SALTING ROUTE OPTIMISATION USING XRWIS AND EVOLUTIONARY COMPUTATION

Lee Chapman[1], Hisashi Handa[2], Xin Yao[1] & John E. Thornes[1]

[1]The University of Birmingham, UK. Email: l.chapman@bham.ac.uk
[2]Okayama University, Japan. Email: handa@sdc.it.okayama-u.ac.jp

ABSTRACT

With limited resources and treatment time constraints, it is imperative that salting routes are planned in advance for efficient and effective winter road maintenance. To aid this process, a salting route optimisation system which combines evolutionary algorithms with the neXt generation Road Weather Information System (XRWIS) has been developed. The system can cope with large-scale instances in the real world within reasonable computation times, to the extent that daily dynamic salting route optimisation can be realised. However, the use of a dynamic system, that each day adjusts salting routes in line with forecast road temperatures, may increase complexity to the extent that user error will occur in the treatment regime. Therefore, a robust (static) solution of salting route optimisation is also presented. Here, the emphasis is placed on thermally ranking optimised routes so that the 'warmer' routes could be left untreated on marginal nights.

Keywords: Salting Routes, XRWIS, Road Surface Temperatures

1. INTRODUCTION

On marginal winter nights, highway engineers often face a difficult decision of whether or not to precautionary salt the road network. The consequences of making a wrong decision are serious, as an untreated network is a major hazard. However, if salt is spread when it is not actually required, there are unnecessary financial and environmental consequences. The decision of whether to salt is taken by consulting a Road Weather Information System (RWIS) which combines weather forecast data with road temperature and condition data. The first generation of RWIS currently in use relies on methods and tools developed in the 1980s, but as technology has moved on, it is now being superseded by the neXt generation RWIS (XRWIS) [7]. Instead of relying on forecast interpolations made by thermal mapping, complete with its inherent limitations, XRWIS models road surface temperatures by considering the influence of the local geography on the road surface [1]. The forecast is displayed in a GIS environment and disseminated direct to the highway engineer via the Internet (Figure 1). In addition to road condition and temperature forecasts, a key feature of XRWIS is an automated decision making algorithm where salting routes are colour coded depending on the required action (Figure 1d). Enormous potential exists for savings to be made by using this route-based forecasting approach by leaving the warmer routes untreated.

Despite the large advance in technology provided by XRWIS, additional tools are required to make the most effective use of resources; if routes can be thermally optimised (i.e. colder road sections grouped together on a single route) or dynamic (new routes developed daily), additional savings can be made. This is where improved salting route optimisation is required. The task of optimising routes driven by winter maintenance fleets is an area open to improvement, as previous work has been limited due to the complex nature of the problem. Optimisation has traditionally been an expensive manual exercise, heavily reliant on local experience and knowledge. Once salting routes have been determined, a ‘static’ and often paper based approach has been taken to try to optimise these for effective use of vehicles, people and material, within the given constraints, to maintain safe road conditions and minimal damage to the environment through road wash off [6].

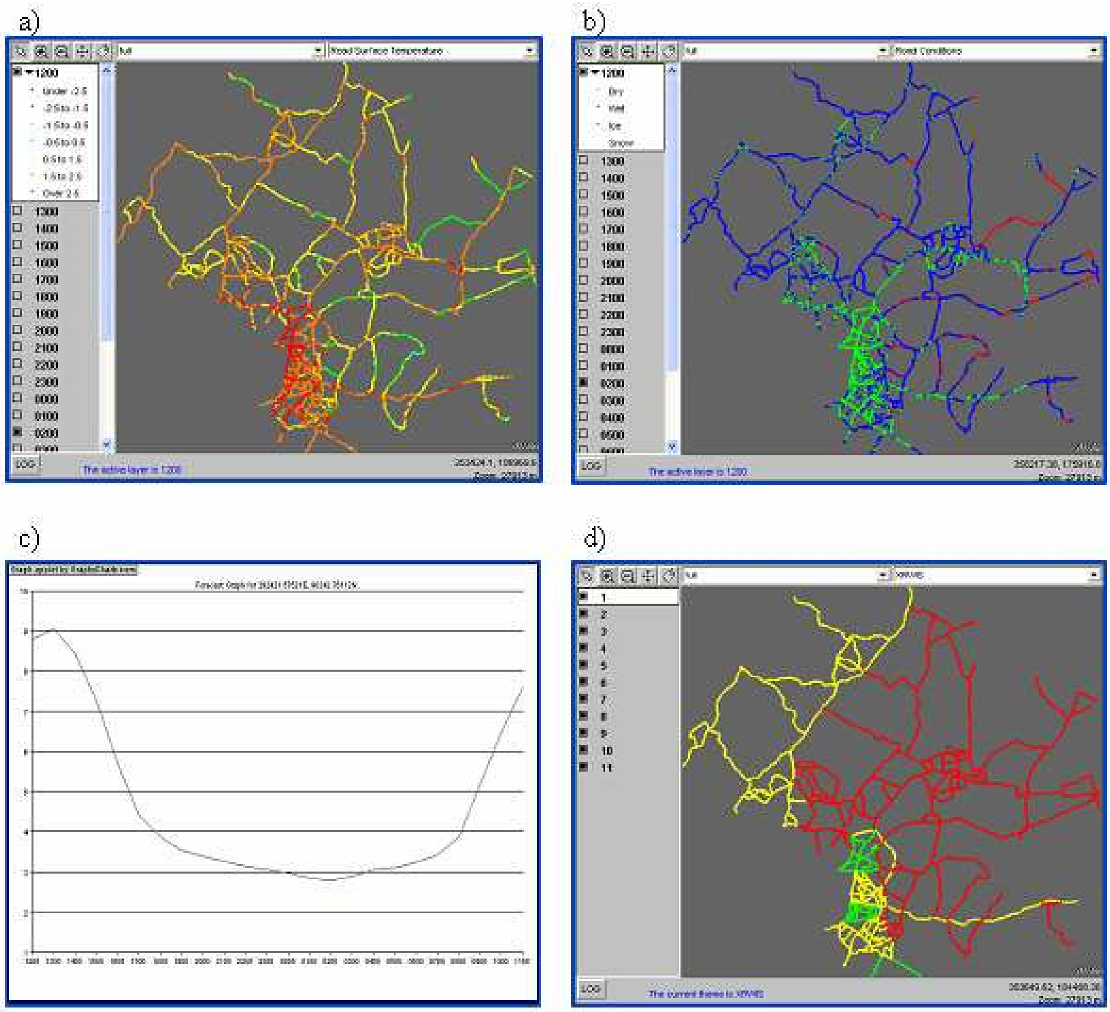

Fig. 1. Sample South Gloucestershire XRWIS output showing a) forecast road surface temperature, b) forecast road surface condition, c) forecast curve showing site specific variation in road surface temperature and d) traffic light colour-coded decisions; Red – route requires treating, Amber – standby (watch forecast) and Green – no action required.

Salting route optimisation traditionally involves maximising efficiency such that the distance treated as a percentage of the total route length is minimised. However, the modelling and optimisation of salting routes is extremely complex and involves numerous interdependent variables (pertaining to the road network, treatment requirements, material characteristics, vehicle numbers and capacities, depot location, availability of staff and limitations of treatment times). The task is better suited to computer analysis combined with accurate road network databases (Figure 2), rather than a manual approach. Successful attempts to solve the problem have been made using computational intelligence (e.g. [3]), but the use of new cutting edge evolutionary algorithms, which have been shown to outperform traditional techniques for similar routing problems, will provide a superior and robust means for route optimisation. Evolutionary algorithms are a generic term used to describe a population-based optimisation algorithm that uses mechanisms inspired by genetics or biological evolution. Such techniques offer a distinct advantage in finding satisfactory solutions (although, not necessarily 100% optimal) quickly, and as such are ideally suited to the potentially time consuming and complex nature of the salting route optimisation problem. This paper discusses how by using a synergy of evolutionary algorithms, XRWIS and geographical information science, new salting route optimisation tools can be developed.

Fig. 2. Sample 'commercial, off the shelf' vector road network data suitable for use in salting route optimisation algorithms. Inset shows the finer detail of traffic restriction orders such as one-way streets (arrows), dual carriageways, intersections and bridge decks (bold).

2. DYNAMIC SALTING ROUTE OPTIMISATION

2.1 Mapping of road temperature as a capacitated arc-routing problem

Salting route optimisation can be regarded as an instance of the Capacitated Arc Routing Problem (CARP) [2-3]. A road network requiring treatment can be modelled as a set of vertices and edges where an edge will incur a cost if there is a need/demand to treat it. There are potentially several vehicles to fill the demands, each with a predefined service capacity. Additionally, all vehicles must depart from defined depot(s) and return there at end of their service tour. The problem is to find a set of tours which have a minimum total cost for all vehicles, ensuring the demands of all required edges are satisfied by at least one vehicle and whilst maintaining that the total services of each vehicle are within capacity. A tour for each vehicle is defined as a sequence of required edges plus the connecting edges not needing treatment but required for travel between demands (deadheading). The total cost of the tour is the sum of required edges in a tour plus the deadheading costs. A distance matrix, calculated by using Dijkstra's Algorithm, is employed between vertices and remains unchanged during evolutionary search. The set of required edges and their demands (i.e. amount of salt), are set by referring to the temperature of prediction points provided by the XRWIS. Figure 3 shows two examples of forecast temperatures on routes in the South Gloucestershire area for two winter nights. These temperature distributions are then translated into CARP instances for the two nights (Figures 3c & 3d). If any point on a particular edge falls below the 0ºC threshold, then the whole edge is classified as requiring treatment. The length of required edges is minimised by initially 'exploding' the entire vector data shown in Figure 2 into the smallest edge lengths possible.

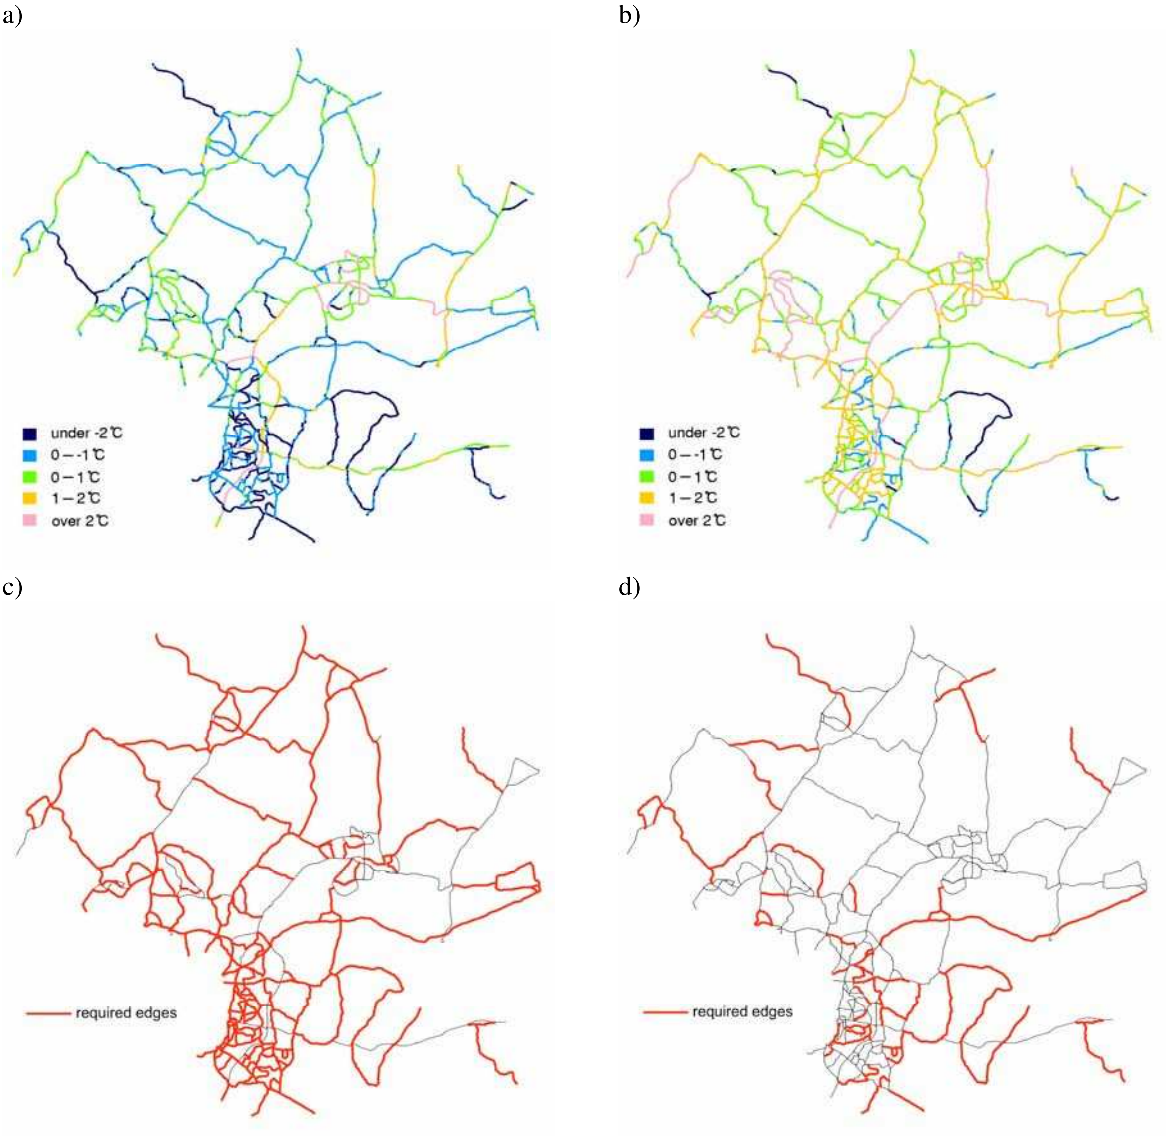

Fig. 3. Temperature distributions in South Gloucestershire, UK on a) a cold winter night and b) a marginal winter night. These temperature distributions are then translated into required edges in c) and d) respectively.

2.2 Evolutionary computation

A new memetic algorithm for solving large-scale salting route optimisation problems has been developed to tackle the salting route optimisation problem. A population-based approach is used that mimics the process of evolution and combines local search heuristics with genetic algorithms. The aim is to find an approximate solution to the problem before using feedback to improve performance. A population of individuals, represented as chromosomes, is created by path scanning (e.g. [2]). These individuals then undergo a process of evolution where they compete for resources. Successful individuals have a higher probability of survival and reproduction (crossover) where information is exchanged between individuals and passed on as a new offspring individual. Due to the large-scale nature of the salting route optimisation problem, the edge assembly crossover operator, developed by [4], was used due to its powerful search abilities (e.g. [5]). However, since this operator was not specifically designed for this problem, it can often yield infeasible solutions. Hence, naive local search methods and a repair operator for offspring individuals are also incorporated in the memetic algorithm.

Whether or not information is exchanged is dependent on the success (fitness *F*) of each individual and is evaluated by using a fitness function. In order to cope with constraints with respect to the capacity of the services for the demands, the following fitness function is used:

$$F = \sum_{i=0}^{m} \left(C_{Ti} + C_p \times E_{Ti} \right) \tag{1}$$

where m is the number of trucks, C_{Ti} denotes the cost of the trucks ith trip T_i, C_p is a predefined coefficient for the penalty term. E_{Ti} indicates the quantity of constraint violation in each truck:

$$E_{Ti} = \begin{cases} D_{Ti} - L_i & \text{if } D_{Ti} - L_i > 0 \\ 0 & \text{Otherwise,} \end{cases} \tag{2}$$

where D_{Ti} and L_i denote the total services for the demands by truck i and the capacity of services for truck i respectively. Once evolutionary search is complete, a chromosome describes the order of required edges within the road network. Special symbols are then used indicating the beginning of tours for each truck into the chromosome. The end result is a series of tours representing the individual salting routes.

2.3 Results

The memetic algorithm was run for the two winter nights translated into CARP instances shown in Figure 3. The acquired solutions and statistics for the two nights are shown in Figure 4 and Table 1 respectively and clearly show the potential of the algorithm in providing dynamic salting route solutions. The financial benefits of the dynamic approach are visibly evident, with the number of routes requiring treatment being reduced to just three on the marginal night.

a) b)

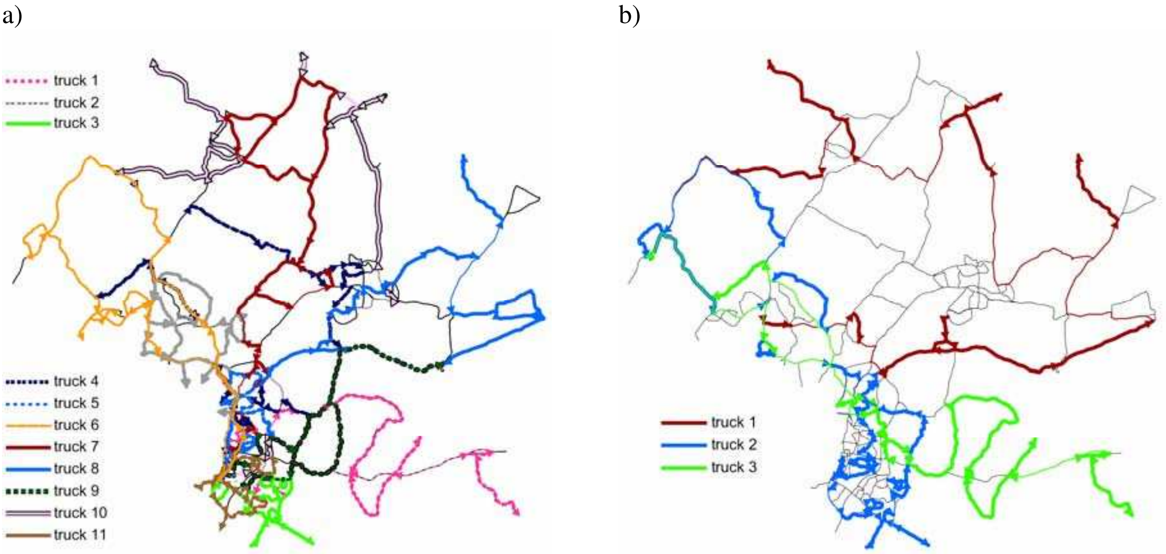

Fig. 4. Acquired solutions from the memetic algorithm for the CARP instances shown in Figure 3 for a) cold night and b) marginal night. A thick line denotes that truck should treat the road where as a thin line denotes the truck travels without service on the road.

	No. of Required Edges	No. of Salting Routes	Approx. Processing Time[1]
Cold Night	385	11	6 hours
Marginal Night	97	3	50 minutes

[1] Times derived using a PC cluster consisting of 22 Itanium 2 processors.

Table. 1. Summary statistics of the salting route optimisation shown in Figure 4.

3. ROBUST (STATIC) SOLUTION

This paper has so far discussed a dynamic solution for salting route optimisation. However, such a system is too advanced for operational use at present. Reasons for this include:

1. Slow computer processing. The South Gloucestershire example used in this paper is a small road network but can still take up to 6-7 hours to produce a solution (Table 1).
2. Compatible satellite navigation technology required on the entire salting fleet coupled with real-time treatment archives.
3. XRWIS technology still undergoing testing and development

To overcome these problems, there is a need to phase in the technology. It is for this reason why a robust (static) solution for salting route optimisation has also been developed. This is a more traditional approach where the same salting routes are used every night throughout the winter season. However, where this approach differs from existing techniques is that routes are specifically designed to concatenate road sections with similar thermal characteristics (e.g. cold bridge decks) so that routes can be given priority treatment. Similarly, routes with relatively warmer thermal characteristics (e.g. heavily trafficked city centres) can be left untreated on the more marginal nights

The robust solution is also derived by using the memetic algorithm and simply assumes that all edges in the network require treatment. However, the robust solution differs from the dynamic routing technique in that multiple scenarios are considered. In this example, initial salting routes are generated based on an archive of past temperature distributions forecasted from XRWIS (where XRWIS is not available, data can be digitised from existing thermal maps). These initial salting routes are denoted as a set of tours for a variable number of trucks that cover the whole area and can be further tested on additional temperature distributions (if available). Figure 5 shows the robust routes generated by the algorithm for South Gloucestershire compared to the existing routes presently in use. A total of 10 training temperature distributions covering a range of marginal scenarios were used to provide the solution. Although, the algorithm has been unable to reduce the number of routes required, there is greater scope for selective salting with the new routes by priority treating the colder rural routes to the south-east of the county. For example, for 50% of the marginal nights in the training dataset, the number of routes requiring treatment were 10 or less. Furthermore, with respect to distance travelled (i.e. deadheading costs) the robust routes are 10% more efficient then the current operational routes.

a) b)

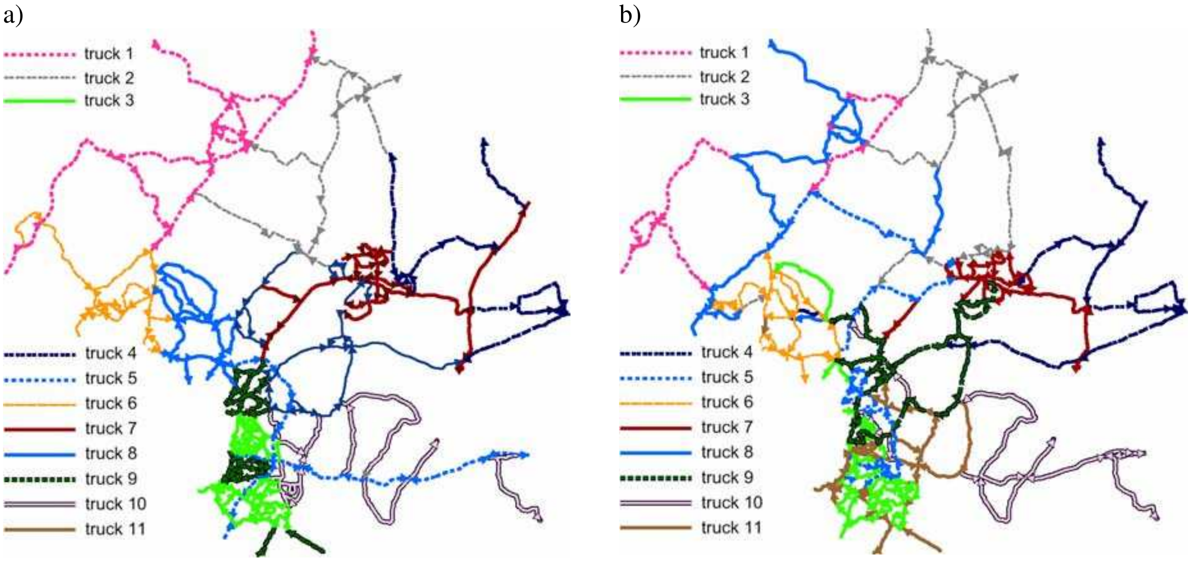

Fig. 5. a) Existing salting routes in operational use in South Gloucestershire and b) the derived robust solution.

Ultimately, robust routes provide a solution that is far simpler for operational use than that of dynamic routing. By using static routes, errors in treatment regimes can hopefully be avoided. Highway engineers then have a series of options (in order of cost-benefit):

1. Treat every edge (road) on every static salting route (the traditional approach)
2. Treat every edge on any static salting route forecast to go below 0ºC
3. Selectively salt only the edges on static salting routes requiring treatment as dictated by XRWIS (reduced salting costs, but high deadheading costs)
4. Full dynamic routing (reduced salting and deadheading costs)

4. CONCLUSIONS

In this paper, a series of new salting route optimisation tools have been proposed which combine XRWIS with evolutionary computation and geographical information science techniques. By using temperature data predicted by XRWIS and by taking account of constraints regarding treatment vehicles and the road network, the algorithm can quickly, and cheaply, optimise a series of robust salting routes for winter road maintenance. These can then be used for daily treatment regimes, with the 'warmer' routes being eliminated on marginal nights. Once this has been achieved, the next step is to begin to phase in selective salting on the robust routes driven by XRWIS forecasts. Once policy catches up with the science, there is no reason why daily dynamic routing cannot become a reality. The financial savings using a robust solution are potentially less than with a dynamic solution, but these are still a considerable advance on existing techniques. Further work is still required in developing the memetic algorithm as the technique is still in its infancy. There is also a need to extend the work to cover a larger area than that used in this paper. Indeed, the larger the road network to be optimised, the greater the potential benefits of using this approach.

ACKNOWLEDGEMENTS

This work was carried out while Hisashi Handa was a visiting research fellow at CERCIA (Centre of Excellence for Research in Computational Intelligence and Applications) at the University of Birmingham. This visit was supported by the Ministry of Education, Science, Sports and Culture, Japan through the Overseas Advanced Educational Research Practice Support Program. Additional thanks to the team at South Gloucestershire County Council, UK whose continued support is very much appreciated.

REFERENCES

[1] Chapman, L. & Thornes, J.E. (In Press) A geomatics based road ice prediction system. *Science of the Total Environment*

[2] Lacomme, P., Prins, C. & Ramdanecherif, W. (2004) Competitive memetic algorithms for arc routing problems. *Annals of Operations Research* **131**:159-185

[3] Li, L.Y.O. & Eglese, R.W. (1996) An interactive algorithm for vehicle routing for winter gritting. *Journal of the Operational Research Society* **47**:217-228

[4] Nagata, Y. & Kobayashi, S. (1997) Edge assembly crossover: a high-power genetic algorithm for the traveling salesman problem. *Proceedings of the 7th International Conference on Genetic Algorithms*, *East Lansing, Michigan, USA*, *July 1997,* 450-457pp

[5] Nagata, Y. (2004) The EAX algorithms considering diversity loss. *Proceedings of the 8th International Conference on Parallel Problem Solving from Nature, Birmingham, UK September 2004 – PPSN VII.* 332-341pp

[6] Thornes, J.E. (1996) The cost benefit of winter road maintenance in the United Kingdom. *Proceedings of the 6th International Road Weather Conference, Birmingham, UK, April 1996. SIRWEC* 1-10pp

[7] Thornes, J.E., Chapman, L. & White S.P (2005) The X-Factor. *Surveyor Magazine* **192**:14-16