Scheduling Planned Maintenance of the National Grid

Research Note RN/3/95

To be presented at: The AISB Workshop on Evolutionary Computation, 3–4 April 1995

W. B. Langdon, Email: W.Langdon@cs.ucl.ac.uk

Dept. of Computer Science,
University College London,
Gower Street, London WC1E 6BT, UK

Switchboard tel:	+44 (0) 171 387 7050 xtn. 3701
Dept. Office Tel:	+44 (0) 171 380 7214
Fax:	+44 (0) 171 387 1397

Modifications:	WBL	8 March 1995	Minor corrections
	WBL	12 January 1995	Version submitted to AISB

Abstract

The National Grid Company Plc. is responsible for the maintenance of the high voltage electricity transmission network in England and Wales. It must plan maintenance so as to minimize costs taking into account:

- location and size of demand,
- generator capacities and availabilities,
- electricity carrying capacity of the remainder of the network, i.e. that part not undergoing maintenance.

This complex optimization and scheduling problem is currently performed manually by NGC's Planning Engineers. Computerized viability checks are performed on the schedules they produce. NGC's Technology and Science Laboratories and UCL aim to generate low cost schedules using Genetic Algorithms. It is hoped this will aid the Planning Engineers.

This paper reports work in progress. So far:

- A four node test problem has been identified
- A fitness function has been devised
- To date work has concentrated upon devising a representation based upon "Greedy" Optimizers.
- The best of these has been incorporated in the QGAME genetic algorithm programming environment and optimal solutions have been readily found.

We are now considering how to scale up. The computational complexity of the best of the greedy optimizers is a concern. Our plans are to consider alternative approaches (such as expansive coding and Genetic Programming) before moving on to a significant section of the whole of the National Grid, such as South Wales.

Keywords: Genetic Algorithms, Scheduling, Optimization, Maintenance Planning, Electricity Transmission Networks.

1 Introduction

In England and Wales electrical power is transmitted by the high voltage electricity transmission network known as the National Grid. It is highly interconnected and carries large power flows. The National Grid 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 optimization 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 minimized.

This paper describes work on the use of Genetic Algorithms (GAs) on this problem to date and how we intend to continue our work.

The work to date is based upon a simplified four node test problem. Section 3 presents the fitness function devised for the four node problem and indicates changes that may be required when considering larger networks. Section 4 describes the four node problem in detail. Section 5 considers various possible Genetic Algorithms and explains why work has concentrated upon using "Greedy" optimizers. Section 6 details the four optimizers that have been tried and the results achieved. The best of these readily produces optimal solutions.

We are now considering how to scale up. The computational complexity of the best of the greedy optimizers is a concern. We have started considering alternative approaches (such as expansive coding and Genetic Programming).

2 Approximating Replacement Generation Costs

NGC use computer tools for costing maintenance schedules however because of their computational complexity, it was felt that these where unsuitable for providing the fitness function. Instead our fitness function is partially based upon estimating the replacement generation costs that would occur if a given maintenance plan were to be used. The estimate is made by calculating the power flows assuming there will be no replacement generation and there are no power flow limits. In practice alternative generators must be run to reduce the power flow through over loaded links in the network. The estimated cost of the alternative generators is assumed to be proportional to the amount of replacement generation needed which is crudely assumed to be equal to the excess power flow.

As a first approximation: issues concerning security of supply following a fault are not covered. We plan to increase the realism of the GA as we progress. Other costs (e.g. maintenance workers' overtime, travel) are not considered as they are of secondary importance.

3 The Fitness Function

The GA's fitness function is composed of two parts; a benefit for performing maintenance plus penalties for exceeding link ratings, isolating nodes and splitting the network.

3.1 Maintenance Benefits

The maintenance requirements of the different components of the National Grid vary both in terms of the number of weeks required to perform them and their urgency. It may be advisable to hold over less urgent maintenance until the following year. We model the urgency by allocating each maintenance activity a benefit term in the fitness function.

During the plan, each part of the network should be maintained a small number of times, typically once. Each maintenance activity takes a few weeks and has a particular fitness benefit. Should a trial maintenance plan schedule the plant item for maintenance for the required number of weeks, that trial plan's fitness is improved by the maintenance benefit. Should more than one maintenance period be required they must not overlap, however there may be an additional fitness benefit if they run consequently (perhaps of about 20%). Once the target number of weeks of maintenance has been reached, there is no additional benefit or penalty from additional maintenance. The total benefit is obtained by summing across all plant items for the whole year.

In the rest of this paper it is assumed that each link in the network can be maintained independently, however there are some cases where this is not so (notably "Tee" junctions). These require several items to be taken out of service simultaneously. As there are relatively few of these they can be treated as special cases.

3.2 Over Loading Costs

For each week of the year the available generators are loaded in strict least cost basis given a predicted demand for that week and the notified generator availabilities. This is known as the merit order dispatch. It is fixed and therefore the same for all trial solutions.

For each week of each plan a "DC load flow"[1] analysis is performed based on the predicted demand and the merit order dispatch described above. An over loading cost proportional to excess power flow will be calculated for each plant item. The total over loading costs will be the sum over each week of the maintenance plan and over all plant items of the over loading costs. The cost will start to become incurred should loading exceed normal operating limits, which are 85% of nominal post fault continuous rating[2].

3.3 Avoiding Isolating Nodes or Splitting the Network

From an operational point of view, no acceptable maintenance schedule would ever isolated a generational or demand node from the rest of the network or split the network. However the GA fitness function must be able to cope with every schedule that is generated. The "DC load flow" algorithm cannot cope with either as they require it to invert a singular matrix. Therefore the fitness function looks for these conditions and defines a fitness for them without calculating power flows.

As both represent highly unfit solutions, weeks of a schedule that cause either contribute a high penalty to the schedule's whole fitness. The penalty for each isolated node is proportional to the load or expected generation at that node. The penalty for a network split is even more severe; it is proportional to the total load across the whole network in that week.

[1] DC load flow analysis is an industry standard approximation.

[2] The nominal post fault continuous rating is figure usually quoted for transmission lines.

3.4 Combined Fitness Measure

The complete fitness measure is expressed in MegaWatt weeks and is given by the following formula:

$$
\begin{aligned}
cost = \quad & \sum_{target} K_t \times (1-\text{maintenance\_scheduled}) \\
+ \quad & \sum_{weeks} \left(\begin{array}{l}
\textbf{If } \text{network\_split } \textbf{then} \\
\qquad S_1 \times \text{ total\_demand} \\
\textbf{Else If } \text{isolated\_nodes} \geq 1 \textbf{ then} \\
\qquad S_2 \times \displaystyle\sum_{isolated\_nodes} \mid \text{demand} \vee \text{generation} \mid \\
\textbf{Else} \\
\qquad \displaystyle\sum_{links} \begin{array}{l} \textbf{If } |\text{flow}| > 0.85 \times \text{rating } \textbf{then} \\ \qquad |\text{flow}| - 0.85 \times \text{rating} \end{array}
\end{array} \right)
\end{aligned}
$$

The first summation being 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 links in the network.

For the four node system (next section) all the constants K_t are equal and S_1 and S_2 are equal to each other. This may not be the case in the full system.

4 The Four Node Demonstration Network

A small power system has been devised. The network is shown in Figure 1. As the problem's network contains four substations (i.e. nodes in the network multi-graph), it is know as the "four node problem".

This power system contains nine generators of differing capacities, availabilities and prices. The substations are linked by pairs of transmission lines (i.e. links in the network multi-graph) to form a square. The transmission lines are of two sorts; 1100MW and 600MW rating but each have the same impedance[3]. Table 1 gives the transmission line capacities. Table 2 shows the predicted demand, which varies from week to week and node to node.

The problem is to devise a nine week maintenance schedule which, maintains each line for at least one week and minimizes predicted line over loading costs whilst not splitting the network or isolating nodes in it. I.e. has minimum cost.

4.1 Determining which Generators to Use

The power flows depend on which lines are being maintained, network impedances, demand for power and which generators are running. The first item we have to determine, the second and third are given (or predicted), the last is given by the merit order method described in Section 2. In each week of the maintenance schedule, the demand for power (Table 2) is met by the

[3]This leads to certain symmetries in this problem. As true power transmission networks do not have them no attempt has been made to exploit them.

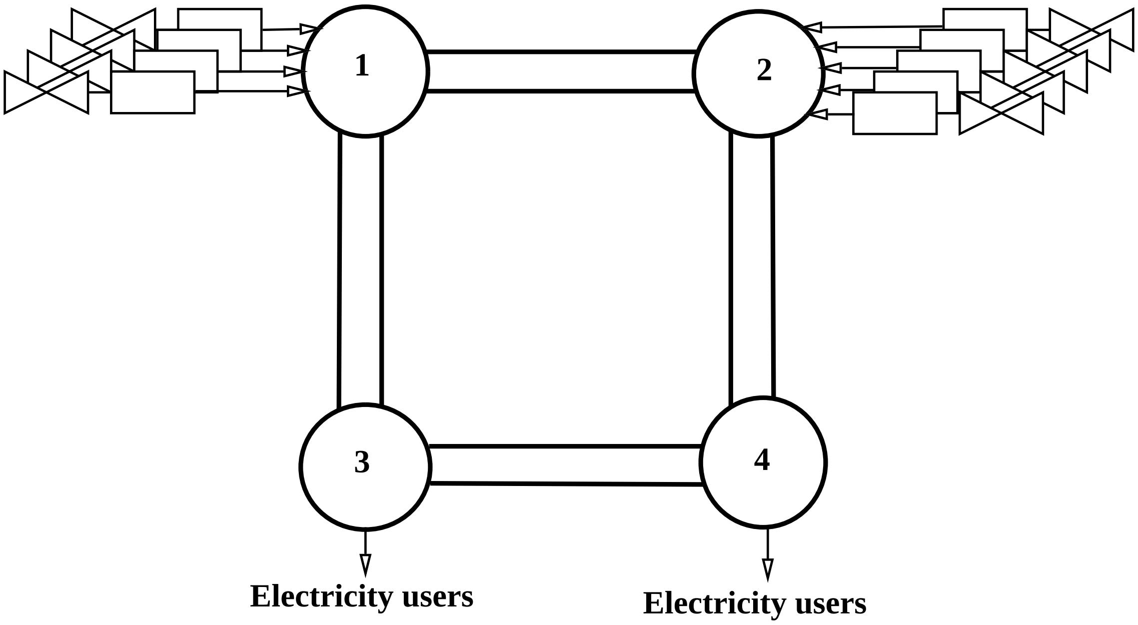

Figure 1: Four node network

Line	Start node	End node	Rating	Impedance
1	1	2	600	1
2	1	2	600	1
3	2	4	1100	1
4	2	4	1100	1
5	4	3	600	1
6	4	3	600	1
7	3	1	1100	1
8	3	1	1100	1

Table 1: Line Capacities (MW) and Impedances

Week	Demand				Node 1 Generators				Node 2 Generators				
	Node 3 MW	Node 4 MW	Total MW		1.1	1.2	1.3	1.4	2.1	2.2	2.3	2.4	2.5
				Price	1.0	2.0	3.0	4.0	3.5	4.5	5.5	6.5	7.5
				Cap	500	500	500	500	400	400	400	400	400
1	1200	1300	2500		x								
2	1200	1600	2800			x							
3	1000	2000	3000				x						
4	1100	2300	3400					x					
5	1100	2400	3500						x				
6	1000	2500	3500							x			
7	900	2600	3500								x		
8	900	2700	3600									x	
9	800	3000	3800										
				x indicates the generator is not available to generate									

Table 2: Predicted Demand and Generator Price, Capacity and Availability

cheapest available generators. Table 3 shows which generators run, and how much they generate in each week. From this the total generation at each node and for each week is calculated.

We now have enough information to calculate the power flows in every power transmission line in the network in every week for any given maintenance schedule. Once we have established the constants K and S, the fitness of *cost* of the schedule can be readily calculated from the known power flows.

Week	Node 1 Generators					Node 2 Generators						Total
	1.1	1.2	1.3	1.4	Total	2.1	2.2	2.3	2.4	2.5	Total	
1	x	500	500	500	1500	400	400	200	0	0	1000	2500
2	500	x	500	500	1500	400	400	400	100	0	1300	2800
3	500	500	x	500	1500	400	400	400	300	0	1500	3000
4	500	500	500	x	1500	400	400	400	400	300	1900	3400
5	500	500	500	500	2000	x	400	400	400	300	1500	3500
6	500	500	500	500	2000	400	x	400	400	300	1500	3500
7	500	500	500	500	2000	400	400	x	400	300	1500	3500
8	500	500	500	500	2000	400	400	400	x	400	1600	3600
9	500	500	500	500	2000	400	400	400	400	200	1800	3800
	x indicates the generator is not available to generate											

Table 3: Merit Order Generator Loading (MW)

4.2 Determining Fitness Function Parameters K and S

The constants of proportional, K and S need to be chosen so that the maintenance benefit dominates. This was done by calculating the line costs associated with many feasible schedules (using the heuristic given in Section 6.1). This gave the highest feasible line cost as being 3870MW weeks. K was set to 4,000 MW weeks, so that any schedule which failed to maintain even one line would have a higher cost that the worst "reasonable" schedule which maintains them all.

S was chosen so that isolating of any node would have a high cost. The lowest demand/generation on a node is 800 MW, setting S to five ensures that any schedule which isolates any node will have a cost higher the worst "reasonable" schedule.

5 The Chromosome

Various ways of representing within the GA the evolved schedules, were considered. When choosing one, we considered not just the four node problem but also the complete National Grid which contains about 300 nodes and about 500 transmission lines which are maintained over a 52 week plan. Table 4 summarizes the representations considered.

Organization	Size	gene	Expected Benefit	Disadvantage
Linear	25,000	bit	Simple	Crossover disruptive, Too long, No structure, Sparse
2D	500 by 52	bit	Simple, Obvious temporal structure	Crossover disruptive, Still very big, No grid structure, sparse
Graph	500 by 52	bit	Realistic	Crossover disruptive, Complex, Little preceding work, Sparse
TSP	25,000	link	Widely studied	"By edge" costs, Sparse
Genetic Programming		S-expression		Complex, Little preceding work
Expansive Coding	$25,000 \times n$	complex	Crossover friendly	Complex, Little preceding work
"Greedy Optimisation"	500	plant id	Good crossover, Compact	Fitness evaluation expensive

Table 4: Possible Chromosomes for the National Grid Scheduling Problem

In the linear, 2D and TSP structures every attempt should be made to keep electrically close plant items close to each other in the representation. The graph representation has the advantage that this can be naturally built into the chromosome.

With a simple linear chromosome, Holland or multi-point crossover, would appear to be too disruptive to solve the problem. A simple example illustrates this. Suppose lines L_1 and L_2 are close together. It will often be the case that good solutions maintain one or the other but neither at the same time. But crossover of solutions maintaining L_1 but not L_2 with those maintaining L_2 but not L_1 will either make no change to L_1 and L_2 or not maintain either or maintain both. The first makes no change to the fitness whilst the second and third make it worse. That is, we would expect crossover to have difficulty assembling better solutions from good "building blocks". The same problem appears in the temporal dimension, and so affects the first three representations given in Table 4. This coupled with the sparseness of the chromosome (about 1 bit in 50 set) and other difficulties, has led to the decision to try a "Greedy Optimisation" representation.

6 Greedy Optimizers

The approach taken so far to solving the power transmission network maintenance scheduling problem has been to split the problem in two; a GA and a "Greedy Optimizer". The Greedy optimizer is presented with a list of work to be done (i.e. lines to be maintained) by the GA. It schedules those lines one at a time, in the order presented by the GA, using some problem dependent heuristic. A variety of heuristics have been tried of increasing sophistication and computational complexity and progressively better results.

6.1 Heuristic 1 – One Week, One Line

If we ensure a schedule maintains all lines and does not isolate nodes its fitness will be independent of K and S. Such schedules are readily devised by only maintaining one line at a time. Further such schedules are reasonable solutions to the problem (indeed it was anticipated that the best solution would be of this type). By testing a number of these we used the results to guide choice of K and S. Figure 2 shows the distribution of line costs lie between 1470 and 3870 MW weeks.

In Figure 2 etc. we plot cumulative frequency against over loading line costs (or fitness). Often GA fitness landscapes are very rough however the roughness or texture of a fitness landscape is a product of the representation used as well as the fitness itself. Using cumulative frequency makes the graph independent of representation and, in this case, makes it smooth.

The least cost solutions have the form 3,7,8,4,5,6,,1,2. Which maintains the two right hand (referring to Figure 1) high rating lines (3, 4) in the first and fourth weeks; the left hand two (7, 8) in weeks two and three; the bottom cross links (5, 6) in the fifth and six weeks, nothing in week seven and the top two cross links (1, 2) in the last two weeks. We shall return to top cross links later.

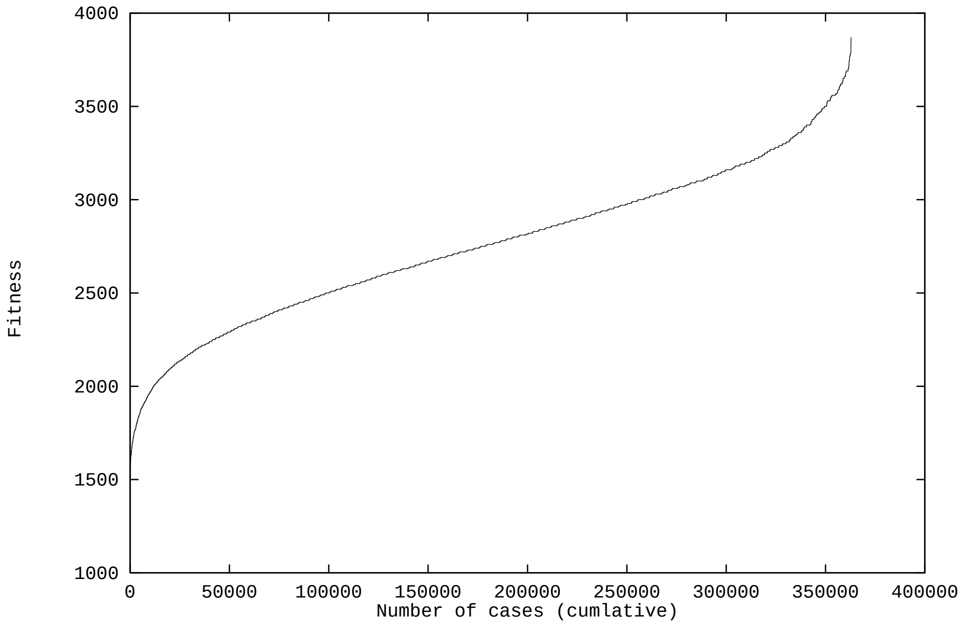

Figure 2: Cumulative frequency v. Line costs (One week, One line)

6.2 Heuristic 2 – Minimum Power Flow

A greedy optimizer was devised which scheduled the maintenance of each line in the week where the power flow through it is a minimum. In the event a tie, the earlier week is chosen. As each line is scheduled, the power flows through the rest of the network are recalculated. As Figure 3 shows, the fitness of the schedules this algorithm produces varies from 1440 to 1710 MW weeks, with 40% being 1480 or less. No schedule causes a node to be isolated or a network split.

Both the best and the worst solutions produced are better than the best and worst given by the one line, one week algorithm. In fact most solutions are better than all but 0.08% of the solutions given by the previous algorithm.

The lower costs have been achieved by maintaining lines when demand (and hence power flows) are least. Even the highest cost schedule maintains the eight lines in the first four weeks, where demand is lowest. This reduces the number of possible schedules so there are many solutions which had identical fitness. This is further increased by the symmetries of the problem. For example there are 736 lowest cost schedules. An example is 4,3,567,128,,,,,. Which maintains the two right hand (referring to Figure 1) high rating lines (4, 3) in the first and second week; the left hand two (7, 8) in weeks three and four; both bottom cross links (5, 6) in the third week and the top two cross links (1, 2) in week four.

These are good solutions which may prove to be acceptable but they are known not to be optimal. Whilst good solutions to the problem as posed, they would probably not be acceptable operationally because the maintenance is too bunched together. This reduces the resilience of the network and leaves it open to single circuit failures. On practice grounds it causes resourcing problems (i.e. may not have enough maintenance gangs available). However these issues are beyond the scope of the four node problem to address.

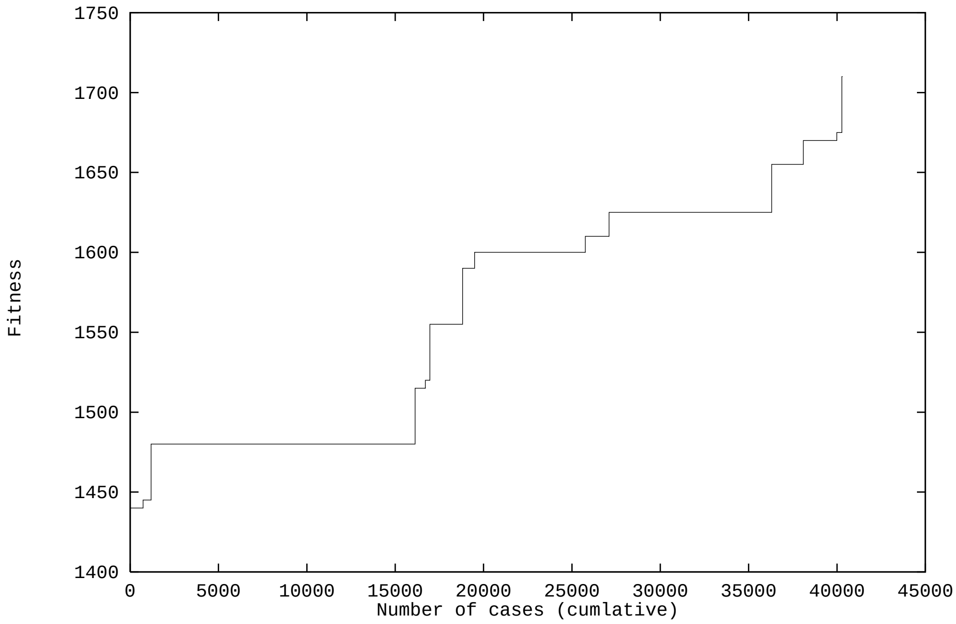

Figure 3: Cumulative frequency v. Fitness (Minimum Power Flow)

6.3 Heuristic 3 – Minimum Line Cost

The greedy lowest power flow optimizer was replaced by the more complex minimum line cost optimizer. This schedules maintenance in the week which currently contributes least to total costs. Where there is a tie, the earliest week is used.

As Figure 4 shows this algorithm produced only a marginal reduction in the best solution found (1411.67 MW weeks) and the cost of the worst schedule is actually worse (2265 MW weeks). Instead of 40% having a fitness of 1480 or better, now only 10% do.

Most schedules again concentrate maintenance in the first four weeks but some, including the best, use the fifth and sixth weeks as well. The best schedules have the form 4,68,57,3,2,1,,,. That is the two right hand high rating links (4, 3) (referring to Figure 1) are maintained in the first and fourth weeks; the left hand pair (8, 7) in the second and third weeks; the lower cross links (6, 5) in weeks two and three and the top cross links (2, 1) in weeks five and six. This more distributed maintenance plan is only a marginal improvement on the previous algorithm.

The worst schedules have the form 7,128,3,4,5,6,,, and suggest there is merit (with this type of greedy optimizer) of delaying consideration of line 1 and 2. We return to this point in the next section.

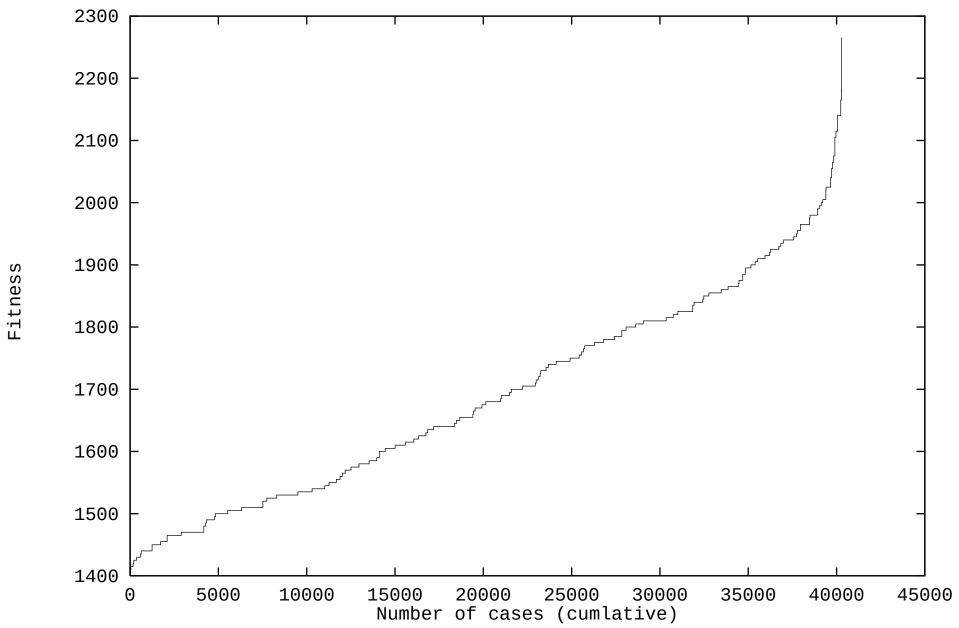

Figure 4: Cumulative frequency v. Fitness (Minimum Line Cost)

6.4 Heuristic 4 – Minimum Increase in Line Cost

The previous greedy optimizers had been based on the assumption that placing a line in maintenance was bound to increase the power flows on the remaining lines and so must increase line costs (or leave them unchanged). Whilst its is theoretically possible for the change in power flows to decrease line costs it was assumed that this would not occur in practice. As this section shows, this assumption was wrong. It is possible to schedule lines 1 or 2 so that they reduce line costs.

Briefly this can happen when one or other of the pairs of high rating links (e.g. 3 and 4) is over loaded but the rest of the network (particularly links 1 and 2) is not and links 1 and 2 are transferring power to node 2. Maintaining line 1 increases the impedance between nodes 1 and 2 and so while the power through line 2 increases, the total power flow between nodes 1 and 2 decreases. As the generation remains unchanged, the power through lines 3 and 4 decreases by the same amount, so reducing the line costs for lines 3 and 4. (The change is met by increased power flow from node 1 to 4 via node 3).

The greedy minimum line cost optimizer was replaced by the more complex least increase in line cost optimizer. This schedules maintenance in the week where maintaining it would lead to the least increase in line costs (or where there is most decrease). Where there is a tie, the earliest week is used. N.b. this optimizer looks one week ahead whereas the others make their decisions using only the links that have already been scheduled.

As Figure 5 shows this algorithm manages to find the optimal solution (1255 MW weeks) and the cost of the worst schedule (1705 MW weeks) is marginally better than the worst found by the least power flow algorithm. Now more than 70% have a fitness of better than 1480 MW weeks. Figure 6 shows the fitness distributions for all four heuristics plotted to the same scale.

Trial solutions now concentrate maintenance in the first three or four weeks except for lines 1 and 2 which are placed in the last two weeks (where demand is heaviest). The best schedules have the form 3,4,567,8,,,,1,2. That is the two right hand high rating links (3, 4) (referring to Figure 1) are maintained in the first and second weeks; the left hand pair (7, 8) in the third and fourth weeks; both the lower cross links (5, 6) in week three and the top cross links (1, 2) in weeks eight and nine. This is the best known solution to the four node problem.

The worst schedules have the form 567,3,4,8,,,,1,2. They suggest there is merit (with this type of greedy optimizer) of considering lines 5 and 6 last. We are now free to consider line 1 and 2 at any point, since the optimizer will almost certainly place them either in week eight or nine regardless.

Table 5 shows the computational complexity of the four heuristics, whose implementation has not been optimized. A large component of them all, is the DC load flow calculation, which assuming the number of nodes $\approx$ no. of lines, is $O(n^3)$. We expect to use sparse matrix and other techniques to speed up the DC load flow calculation. But even so, considerable growth in the fitness evaluation time must be anticipated, as the size of the problem (i.e. the number of nodes) is increased.

Heuristic	Order of Computational Complexity	
One week, One Line	weeks$\times$DC	weeks$\times$nodes^3
Min Power Flow	(lines+weeks)$\times$DC	(lines+weeks)$\times$nodes^3
Min Line Cost	(lines+weeks)$\times$DC	(lines+weeks)$\times$nodes^3
Min Increase Line Cost	weeks$\times$lines$\times$DC	weeks$\times$lines$\times$nodes^3

Table 5: Computational Complexity of Fitness Function

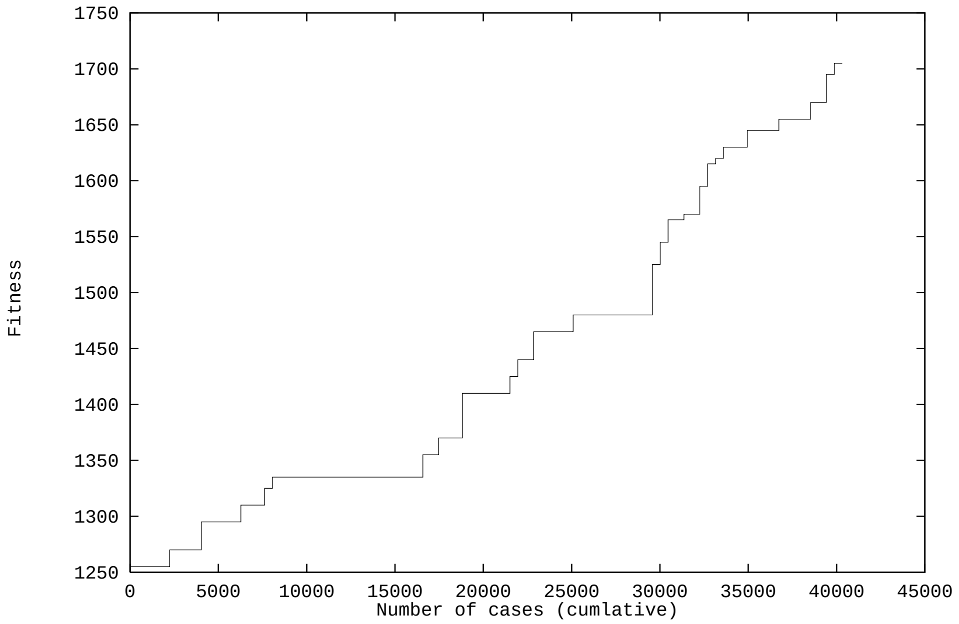

Figure 5: Cumulative frequency v. Fitness (Minimum Increase in Line Cost)

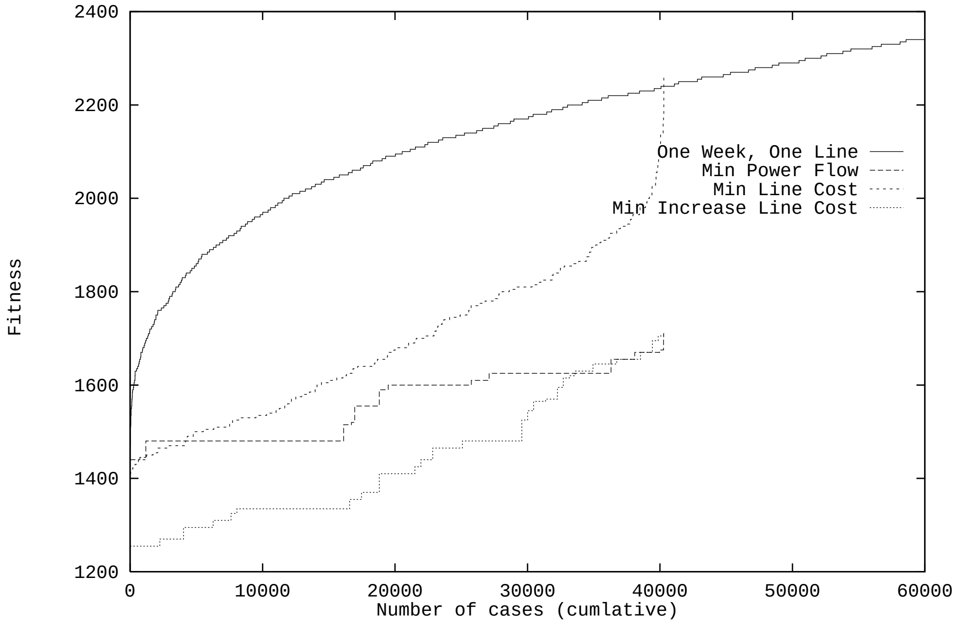

Figure 6: Cumulative frequency v. Fitness (Greedy Optimizers)

7 Using QGAME

Using a chromosome which consisted of a permutation of the lines to be scheduled, PMX crossover and a fitness function which used the minimum increase in line costs greedy optimizer (Section 6.4) the GA QGAME had no difficulty in finding the optimal solution. With a population of 20 and default parameters, optimal solutions where found by generation four in a couple of seconds.

By generation 12 approximately half the population consisted of the same optimal solution, however there were a number of slightly different chromosomes which also gave rise to optimal schedules. By generation 20 almost all the population consists of optimal solutions but there are now at least two different but equivalent optimal schedules and there are at least four different but optimal chromosomes. The population now appears to be dynamically stable, with no appreciable change in either chromosomes (genotype) or schedules (phenotype) by generation 100.

8 Discussion

Given the power of the minimum increase in line costs greedy optimizers (Section 6.4) the good results produced by the GA are not surprising. However using any GA with any of the other three greedy optimizers could not possibly have found any of the optimal solutions.

The four node problem can be represented as a search in a fitness landscape of $2^{lines \times weeks} = 2^{8\times 9} = 2^{72} = 4.7\ 10^{21}$ points. The most promising $2.1\ 10^{-14}$ of this landscape is plotted in Figure 7. This graph was produced by considering only schedules which did not do more maintenance than required[4]. This was felt to be sure to contain the global optima. Whilst, as Section 6.4 shows, lower cost solutions might be found by gratuitously taking lines out of service when they don't require maintenance, this would never happen in practice as this reduces the security of supply.

The four greedy optimizers force the GA to consider only solutions where each line is maintained exactly once. By itself this reduces the search space by a factor of $9.1\ 10^{-15}$, by using heuristics to guide where to place the eight weeks of maintenance the total search space is reduced by $8.5\ 10^{-18}$. However these heuristics require computation and (as the first three show) can misguide the search.

In addition to reducing the size of the search space, the successful greedy algorithm, considerably enhances the chance of finding the optimal solution because there are many (2240) solutions which produce it. This is a significant fraction of the search space and means a GA is unnecessary. Whilst some of the multiple optimal solutions arise from the symmetry of the original problem[5] others come from the power of the greedy algorithm itself.

[4] Fitness values above 10,000 MW weeks is shown dotted as only the maximum values were recorded. The values below 10,000 are plotted on an expanded scale in Figure 8.

[5] In the problem search space there are 32 optimal solutions.

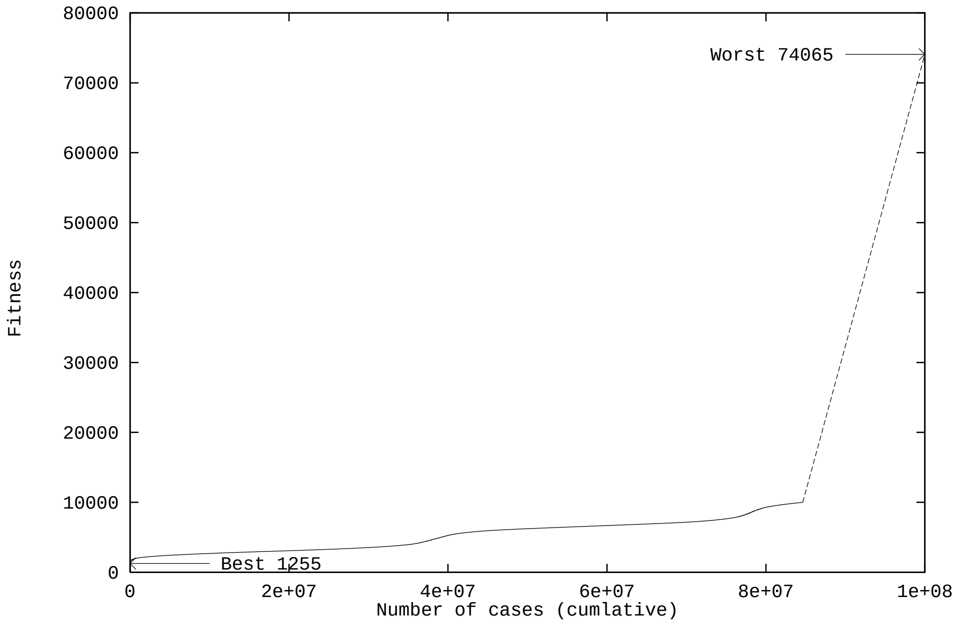

Figure 7: Cumulative frequency v. Fitness (No line maintained more than once)

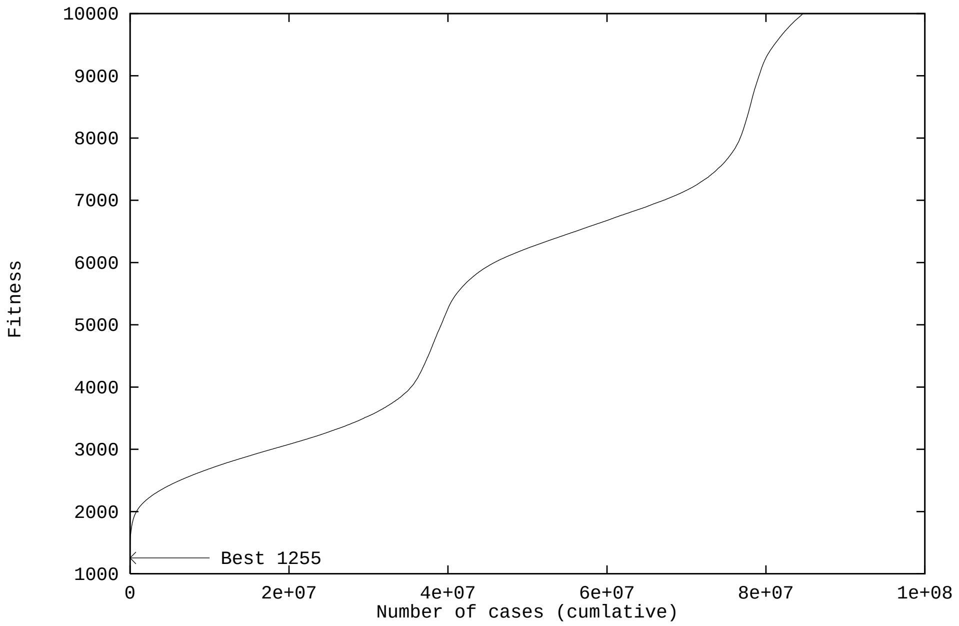

Figure 8: Cumulative frequency v. Fitness (No line maintained more than once)

9 Conclusions

This paper has presented a series of experiments which show genetic algorithms in conjunction with a "greedy optimizer" can readily solve the four node problem and, provided the growth in fitness evaluation time can be contained, gives reason to belief that, if carefully applied, this approach will be suitable to scheduling the maintenance of the National Grid.

10 Further Work

We plan to consider alternative chromosomes, such as expansive coding and genetic programming before considering scheduling the maintenance of a significant part of the National Grid (possibly South Wales; a network of about 30 nodes).

So far we have not considered the problem of ensuring maintenance does not unduly effect the resilience of the network. This will be left to a later stage.

The computation cost of the fitness evaluation and greedy optimizer can be reduced by using "sparse matrix" techniques. Further speed ups may be achieved by using approximate rather than exact calculations.

11 Acknowledgments

W. B. Langdon is funded by the EPSRC and the National Grid Company Plc.

I would like to thank Mauro Manela and my supervisors (M. Levene and P. C. Treleaven) for their critisims and ideas, Mike Calviou and Maurice Dunnett for the fitness function and the four node problem, and Laura Dekker for assistance with setting up QGAME.