1323	45 255	202	480	179
WOAL	23	1433	47 270	211	542	188
WOBL	21	1350	44 100	196	498	188
WOABL	27	1692	55 440	247	598	216

$E_{\text{min\_ratio}}$. Note that both procedures do not take the vehicle capacity Q into account. The quality of a partition is measured by the sum of the CARP deadheading distances in each district of a partition. Furthermore, six values of Q are explored: $Q = 225$, 300, 375, 450, 600 and 900. An upper bound on the CARP deadheading is calculated by a local search heuristic[20] and, to obtain a lower bound, we have implemented a cutting plane approach, based on the so-called supersparse formulation for the CARP.[21] We remark that in the calculation of these bounds, the vehicles can cross the district borders for deadheading. A second set of experiments involves the comparison of $C_{\text{min\_ratio}}$ and the three C_{ILP} variants. We investigate the integer linear programming approach in 17 instances (A–WOB in Table 1) for $Q = 600$ and 900, and for three values of $R_{lim,i}$: 1.3, 1.4 and 1.5. Some valid inequalities[6] are added to the initial formulations and the models are then solved by the standard branch and bound algorithm of CPLEX 6.5 (after tuning the parameters). Both the deadheading distance and the number of scheduled vehicles are used to evaluate the partitions. Finally, we compare the upper bounds on the routing cost for problems A–WOB and the six values of Q, with a multi-depot CARP cutting plane lower bound,[6] the latter approach being adapted from the single depot CARP lower bound procedure.[21] All the algorithms are coded in C\C++ and run on a personal computer (Pentium II, 500 MHz).

Table 2 evaluates $C_{\text{min\_ratio}}$ and $E_{\text{min\_ratio}}$ partitions for $Q = 450$. For every graph, we report the lower and upper bound on the deadheading (lb, ub) and the gap $(ub-lb)/lb$. The last three columns compare the results of both districting procedures and contain the percentage difference

Table 2 Evaluation of the $C_{\text{min\_ratio}}$ and $E_{\text{min\_ratio}}$ partitions for $Q = 450$

	C_{min_ratio}			E_{min_ratio}			Δ		
Problem	lb_C	ub_C	gap (%)	lb_E	ub_E	gap (%)	$(lb_E-lb_C)/lb_C$	$(ub_E-ub_C)/ub_C$	$(lb_E-ub_C)/ub_C$
A	5000	5160	3.20	5890	5990	1.70	17.80	16.09	14.15
B	4800	4865	1.35	5395	5515	2.22	12.40	13.36	10.89
L	4595	4750	3.37	5125	5190	1.27	11.53	9.26	7.89
O	6050	6215	2.73	6215	6270	0.88	2.73	0.88	0.00
W	5190	5305	2.22	5950	6000	0.84	14.64	13.10	12.16
AB	8880	9130	2.82	10 795	10 960	1.53	21.57	20.04	18.24
AL	8945	9370	4.75	10 835	11 015	1.66	21.13	17.56	15.64
BL	9010	9260	2.77	10 130	10 245	1.14	12.43	10.64	9.40
OA	10 965	11 150	1.69	12 050	12 190	1.16	9.90	9.33	8.07
OB	9800	10 035	2.40	11 120	11 285	1.48	13.47	12.46	10.81
WO	10 160	10 545	3.79	12 275	12 390	0.94	20.82	17.50	16.41
ABL	12 515	12 930	3.32	15 190	15 420	1.51	21.37	19.26	17.48
OAB	14 410	14 640	1.60	16 440	16 640	1.22	14.09	13.66	12.30
OAL	14 910	15 370	3.09	16 995	17 225	1.35	13.98	12.07	10.57
OBL	14 015	14 425	2.93	15 855	16 040	1.17	13.13	11.20	9.91
WOA	15 075	15 480	2.69	18 110	18 310	1.10	20.13	18.28	16.99
WOB	14 280	14 685	2.84	17 180	17 405	1.31	20.31	18.52	16.99
OABL	17 335	17 885	3.17	20 835	21 100	1.27	20.19	17.98	16.49
WOAB	18 520	18 980	2.48	22 500	22 760	1.16	21.49	19.92	18.55
WOAL	19 020	19 705	3.60	23 055	23 345	1.26	21.21	18.47	17.00
WOBL	18 490	19 075	3.16	21 915	22 160	1.12	18.52	16.17	14.89
WOABL	21 990	22 795	3.66	26 895	27 220	1.21	22.31	19.41	17.99
Average			2.89			1.30	16.60	14.78	13.31

in deadheading between $E_{\text{min\_ratio}}$ and $C_{\text{min\_ratio}}$, based, respectively, on the lower bounds $(lb_E{-}lb_C)/lb_C$, on the upper bounds $(ub_E{-}ub_C)/ub_C$ and on the lower bound in $E_{\text{min\_ratio}}$ and the upper bound in $C_{\text{min\_ratio}}$ districts $(lb_E{-}ub_C)/ub_C$.

Quite accurate solutions are obtained for the CARP in the districts within each graph, the average gaps counting 2.89 and 1.30% in partitions generated by $C_{\text{min\_ratio}}$ and $E_{\text{min\_ratio}}$. The $C_{\text{min\_ratio}}$ districts definitely allow a better routing than those constructed by $E_{\text{min\_ratio}}$: on average, the lower bound on the deadheading in an $E_{\text{min\_ratio}}$ partition is 16.60% higher than in a $C_{\text{min\_ratio}}$ partition, the upper bounds are 14.78% higher and by comparing the lower bound in $E_{\text{min\_ratio}}$ partitions with the upper bound in $C_{\text{min\_ratio}}$ partitions, we observe that on average, at least 13.31% can be saved in deadheading by applying the cycle approach. Table 3 presents the average results for the six values of Q. We report the average gaps and the average differences in deadheading, in percent.

Clearly, the CARPs are solved more accurately for large Q. More importantly, observe that the $C_{\text{min\_ratio}}$ partitions systematically allow a better routing than the $E_{\text{min\_ratio}}$ partitions. The larger Q, the more $C_{\text{min\_ratio}}$ outperforms $E_{\text{min\_ratio}}$. For $Q = 900$, up to 28% less deadheading is obtained in the partitions based on the cycles. For $Q = 225$, the routing is still slightly better in $C_{\text{min\_ratio}}$ districts, but it is clear that both procedures start competing with each other. This performance is easily explained from the observations made in a previous section: When Q is (very) small, only a few edges can be serviced in a vehicle tour and the radial travel distance will dominate in the routing cost. It is, therefore, a good policy to focus on the construction of compact districts by assigning individual edges to the nearest facility. For larger Q on the other hand, the distances for travelling from one edge to another gain in importance and these local routing cost aspects seem, for arc routing, well accounted for by pre-clustering the edges into cycles.

Table 4 presents the upper bounds on the deadheading (ub) and the actual number of vehicles K used for $Q = 900$, in districts obtained by $C_{\text{min\_ratio}}$ and $C_{\text{ILP\_2}}$, $R_{lim,i} = 1.5$. We show the differences in the upper bound (in percent) and in the number of scheduled vehicles. $K_{lb}/K_{\text{part-wc},lb}$ refers to the continuous lower bound on the number of vehicles before partitioning and after partitioning in the worst case.

Notice that the number of vehicles dispatched ($Q = 900$) in $C_{\text{min\_ratio}}$ partitions always lies in the range $[K_{lb}, K_{\text{part-wc},lb}]$ — most of the time somewhat in the middle. A similar performance, although not reported, has been observed for the other Q values and with respect to the number of vehicles, there was on average no significant difference between $C_{\text{min\_ratio}}$ and $E_{\text{min\_ratio}}$ partitions. Table 4 shows that designing districts according to the ILP approach reduces the number of scheduled vehicles considerably. In fact, with $C_{\text{ILP\_2}}$, we find for each problem a partition where the routing can be carried out with slightly more than K_{lb} vehicles. Moreover, compared with the routing in $C_{\text{min\_ratio}}$ districts, for all but one of the instances the upper bound on the deadheading is also lower. On average, for $Q = 900$, the $C_{\text{min\_ratio}}$ partitions incur 2.74% more deadheading than those found by $C_{\text{ILP\_2}}$, $R_{lim,i} = 1.5$. The performance of the three ILP versions is presented in Table 5. We compare with $C_{\text{min\_ratio}}$ and report the average percentage difference in the upper bound and the sum of the number of vehicles used in the 17 instances ($\sum K$).

Considering the number of vehicles, significant savings are obtained compared with $C_{\text{min\_ratio}}$ even when only few units remain to be allocated in the ILP models ($R_{lim,i} = 1.3$). By increasing $R_{lim,i}$, more combinations are explored and (slightly) better partitions are found. With respect to the upper bound on deadheading, we observe that for $Q = 900$, the three C_{ILP} versions improve upon $C_{\text{min\_ratio}}$, while this is not the case any more for $Q = 600$ — the differences in deadheading however being small. Based on the deadheading, it is more difficult to select the best approach among the three ILPs, although in general, it is better to penalise for non-compactness in the objective function instead of minimising the number of vehicles ($C_{\text{ILP\_1}}$) only. This shows that having the edges near the facilities from where they are serviced remains rather important even for large Q. The entry 'best' C_{ILP} in Table 5 reports the performance if for each instance the best partition (lowest ub) among the nine trials is selected.

The ILPs have also been tested for small Q values. Most of the time, the models do find partitions with $K_{\text{part},lb}$ close

Table 3 Evaluation of the $C_{\text{min\_ratio}}$ and $E_{\text{min\_ratio}}$ partitions for different Q values

Q	C_{min_ratio} gap (%)	E_{min_ratio} gap (%)	Δ $(lb_E{-}lb_C)/lb_C$	$(ub_E{-}ub_C)/ub_C$	$(lb_E{-}ub_C)/ub_C$
225	5.18	4.02	2.30	1.17	−2.74
300	4.11	2.27	6.36	4.48	2.16
375	3.63	2.00	11.68	9.92	7.77
450	2.89	1.30	16.60	14.78	13.31
600	1.87	0.45	22.29	20.58	20.04
900	1.43	0.39	29.33	27.99	27.49

Table 4 Evaluation of the $C_{\min\_ratio}$ and $C_{\mathrm{ILP\_2}}$, $R_{lim,i}=1.5$ partitions for $Q=900$

		C_{min_ratio}		$C_{ILP\_2,1.5}$		Δ	
Problem	$K_{lb}/K_{part\text{-}wc,lb}$	ub_C	K_C	ub_{ILP}	K_{ILP}	$(ub_C-ub_{ILP})/ub_{ILP}$	K_C-K_{ILP}
A	13/18	3685	17	3515	14	4.84	3
B	10/13	3025	11	3050	10	−0.82	1
L	12/17	3105	14	3035	12	2.31	2
O	14/18	3995	17	3825	15	4.44	2
W	15/20	3510	19	3440	15	2.03	4
AB	22/31	5955	28	5905	23	0.85	5
AL	24/35	6520	31	6275	25	3.90	6
BL	21/30	5800	25	5615	22	3.29	3
OA	27/37	7510	34	7320	28	2.60	6
OB	24/32	6220	27	6020	25	3.32	2
WO	29/39	6635	33	6365	32	4.24	1
ABL	33/48	8350	42	8310	35	0.48	7
OAB	36/50	9145	45	8750	38	4.51	7
OAL	38/54	10 345	48	10 030	40	3.14	8
OBL	35/49	8995	41	8615	37	4.41	4
WOA	42/58	10 150	50	10 035	44	1.15	6
WOB	38/52	9010	44	8840	39	1.92	5
	Σ433/601		Σ526		Σ454	Average 2.74	

Table 5 Evaluation of the C_{ILP} partitions for $Q=600$ and 900

		$Q=600$		$Q=900$	
Procedure	$R_{lim,i}$	$(ub_C-ub_{ILP})/ub_{ILP}$	ΣK	$(ub_C-ub_{ILP})/ub_{ILP}$	ΣK
$C_{\min\_ratio}$	—	—	725	—	526
$C_{\mathrm{ILP\_1}}$	1.3	−0.58	683	1.78	463
	1.4	−1.30	679	1.60	460
	1.5	−2.54	673	1.27	452
$C_{\mathrm{ILP\_2}}$	1.3	0.21	683	1.93	458
	1.4	−0.32	675	2.68	458
	1.5	−0.60	673	2.74	454
$C_{\mathrm{ILP\_3}}$	1.3	0.14	684	1.66	463
	1.4	−0.25	673	1.88	455
	1.5	−0.81	671	1.66	454
'best' C_{ILP}		1.92	674	3.64	457

to K_{lb}. Since the majority of these districts however, have a total demand equal (or almost equal) to an integer multiple of Q (involving 6–12 vehicles), good routing schemes with only $K_{\mathrm{part},lb}$ vehicles are rarely obtained. Moreover, we usually end up with partitions performing worse in the number of vehicles used and in deadheading than the $C_{\min\_ratio}$ partitions. As could be expected, the ILPs are therefore most suited when Q is relatively large.

With respect to the computation times, we note that $C_{\mathrm{ILP\_1}}$ solves faster than $C_{\mathrm{ILP\_2}}$, while solving $C_{\mathrm{ILP\_2}}$ requires less time than $C_{\mathrm{ILP\_3}}$. The computation times increase for increasing problem size (a larger network or a larger $R_{lim,i}$) and for decreasing Q. Networks spanning one province are solved almost immediately or within a few seconds; two and three provinces mostly require several seconds to a few minutes to solve. In three of the 306 ILP problems, the solution time exceeded one hour to reach and prove optimality.

Finally, Table 6 evaluates the quality of the routing in the $C_{\min\_ratio}$ partitions by comparison with a multi-depot CARP cutting plane lower bound. We report the average, worst and best gap over 17 instances (A–WOB) for the six

Table 6 Gaps between the routing in $C_{\min\_ratio}$ partitions and a multi-depot CARP lower bound

	Q					
Gap (%)	*225*	*300*	*375*	*450*	*600*	*900*
Average	33.49	34.35	32.57	28.17	23.32	16.92
Worst	38.20	41.05	42.04	34.00	33.58	22.81
Best	19.43	18.55	25.42	12.49	15.55	9.60

values of Q. As in the other tables, the gaps are based on the deadheading only, acknowledging that the length on the required edges is a constant. If we include the length on these edges, the average gaps for $Q = 225$ and 900 would read 14.76 and 3.42%, instead of 33.49 and 16.92%.

Table 6 shows that the degree of suboptimisation due to the districting remains rather limited. Indeed, the multi-depot CARP does not impose any restriction (connectivity, no overlap) on the subgraphs induced by the edges that are serviced from the same facility. The gaps between the routing in the $C_{\text{min\_ratio}}$ partitions and the multi-depot CARP lower bound are, therefore, fairly reasonable, but undoubtedly, good multi-depot CARP algorithms can improve upon the routing in the district partitions. Comparing their performance with the lower bound would allow a finer evaluation of the quality of the lower bound and of the routing within the partitions.

Recommendations for practice

The need for districting models and for more insight into the interaction between districting and routing are evident, as the Flemish Administration has hired a consultancy firm in the past for reconfiguring the district boundaries and the routes. These new routes and districts are partly based on climatical considerations and on thermal characteristics of the roads so that when the ground temperature fluctuates around the freezing point, no gritting must be performed in the warmer zones and some routes can be excluded. Apart from these thermal data, no other aspects related to efficient routing were considered in the new design and the current district borders coincide to a large extent with irrelevant province and commune boundaries. From the experiments in the previous section and by taking the operational characteristics for preventive and curative gritting into account (a truck can spread about 70 and 35 km in a preventive and curative intervention, and the average edge length equals 3.3 km), districting based on cycles would almost surely incur much less deadheading. For preventive and curative operations, the $C_{\text{min\_ratio}}$ partitions are estimated to incur 23–25 and 8–10% less deadheading compared to $E_{\text{min\_ratio}}$ partitions. Alternatively, the ILP approach can be used, for instance with the operational settings for preventive gritting ($Q = 700$) and rather low $R_{lim,i}$ values. The resulting districts would perform for preventive operations similarly in routing quality as $C_{\text{min\_ratio}}$ districts, although fewer vehicles would be required. For low $R_{lim,i}$ values, the districts are compact in shape so that good curative tours can be constructed as well. Furthermore, important savings can be achieved by looking across the province borders: for $Q = 450$ (see Table 2), applying $C_{\text{min\_ratio}}$ on the road network of each province separately ($A + \ldots + W$) and on the entire network in Flanders (WOABL), yields upper bounds on the deadheading of 26 295 and 22 795, ie a difference of 15.35%. For preventive and curative interventions ($Q = 700$ and 350), the additional savings in deadheading are estimated to be around 22 and 12%, respectively. Nevertheless, it might be recommended to adapt the procedures and to address climatical considerations as well. The network edges and/or the units can be subdivided into classes according to their thermal characteristics. This information can then be exploited in the edge exchange heuristic[6] or in the districting procedures in order to determine more thermally homogeneous cycles or districts and to reach a superior district partition eventually.

Conclusions

In this paper, we have addressed the problem of district design for the organisation of arc-routing activities. Districting is viewed as an intermediate step between the location of the facilities and the determination of the routes. Since well-designed districts should be able to support efficient routing, it is recommended to include, along with other criteria, routing aspects in a districting procedure. We propose and test three heuristics for district design, each procedure being inspired by a number of guidelines in order to perform well in a certain value range for the vehicle capacity Q. When Q is very small, only a few edges are serviced in a tour and the radial travel cost will dominate. It is, therefore, a good policy to allocate individual edges to the nearest facility ($E_{\text{min\_ratio}}$). For larger Q, we pre-cluster the edges into small cycles, recognising that the local routing cost becomes important. In $C_{\text{min\_ratio}}$, cycles are assigned to the closest facility. Thus, aspects related to both the radial and the local travel costs are taken into account, and the procedure performs well for average Q values. In C_{ILP}, the focus is mainly on the local routing cost and the cycles are aggregated into districts through the solution of an ILP model, minimising a lower bound on the number of scheduled vehicles. This last approach is particularly suited when Q is large.

The performance of the three procedures is validated through extensive computational experiments and by comparison with a multi-depot CARP lower bound, the latter showing that the degree of suboptimisation due to districting remains rather limited. Finally, some recommendations for improving the efficiency of the gritting operations in Flanders are given.

References

1 Mehrotra A, Johnson EL and Nemhauser GL (1998). An optimization based heuristic for political districting. *Mngt Sci* **44**: 1100–1114.
2 Fleischmann B and Paraschis JN (1988). Solving a large scale districting problem: a case report. *Comput Opns Res* **15**: 521–533.

3 Pezzella F, Bonanno R and Nicoletti B (1981). A system approach to the optimal health care districting. *Eur J Opl Res* **8**: 139–146.
4 Ferland JA and Guénette G (1990). Decision support system for the school districting problem. *Opns Res* **38**: 15–21.
5 D'Amico SJ, Wang S, Batta R and Rump CM (2002). A simulated annealing approach to police district design. *Comput Opns Res* **29**: 667–684.
6 Muyldermans L (2003). *Routing, districting and location for arc traversal problems*. PhD thesis, Centre for Industrial Management, Katholieke Universiteit Leuven.
7 Grilli di Cortona P *et al* (1999). *Evaluation and Optimization of Electoral Systems*. SIAM Monographs on Discrete Mathematics and Applications: Philadelphia, PA.
8 Eglese RW (1994). Routing winter gritting vehicles. *Discr Appl Math* **48**: 231–244.
9 Li LYO and Eglese RW (1996). An interactive algorithm for vehicle routing for winter gritting. *J Opl Res Soc* **47**: 217–228.
10 Golden BL and Wong RT (1981). Capacitated arc routing problems. *Networks* **11**: 305–315.
11 Dror M (ed) (2000). *Arc Routing: Theory, Solutions and Applications*. Kluwer Academic Publishers: Dordrecht.
12 Lotan T, Cattrysse D and Van Oudheusden D (1996). Winter gritting in the province of Antwerp: a combined location and routing problem. *Belg J Opns Res Statist Comput Sci* **36**: 141–157.
13 Cattrysse D, Van Oudheusden D and Lotan T (1997). The problem of efficient districting—improving distributed operations depends on efficient partitioning of a service area into districts. *OR Insight* **10**: 9–13.
14 Simchi-Levi D (1992). Hierarchical planning for probabilistic distribution systems in Euclidean spaces. *Mngt Sci* **38**: 198–211.
15 Muyldermans L, Cattrysse D, Van Oudheusden D and Lotan T (2002). Districting for salt spreading operations. *Eur J Opl Res* **139**: 521–532.
16 Edmonds J and Johnson EL (1973). Matching, Euler tours and the Chinese postman problem. *Math Program* **5**: 88–124.
17 Fleischner H (1990). *Eulerian Graphs and Related Topics, Part 1 Volume 1. Annals of Discrete Mathematics Volume 45*. North-Holland: Amsterdam.
18 Male JW and Liebman JC (1978). Districting and routing for solid waste collection. *J Environ Eng Div ASCE* **104**: 1–14.
19 Okabe A, Boots B and Sugihara K (1992). *Spatial Tessellations. Concepts and Applications of Voronoi Diagrams*. Wiley: Chichester.
20 Beullens P, Muyldermans L, Cattrysse D and Van Oudheusden D (2003). A guided local search heuristic for the capacitated arc routing problem. *Eur J Opl Res* **147**: 629–643.
21 Belenguer JM and Benavent E (1998). The capacitated arc routing problem: valid inequalities and facets. *Comput Optim Appl* **10**: 165–187.

Received April 2002;
accepted July 2003 after one revision