Scheduling Planned Maintenance of the South Wales Region of the National Grid

W. B. Langdon*

School of Computer Science, University of Birmingham,
Birmingham B15 2TT

22 November 1996

Abstract

The maintenance of the high voltage electricity transmission network in England and Wales (the National Grid) is planned so as to minimise costs taking into account:

- location and size of demand for electricity,
- generator capacities and availabilities,
- electricity carrying capacity of the remainder of the network, i.e. that part not undergoing maintenance,
- resilience of the network to faults (contingencies).

This complex optimisation and scheduling problem is currently performed manually (albeit with some computerised assistance). This paper reports recent work aiming to automatically generate low cost schedules using genetic algorithms (GA) using the South Wales region as a demonstration network.

The combination of a "greedy optimiser' with a permutation GA, which has been demonstrated on a small network, was successfully applied to the South Wales network.

1 Introduction

In England and Wales electrical power is transmitted by a high voltage electricity transmission network which is highly interconnected and carries large power flows. It is owned and operated by The National Grid Company plc. (NGC) who maintain it and wish to ensure its maintenance is performed at least cost, consistent with plant safety and security of supply.

There are many components in the cost of planned maintenance. The largest is the cost of replacement electricity generation, which occurs when maintenance of the network prevents a cheap generator from running so requiring a more expensive generator to be run in its place.

The task of planning maintenance is a complex constrained optimisation scheduling problem. The schedule is constrained to ensure that all plant remains within its capacity and the cost of replacement generation, throughout the duration of the plan is minimised. At present maintenance schedules are initially produced manually by NGC's Planning Engineers (who use computerised viability checks and optimisation on the schedule after it has been produced). This paper describes work by NGC's Technology and Science Laboratories and University College London which investigates the feasibility of generating practical and economic maintenance schedules using genetic algorithms (GAs) [Hol92].

Earlier work [Lan95] investigated creating electrical network maintenance plans using Genetic Algorithms on a demonstration four node test problem devised by NGC [Dun93]. [Lan95] showed the combination of a Genetic Algorithm using an order or permutation chromosome combined with hand coded "Greedy" Optimisers can readily produce an optimal schedule for this problem.

[1] *email:* W.B.Langdon@cs.bham.ac.uk

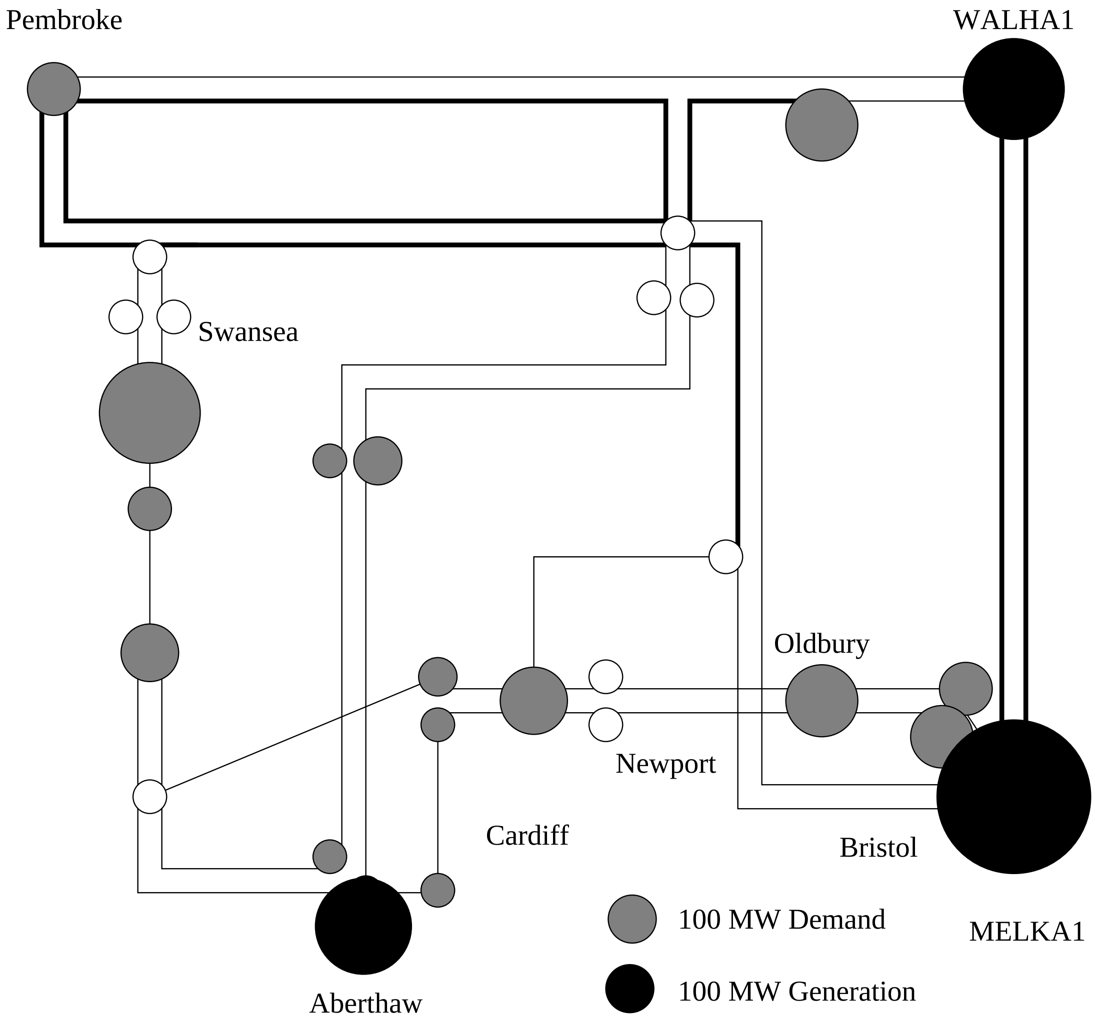

Figure 1: South Wales Region High Voltage Network

In this paper we report work showing the combination of GAs and greedy optimisers described in [Lan95] being applied to the South Wales region of the National Grid. Other work on the South Wales problem using Genetic Programming is reported in [LT97].

Section 2 describes the South Wales region high voltage network, Section 3 describes the fitness functions used in Sections 4 and 5 in which low cost schedules are produced, firstly without the network resilience requirement and secondly (in Section 5) when including it. Details of the GA parameters used are given in Appendix A.

2 The South Wales Region of the UK Electricity Network

The South Wales region of the UK electricity network carries power at 275K Volts and 400K Volts between electricity generators and regional electricity distribution companies and major industrial consumers. The region covers the major cites of Swansea, Cardiff, Newport and Bristol, steel works and the surrounding towns and rural areas (see Fig. 1). The major sources of electricity are infeeds (2) from the English Midlands, coal fired generation at Aberthaw, nuclear generation at Oldbury and oil fired generation at Pembroke. Both demand for electricity and generation change significantly through the year (See Figs. 2 and 3).

The representation of the electricity network used in these experiments is based upon the engineering data available for the physical network; however a number of simplifications have to be made. Firstly the regional network has been treated as an isolated network; its connections to the rest of the network

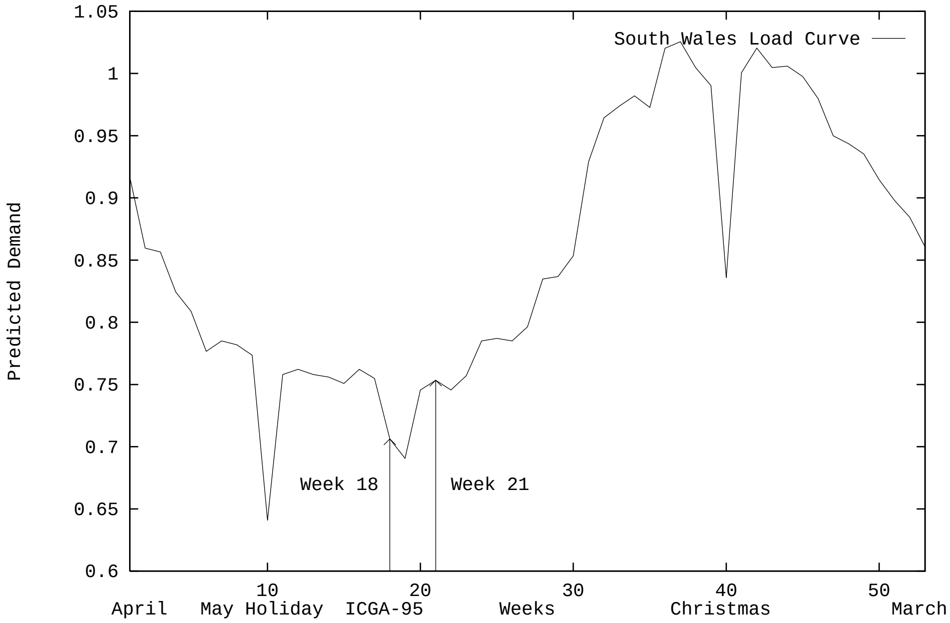

Figure 2: Predicted Demand in South Wales Region

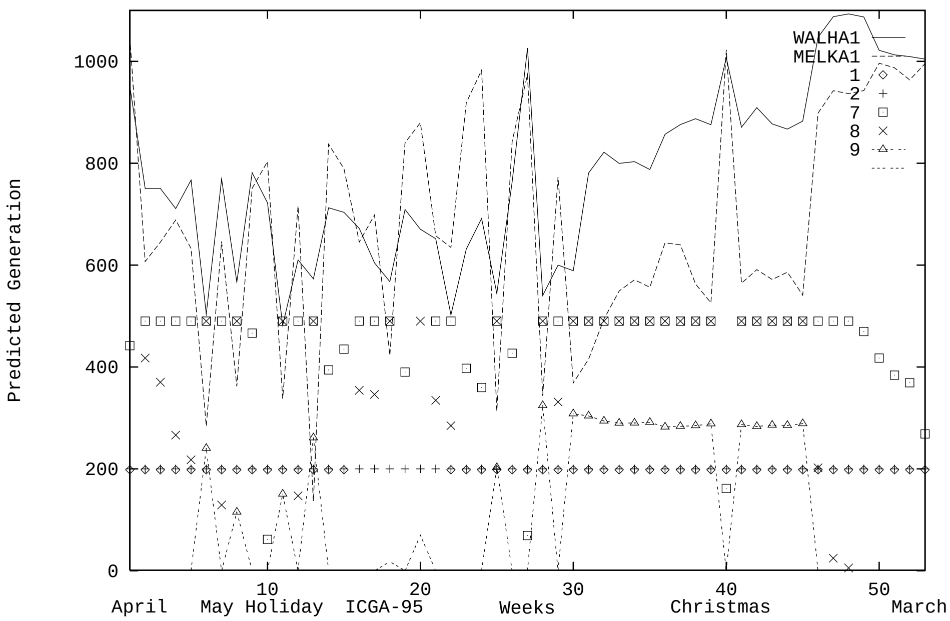

Figure 3: Predicted Generation in South Wales Region

have been modelled by two sources of generation connected by a pair of low impedance high capacity conductors. Secondly the physical network contains short spurs run to localised load points such as steel works. These "T" points have been simplified (e.g. by inserting nodes into the simulation) so all conductors connect two nodes. The industry standard DC load flow approximation is used to calculate power flows through the network.

In the experiments reported in this paper the maintenance planning problem for the South Wales region has been made deliberately more difficult than the true requirement. In these experiments:

- All lines must be maintained during the 53 week plan (1995 had 53 weeks rather than 52). Typically about one third of the lines are maintained in any one year.
- All maintenance takes four weeks. Typically scheduled outage of a line is between a week and $1\frac{1}{2}$ months.
- We consider two versions of the problem. In the first the requirement that the network must remain secure against faults in the network was ignored. To compensate for this, the problem was artificially increased in difficulty by reducing the capacity of all lines by 50%. When satisfactory solutions to this version of the problem have been evolved, we then considered evolving maintenance plans which contain a degree of network resilience (cf. Section 5). In this second problem the true conductor rating were used.

Considering potential network faults is highly CPU intensive and so the Genetic programming approach described in [LT97] did not attempt to solve the second version of the South Wales problem.

3 The Fitness Function

The fitness function used by both GA and genetic programming (GP) approaches to scheduling maintenance in the South Wales region is based on that used in the four node problem. To summarise;

- The demand for electricity for the coming year is predicted (cf. Fig. 2).
- The cheapest electrical generators are scheduled to meet this demand (cf. Fig. 3).

  To optimise costs across the whole network, these two steps are performed nationally. Imbalances between demand and generation in each region result in power transfers between regions and in the case of the South Wales region for this year, large inward power flows were planned (cf. the top two curves in Fig. 3).
- The fitness function includes a maintenance benefit for each maintenance task completed (i.e. for each line maintained for at least four consecutive weeks).
- If maintenance were to cause power flows to be redistributed in such a way as to overload part of the network it would be necessary to re-arrange the network to avoid this. In practice this may entail changes to the pattern of electricity generation. As we started with the cheapest generators, this will be more expensive. In practice this is the major cost of planned network maintenance. The fitness function models this cost with a penalty given by the amount each line is overloaded.
- Large penalties are also included to punish schedules which isolate nodes. (However nodes with no demand or generation connected directly to them are excluded from these penalties). When nodes are isolated or the network is split, the penalty for each node in the disconnected part of the network are summed and then the usual loadflow and fitness calculations are performed on the remaining connected network (i.e. the connected network containing the most nodes). In the South Wales problem there is no separate penalty for splitting the network.

$$
\begin{aligned}
\text{cost} = & \sum_{\text{target}} K_t \times (1-\text{maintenance\_scheduled}) \\
& + \sum_{\text{weeks}} \left(\begin{array}{l} S_2 \times \sum\limits_{\text{isolated\_nodes}} \mid \text{demand} \vee \text{generation} \mid \\ + \\ \sum\limits_{\text{remaining\_lines}} \begin{array}{c} \textbf{If } |\text{flow}| > \text{rating } \textbf{then} \\ |\text{flow}| - \text{rating} \end{array} \end{array} \right) .
\end{aligned} \tag{1}
$$

Equation 1 gives the fitness function used in Section 4, i.e. excluding consideration of contingencies (network faults). In Equation 1 the first summation is over all target maintenance (n.b. the trial plan's cost is increased by K_t if the corresponding maintenance is not scheduled). The second outer summation being over each week of the maintenance plan; the first inner one, being over all isolated nodes, and the second, over all lines in the network.

Where contingencies are to be considered, the second part of Equation 1 is extended to include a penalty for each contingency which is calculated as if the contingency had occurred but weighted by an arbitrary scaling factor based upon the severity of the contingency (cf. Equation 2). During the contingency part of the fitness calculation the post fault rating on all the conductors is used. This is about 17% higher for each conductor than the corresponding rating used in the first part.

$$
\begin{aligned}
\text{cost} = & \sum_{\text{target}} K_t \times (1-\text{maintenance\_scheduled}) \\
& + \sum_{\text{weeks}} \left(\begin{array}{l} S_2 \times \sum\limits_{\text{isolated\_nodes}} \mid \text{demand} \vee \text{generation} \mid \\ + \\ \sum\limits_{\text{remaining\_lines}} \begin{array}{c} \textbf{If } |\text{flow}| > \text{rating } \textbf{then} \\ |\text{flow}| - \text{rating} \end{array} \end{array} \right) . \\
& + \sum_{\text{contingency}} \text{scaling} \times \sum_{\text{weeks}} \left(\begin{array}{l} S_2 \times \sum\limits_{\text{isolated\_nodes}} \mid \text{demand} \vee \text{generation} \mid \\ + \\ \sum\limits_{\text{remaining\_lines}} \begin{array}{c} \textbf{If } |\text{flow}| > \text{rating}_{+17\%} \textbf{then} \\ |\text{flow}| - \text{rating}_{+17\%} \end{array} \end{array} \right)
\end{aligned} \tag{2}
$$

For the South Wales problem the same values of K_t and S_2 as the four node system where used. I.e. K_t is 4,000 MW and $S_2 = 5$. [Gor95] verified the values used for the four node problem are applicable to the South Wales region.

The South Wales network is obviously considerably larger and more complex than the four node network. Correspondingly, both the time to perform network connectivity checks and to calculate powerflows within the network are greatly increased. Since both are fundamental parts of some greedy schedulers and of the fitness function, the run time of the South Wales problem is far greater than that of the four node problem.

In later experiments the sparse matrix package SDRS2 was used to reduce the CPU requirements of the loadflow calculation. While successful, the reduction was only by a factor of two. A larger ratio is anticipated with larger networks, such as the National Grid. More success was had with the "rankone" technique, which in some cases produced a ten fold speed up in Section 5.

Analysis indicates that performing the network connectivity checks can be a bottle neck. A significant reduction in run time was achieved by replacing the original algorithm. Additional reduction was achieved by maintaining a cache of previously evaluated network connectivities. It is felt that caching techniques could produce worthwhile reductions in run time if also applied to the loadflow calculation.

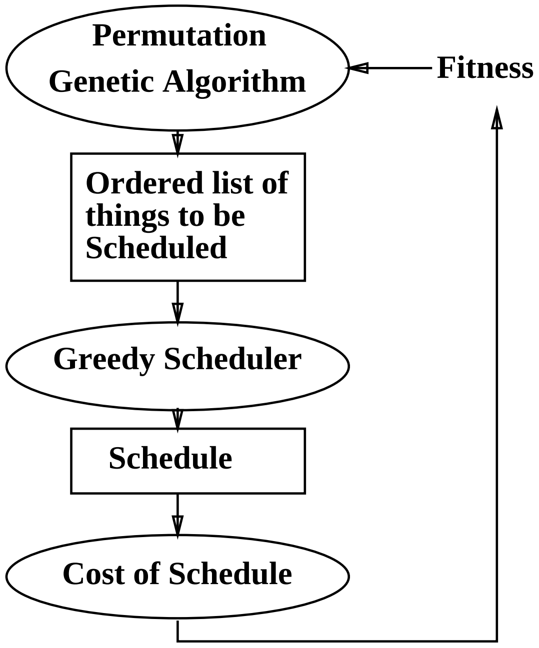

Figure 4: Hybrid GA and "greedy optimiser"

4 South Wales Problem without Contingencies

4.1 Greedy Scheduler

Briefly the greedy scheduler (cf. Fig. 4) is presented with each line in the order determined by the GA chromosome. For each line it determines which four weeks will cause the overall cost of the schedule to increase the least (or decrease the most). The line is then scheduled in those four weeks and the next line is processed.

This greedy scheduler is the minimum increase cost greedy scheduler that solved the four node problem (Heuristic 4 in [Lan95]) but modified to take into account maintenance takes four weeks rather than one.

4.2 Results

Fig. 5 contains the least cost schedule produced using the minimum increase cost greedy scheduler using the same GA that was used to solve the four node problem with the same population size (cf. Appendix A).

Not only is this a low cost schedule (it has a cost of 616MW weeks) but it is a good one in that it contains several "sensible" decisions, for example performing maintenance on lines and the transformers which supply them simultaneously. For example line A819 and transformer F819, which supply Swansea, are both maintained in weeks 4–7.

4.3 Discussion

While still better performance could be obtained using modified greedy schedulers, these results were sufficiently encouraging as to warrant including the fault resilience requirements by explicitly including contingencies. This is described in the next section.

The run time of these experiments (which did not use SDRS2 or rankone modification algorithms) are consistent with predictions made by Table 12 in [Lan95, page 148]. It is encouraging that reasonable results have been obtained using essentially the same technique as was used on the four node problem, despite this problem's search space being 2.1 10^{60} fold larger.

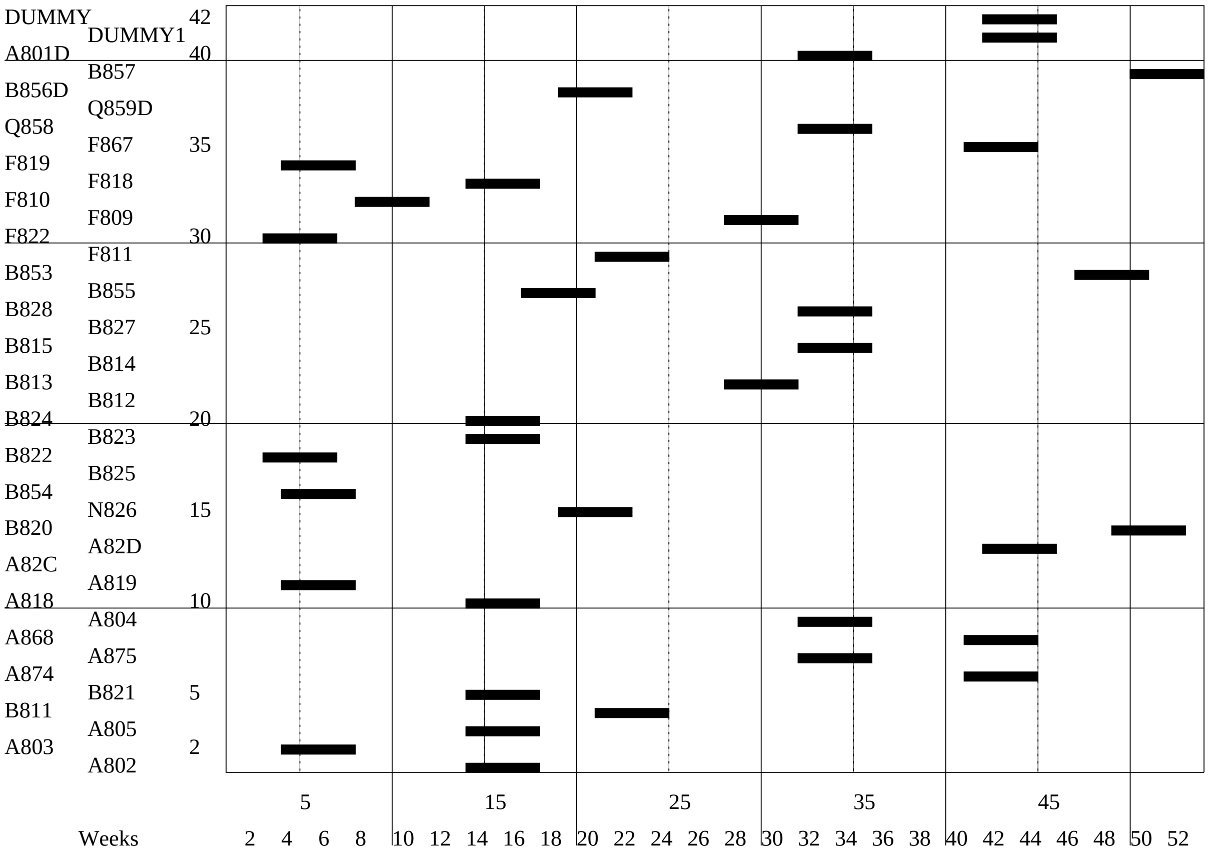

Figure 5: Schedule for South Wales without Contingencies

Another heuristic, which produced improved results, replaced the line overloading component of the minimum increase in cost greedy scheduler with the square of the powerflow. It is felt this produced improved performance because in the case of the South Wales network, unlike the four node problem, many conductors are not near their ratings and so the minimum increase in cost heuristic is relatively independent of changes to their powerflows. By using power flow squared, the greedy scheduler becomes more sensitive to changes but particularly to changes in lines near their capacities.

On this version of the South Wales problem and using a population with seeded with the same (and other) heuristics, GP was able to find schedules with a cost of 388MW weeks albeit with far greater resources (e.g. the population was 1000) [LT97].

[Gor95] used a linear chromosome with non-binary alleles [Ros94] to solve the four node problem but was less successful on this version of the South Wales problem.

5 South Wales Problem with Contingencies

In the experiments reported in this section all 52 contingencies (i.e. potential faults which prevent a line carrying power) were included. This greatly increases run time. The 52 contingencies comprise; 42 single fault contingencies (one for each line in the network and twelve double fault contingencies (for cases where lines run physically close together, often supported by the same towers). The scaling factors for the single fault contingencies were 0.1, with 0.25 for the double fault contingencies, except the most severe where it was 1.0.

5.1 Greedy Scheduler

A number of "Greedy optimisers" were tried with the permutation GA (described in Appendix A). Satisfactory schedules where obtained when considering all 52 contingencies, using a two stage heuristic.

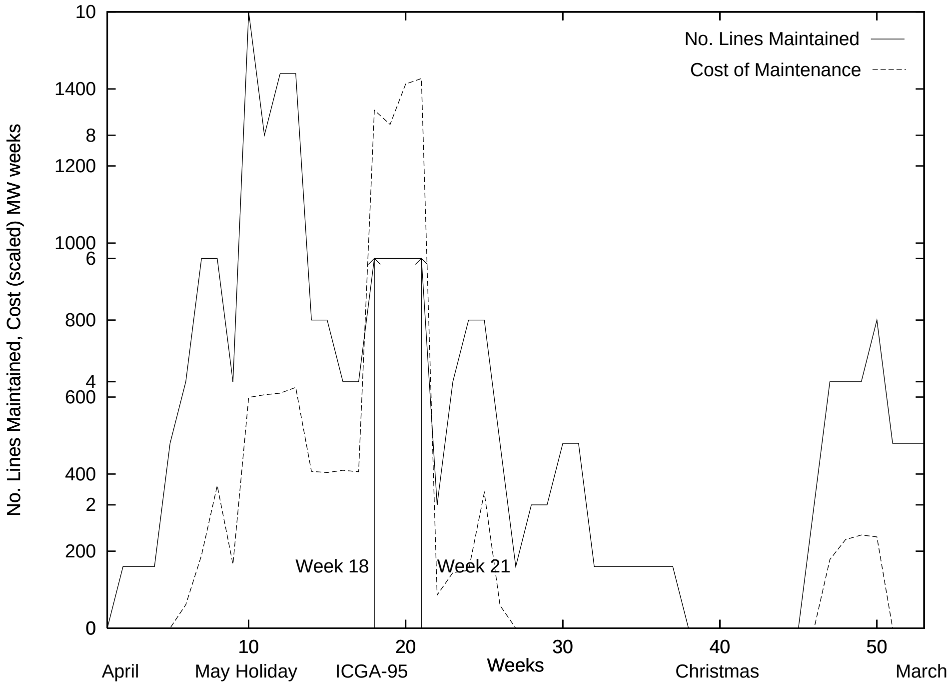

Figure 6: South Wales with Contingencies Maintenance Schedule

The choice of which four weeks to schedule a line's maintenance in is performed in two stages. In the first stage, we start with the maintenance schedule constructed so far and calculate the effect on network connectivity of scheduling the line in each of the 50 available weeks if each of the possible contingencies occurred simultaneously with the line's maintenance. (NB one contingency at a time). To reduce computational load, at this stage we only consider if the result would cause disconnections, i.e. we ignore line loadings. In many cases none of the contingencies would isolate nodes with demand or generation attached to them and all 50 weeks are passed onto the second stage. However as the maintenance schedule becomes more complete there will be cases where one or more contingencies would cause disconnection in one or more weeks. Here the sum over all contingencies, of the load disconnected, weighted by the each contingency's' scaling factor, is calculated and only those weeks where it is a minimum are passed to the second stage.

In the second stage the modified version of the minimum increase in line cost heuristic, used in Section 4.1, is used to select the weeks in which to perform maintenance. NB the second stage can only choose from weeks passed to it by the contingency pre-pass. By considering contingencies only in the first stage, run time is kept manageable.

5.2 Results

The least cost schedule found in generation 18 after 78 minutes of one run of QGAME on a 50 MHz SUN server 690MP is shown in Fig. 6. The cost of the schedule is dominated by the contingency costs incurred in weeks 18–21. Four of the six lines maintained during these four weeks will always incur contingency costs, because once taken out of service nearby demand nodes are only connected to the remaining network by a single line (see Fig. 7). By placing these in weeks 18–21 their contingency costs have been minimised as these weeks correspond to minimum demand.

5.3 Discussion

While the schedules found do represent low cost solutions to the problem presented to the GA, unlike in the four node problem, they may not be optimal. With problems of this complexity, low cost solutions are all that can reasonably be expected. Computer generated schedules have been presented to NGC's planning engineers, while expressing some concerns, they confirm they are reasonable.

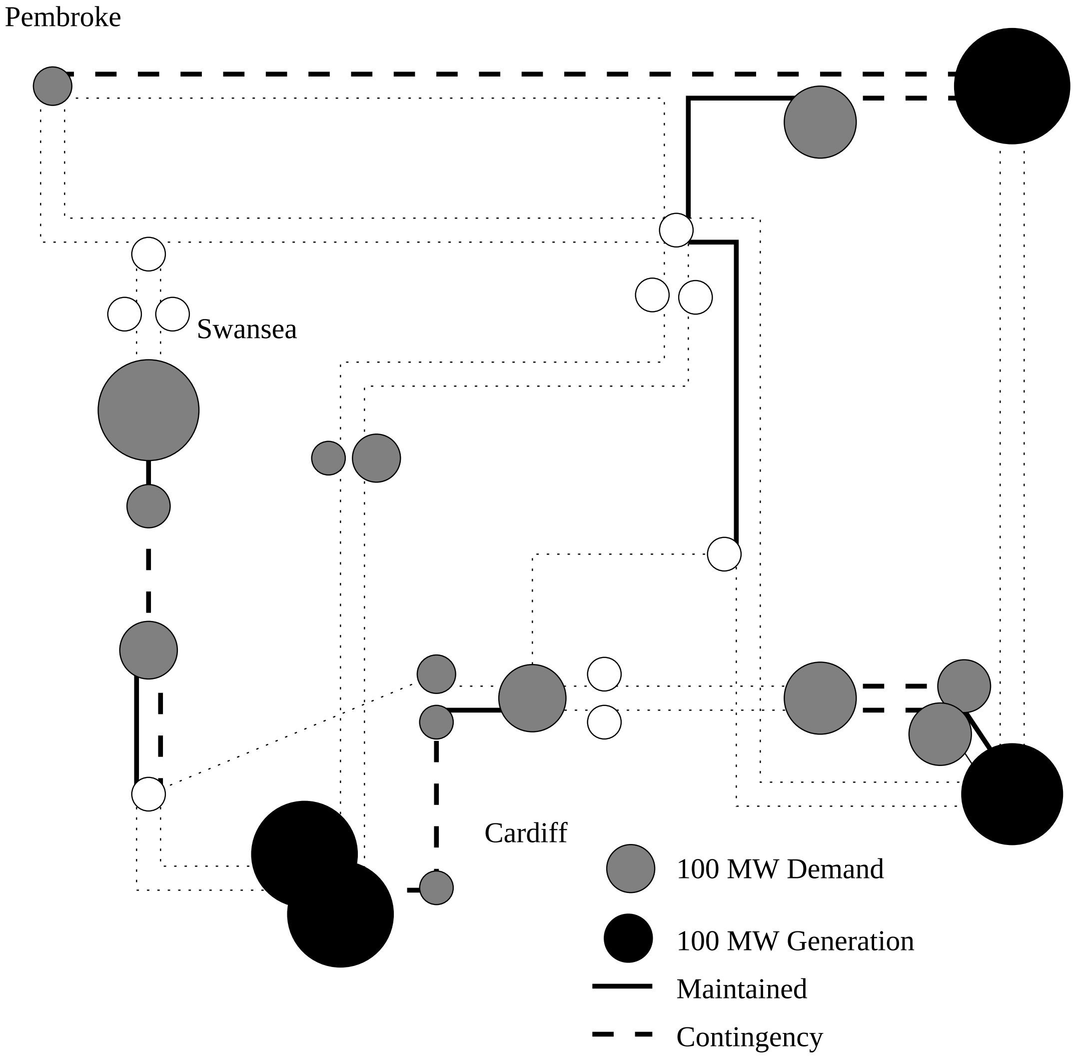

Figure 7: Weeks 18–21 of South Wales with Contingencies Maintenance Schedule

The two principle concerns were; firstly the GA appears to have been concerned to optimise the contingency costs associated with small demand nodes with only two connections. Since contingency costs associated with them are inevitable, currently planning engineers do not concentrate on minor improvements to them.

The greedy optimiser approach, schedules each line one at a time. There is a concern that there may be circumstances where the optimal solution requires two lines to be maintained at times which do not correspond to the local best for either of them. For example suppose it is planned that a generator will be run at half power for eight weeks and that it is connected by several lines, two of which are to be maintained. It would make sense to maintain them one after the other but the greedy scheduler could decide the best time to schedule one of them (looking just at the best for that line) was in the second of these eight weeks, so preventing the optimal schedule from being found. This circumstance does not arise in the South Wales problem, investigations of the national network are required to see if it occurs elsewhere and how serious a problem it is.

Since this work began. NGC's manual methods have been augmented by a computerised single pass hill climbing technique, which has produced considerable financial savings.

6 Conclusions

We have described the complex real world problem of scheduling preventive maintenance of a very large electricity transmission network. The combination of a GA and hand coded heuristic has been shown to produce acceptable schedules for a real regional power network when including consideration of network robustness to single and double failures. However consideration of such contingencies considerably increases run time and so the production of schedules for the entire national network has not yet been demonstrated.

The time taken to perform GA fitness evaluations and with it program run time, grows rapidly with problem size and number of potential failures that must be considered. It is anticipated that running on parallel machines will be required to solve the national problem using a GA or GP approach. However there are a number of techniques which could be used to contain run time.

A Using QGAME

QGAME [RT94] was used in all the GA experiments on grid scheduling problems. In every case the genetic algorithm contain a single population of 20 individuals, each contained a chromosome which consisted of a permutation of the lines to be scheduled. Permutation crossover and default parameters were used (see Table 1).

Table 1: QGAME genetic algorithm parameters

Name	Value	Meaning
Problem_type	Combinatorial	Combinatorial crossover and mutation operators are used on the chromosome, which holds a permutation of the lines to be scheduled
GeneType	G_INT	Line identification number
POPULATION	20	Size of population
CHROM_LEN	8	Eight genes (lines) per chromosome
OptimiseType	Minimise	Search for lowest fitness value
PCROSS	0.6	On average each generation $0.6 \times$POPULATION/2 pairs of individuals are replaced by two new ones created by crossing them over
PMUT	0.03	On average each generation $0.03 \times$POPULATION individuals are mutated by swapping two genes
Selection	Truncated	The best TRUNCATION$\times$POPULATION of the population are passed directly to the next generation. The remainder are selected using normalised fitness proportionate (roulette wheel) selection. Fitness is linearly rescaled so the best in the population has a normalised fitness of 1.0 and the worst zero
TRUNCATION	0.5	The best 50% passed directly to the next generation
POOLS	1	Single population
CHROMS	1	Individuals use single chromosome
MAX_GEN	100	Run for 100 generations

Acknowledgements

This work was carried out at UCL Computer Science department, while W. B. Langdon was funded by the EPSRC and The National Grid Company plc.

I would like to thank Ursula Bryan and Daniel Waterhouse for the South Wales problem and for the use of fast loadflow calculation software.

The four_node problem definition and QGAME are available via anonymous ftp, site `cs.ucl.ac.uk` directory `genetic/four_node`.

References

[Dun93] R. M. Dunnett. A proposal to use a genetic algorithm for maintenance planning. PSBM note, National Grid, Technology and Science Laboratories, July 1993. Private communication.

[Gor95] T. G. W. Gordon. Schedule optimisation using genetic algorithms. Master's thesis, University College, London, October 1995.

[Hol92] John H. Holland. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence.* MIT Press, 1992. First Published by University of Michigan Press 1975.

[Lan95] W. B. Langdon. Scheduling planned maintenance of the national grid. In Terence C. Fogarty, editor, *Evolutionary Computing*, number 993 in Lecture Notes in Computer Science, pages 132–153. Springer-Verlag, 1995.

[LT97] W. B. Langdon and P. C. Treleaven. Scheduling maintenance of electrical power transmission networks using genetic programming. In K. Warwick, R. Aggarwal, and A. Ekwue, editors, *Applied Artificial Intelligence in Power Systems*, chapter 10. IEE, 1997.

[Ros94] Peter Ross. *About PGA 2.8*, 1994. Available via ftp `ftp.dai.ed.ac.uk` directory `pub/pga-2.8`.

[RT94] J. L. Ribeiro Filho and P. Treleaven. GAME: A framework for programming genetic algorithms applications. In *Proceedings of the First IEEE Conference on Evolutionary Computing – Proceedings of the 1994 IEEE World Congress on Computational Intelligence*, volume 2, pages 840–845, Orlando, USA, 26 June-2 July 1994. IEEE Press.