Scheduling Generator Outages to Minimise Financial Impact

Phil Sheppard
ECNZ, Northern Generation
Hamilton, New Zealand
phil.sheppard@ecnz.co.nz

Abstract

Many papers have been written on short and medium term scheduling of hydro-electric generation on a river system so as to maximise operating profit. A related area of growing importance is the scheduling of generator maintenance outages to minimise their impact on the profit objective.

Historically, the co-ordination of generator outages on a river system has been performed on a "manual" basis using "rules of thumb" without knowledge of the financial impact. Determining the exact financial impact of a historical outage is a complex task, the market equilibrium may be very different with and without the outage. For example, the market clearing price may change significantly and the resulting distribution of generation among market participants may be quite different. However, the financial impact of a set of outages can be approximated with sufficient accuracy to provide sensible data for decision making.

A methodology to (i) cost a set of outages, and (ii) minimise their total cost is proposed. The total cost of a set of outages is adjusted by the timing of the outages, and a genetic algorithm is used to identify an improved solution. It is shown that the financial impact of a set of outages can be reduced given some freedom in outage timing.

1. Introduction

Maintenance of generating plant generally requires plant to be shutdown and released from service. There are normally several hundred planned outages over the 39 generating units on the Waikato River each year. Each outage may impact financially on ECNZ and a well co-ordinated outage programme can make a significant difference to the total cost.

The introduction of the Electricity Market has fundamentally changed the operation of ECNZ from one where revenue was fixed and total costs of generation were to be minimised to one where revenue is uncertain, making the profit objective more complex. Prior to the first generation split-up of ECNZ (into ECNZ and Contact Energy) any lost generation resulting from a generator outage would be “picked up” at another ECNZ station, so the revenue would not change but the cost of generation may increase.

Energy “lost” due to a generator outage may be replaced by another Waikato generator (with a less profitable operating schedule), by another generator within ECNZ, by a competitors generator, or by a combination of the above. In the first two cases the energy lost incurs increased costs while in the third case the energy lost incurs loss of market share and revenue. Loss of revenue can be represented as a cost and throughout this paper the financial impact of an outage is referred to as a cost whether it arose from increased costs or loss of revenue.

In order to extract financial benefits from effective outage co-ordination it is necessary to estimate the cost of future outages and then schedule each outage so as to minimise the total cost of all outages. Historically, the co-ordination of outages has been performed on a manual basis using rules of thumb without knowledge of the financial impact. The objective of this investigation then, is to develop methods for:

1. Costing a set of future planned outages
2. Minimising the total cost of a set of future planned outages by moving outage start dates within early and late start bounds.

Consistent with many decision support tools in the electricity industry, the method presented in this paper uses discrete half hour time steps and this is the definition of a period. Although outages can start and end at any time, outage start and end times are rounded to the nearest half hour.

2. Costing a Set of Outages

The catch with outage costing is to estimate what the station would have generated had the unit been available. In some cases the generation at the station would have been no different, in others the unavailability of a generating unit may change the optimal basis significantly (in general hydro-electric schemes are scheduled by a linear program). Determining exactly what the station would have generated had the unit been available is a complex task in a market environment. There are many variables including the possible loss of market share to a competitor and the unpredictability of competitor behaviour.

Just prior to the commencement of the New Zealand Electricity Market, ECNZ developed an application called PCSchn which was used to schedule generation across all its stations. This application was based on a linear program using a Cplex solving engine. The objective function for this LP was to minimise costs (assuming a fixed demand and revenue stream) and an outage could be costed by solving the LP with and without the outage [1]. The linear program represented each hydro station with a piecewise linear power/flow curve, and a schedule would solve in approximately 5 to 10 minutes. A single outage could therefore be costed in approximately 15 minutes. A batch process was used at Northern Generation to calculate the previous days outage costs taking approximately 3 hours per day.

This method described the picture reasonably accurately in an environment where ECNZ represented a large majority of the NZ generation picture. However, once NZEM commenced, competitor data was not available and the objective function was no longer relevant in an environment where generation and revenue is not fixed.

At Northern Generation we decided to investigate the feasibility of developing a more approximate but faster method for costing an outage. The rational for this decision was:

1. The uncertainties surrounding the inputs are such that there is little benefit to be gained in employing a highly accurate method. In addition, the behaviour of other market participants cannot not be predicted with any degree of certainty.
2. To provide station and project management staff the ability to cost outages without having to be trained in a complex application. A key objective at Northern Generation is to promote commercial awareness throughout the Group. For a tool to be used consistently by a range of staff it needs to be easy to use and provide quick results.
3. Ultimately a fast outage costing algorithm is necessary to support a multiple outage co-ordination process.

Scheduling the Waikato scheme is essentially a daily process as there is no significant day to day storage. The flow transit time from Lake Taupo to Karapiro Power Station is approximately 17 hours with a 13 hour transit time to the second lake, Ohakuri. Lake Ohakuri is the prime storage lake on the river and water is released during the night from Lake Taupo through Aratiatia Power Station to prepare Lake Ohakuri for the following day.

A detailed investigation into generation patterns on the Waikato River revealed strong correlation between the daily total generation across the river and the average power at each station during each half hour (see chart 1 below). In view of the small daily storage this result is not surprising and essentially it says that the today's generation distribution across the Waikato stations is determined largely by the days' expected generation. Therefore we can use daily Waikato generation values and the linear relationships between daily generation and average station power to estimate the expected generation at each station during each half hour.

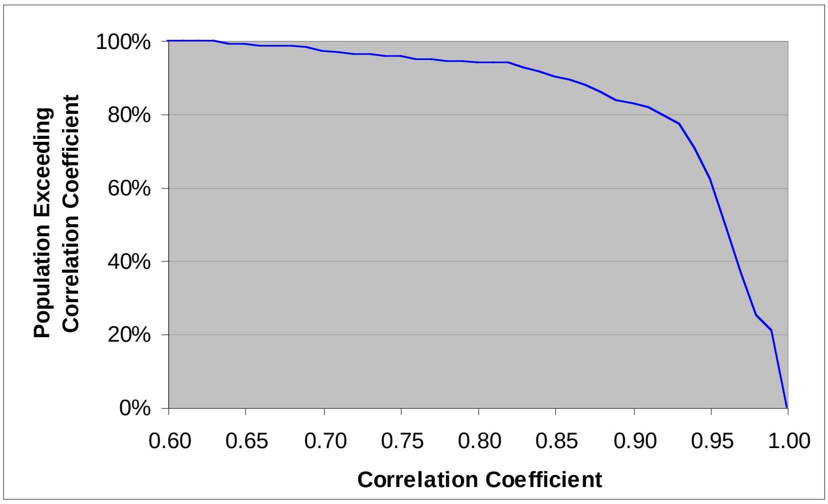

Chart 1: Exceedence Curve for Station and Total Generation Correlation Coefficients

A key assumption of the above method is that the lost energy from a days' outages at single station is small in comparison to the daily total generation on the river.

With a method to estimate what a station would have generated had the unit been available, what remains is to determine the shortfall (if any) and where the energy would be replaced. The shortfall of the available generation from the expected generation (if positive) provides the lost generation for the period. The lost generation can be replaced in one or more of the following three ways: by another station on the Waikato River (with a less efficient operating schedule), by another station within ECNZ, or by a competitor in the spot market. At this point the method gets more imprecise and three tranches are defined to approximate the replacement of energy arising from an outage. The first tranche estimates the daily quantity of energy that could reasonably by replaced within the Waikato chain and the cost associated with that. The second tranche estimates the daily quantity of energy that could reasonably be replaced within ECNZ and the cost associated with that. The third tranche assumes that the energy is replaced by a competitor at the spot price. The daily sum of the lost generation across the river is applied to the replacement energy function to determine the total cost of replacing the lost energy.

The first step in costing a set of outages involves determining the number of units unavailable at each station for each period of the year, this sets the availability. This process takes care of overlapping outages at each station. Later in the process it is necessary to check for overlapping outages across to river to correctly cost the replacement energy according to the price/quantity tranches.

Mathematical Representation of the Outage Costing Algorithm

$$Availability_{s,t} = LookupTable\,(NumUnitsOut_{s,t},\, s)$$

$$Lenergy_d = \sum_{\text{all } p,\, s} Max(DailyGen_d.Slope_{s,p} + Intercept_{s,p} - Availability_{s,t}\ , 0)$$

$$Cost = \sum_{\text{all } d} \begin{cases} Lenergy_d.P^1_d & Lenergy_d \le Q^1_d \\ (Lenergy_d - Q^1_d).P^2_d + P^1_d.Q^1_d & Q^1_d < Lenergy_d \le Q^2_d \\ (Lenergy_d - Q^2_d).P^3_d + P^2_d.Q^2_d + P^1_d.Q^1_d & Lenergy_d > Q^2_d \end{cases}$$

$NumUnitsOut_{s,t}$ is the number of units unavailable for generation at time t and station s.
$Lenergy_d$ is the daily total energy lost due to all generator outages within that day.
$DailyGen_d$ is the expected daily generation profile for all the Waikato stations.
Q^n and P^n represent the expected three tranche daily replacement energy function.

Subscripts:
s station (Aratiatia to Karapiro, 1 to 8)
t time (in half hour increments over time horizon)
p period (particular half hour of the day)
d day

3. Minimising the Cost of a Set of Outages

The second objective of the investigation is to minimise the total cost of a schedule of outages. This section describes the objective function, decision variables and constraints, and following this alternative search methods are considered.

The objective function is to minimise the total cost across the time horizon (one year) of all the outages. Alternative solutions are gained by moving the *start date* of one or more outages subject to the following constraints:

1. The duration of an outage remains fixed. Optimising the trade-off between reducing outage duration and increasing resources is outside the scope of this investigation.
2. The start time of each outage is bounded with *early* and *late start times*. Early and late start times must be set such that the entire outage is within the time horizon with any allowed start time.
3. In addition to early and late start, the *lead time* is included which can be described as the time required to prepare for the outage. Once the time between the start date and now is less than the specified lead time the outage becomes fixed and can no longer be moved by the optimisation. An outage cannot be moved such that the time between the start time and now is less than the specified lead time.
4. Other constraints to provide linkage from one outage to another can be introduced. For example if four generators at one station are to be maintained in sequence then the start time of subsequent outages can be linked by a fixed delay to the start time of a prior outage, leaving only the start time of the parent outage to be moved by the optimisation. This type of constraint is discouraged as it defeats the purpose of the optimisation. However, the schedule cost with and without the linkage constraint can be found to determine the cost of including the constraint.

The time of day that the outage commences is not changed by the optimisation. If the outage duration is more than a couple of days then the impact of varying the starting time in the day is small. For short outages (less than a couple of days in duration) the starting time of day is assumed to have been optimised prior to this process (tools are available for this).

Unfortunately the conversion process of outage start and end times to a half hourly availability profile and the step-wise replacement energy function mean that approximation for a linear program would be very complex. The solution space for a combinatorial approach would be very large requiring approximately D^n solutions to be investigated where D is the average number of days available within which an outage can be moved (average late start - early start) and n is the number of outages. Typically D is 50 and n is at least 200 giving at least 10^{339} solutions!

The problem appears ideally suited to a heuristic such as a genetic algorithm using an array of outage start dates as the decision variables. Tabu search appears to be another suitable candidate and an investigation to compare this against the genetic algorithm would be useful to build confidence in the optimisation. Reeves and Beasley [2] put it well, *"should we prefer an exact solution of an approximate model, or an approximate solution of an exact model"*.

The next section discusses implementation of the genetic algorithm framework, including creation of the initial population and the selection, crossover and mutation operators.

4. Implementation of Genetic Algorithm

The key idea behind the basic genetic algorithm is that the next generation (population) of solutions is “fitter” than the previous generation through the process of selection, crossover and mutation. The functions underlying a genetic algorithm are modelled on biological genetic theory. The simplest single-stranded chromosome organism (haploid) with genes representing the start day of each outage is used in this investigation.

Noting Goldbergs principle of minimal alphabets [3]: “*The user should select the smallest alphabet that permits a natural expression of the problem*” the starting date of each outage was initially represented as a integer between 1 and 365. A binary alphabet was also investigated as a comparison and it was found that performance of the genetic algorithm improved significantly. It is interesting to note that the binary alphabet most closely models genetic codes with genes taking one of two states, dominant or recessive.

Thus, starting dates for all the outages are converted into binary and concatenated into a single string or chromosome. A population of members is created and then the genetic operators are applied to create the next generation. Each member represents a schedule of outages with a starting date chromosome and an associated fitness (or cost). To build the initial population, members are derived from the initial schedule by randomly moving the start time of each outage within the constraints.

Once the population is created each member is evaluated and then the selection function chooses members for the next generation by random selection weighted by the fitness of each member. In this case the lower the cost of a member, the greater its probability of being selected.

The crossover and mutation functions are not invoked in the creation of every new member in the next population, rather they are invoked according to a defined probability. Typically crossover is invoked with a 50 % probability and mutation is invoked with a 5 % probability. The effect of varying the probability of mutation is shown in the next section.

The crossover operator involves taking two selected members (A and B say) and choosing a random crossover location in the concatenation of start times. The first member takes the outage start times of member A up to the crossover location and then takes the remaining outage start times of member B to compete the schedule. The second member does the same with member B first then member A.

The mutation operator involves moving through each bit in the binary string and with a 50 % probability changing the state of the bit. The mutation operator is most affected by the change from the 1 to 365 integer alphabet to the binary alphabet.

The outage cost functions and genetic algorithm were implemented in C primarily for speed. Recently a user interface has been implemented in Delphi which calls the C functions as a dynamic link library.

5. Results

The outage plan for the 98/99 financial year has been used as a test input to the optimisation process. This plan contains 180 outages over 39 generating units on the Waikato hydro scheme ranging in duration from 2 hours to 17 weeks. Approximately half of the outages were constrained by the “linkages” to other outages.

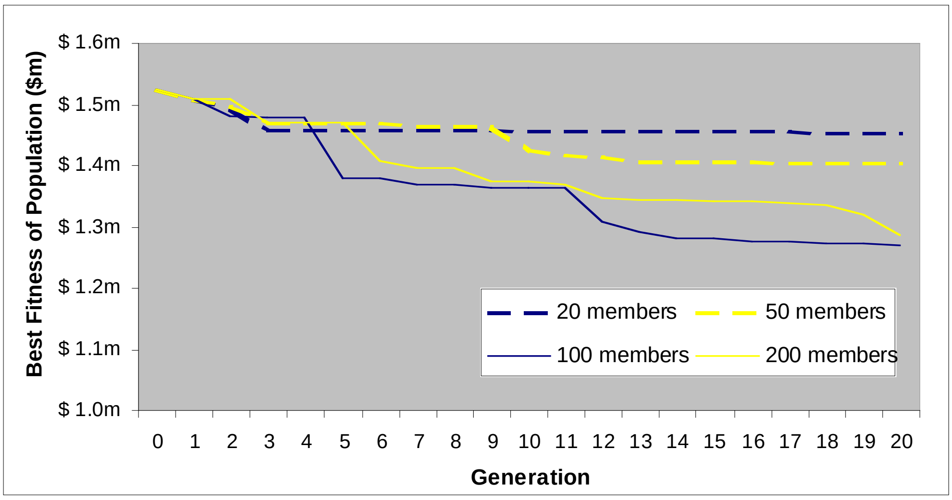

Chart 2: Best fitness found with varying population size

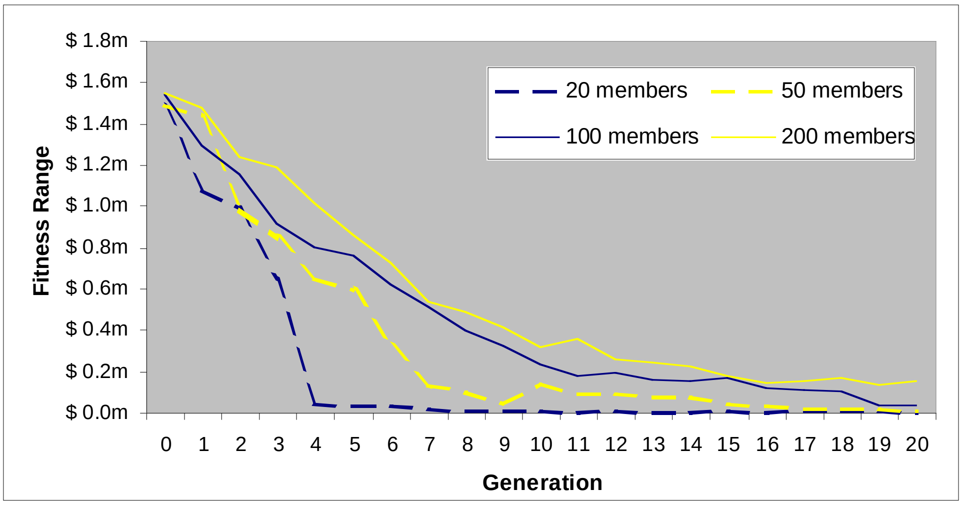

Chart 3: Fitness range found with varying population size

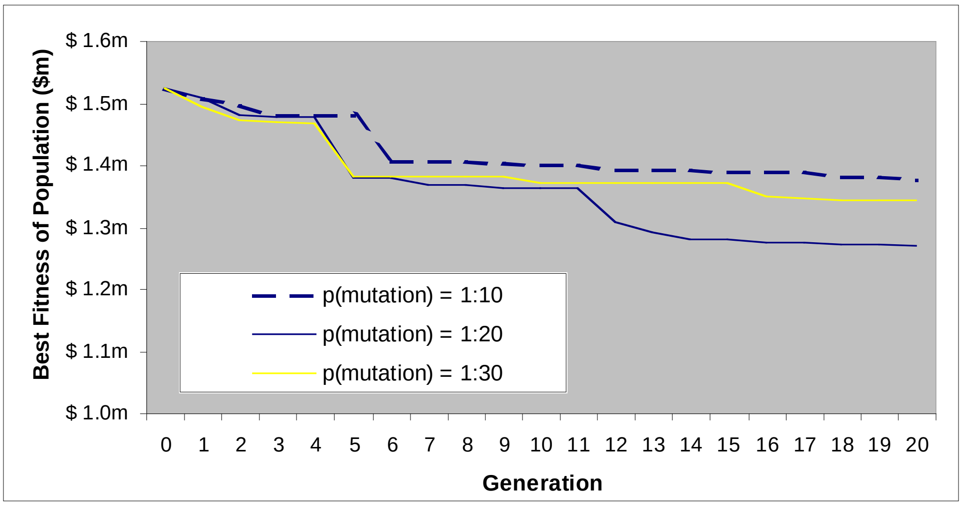

Chart 4: Best fitness with varying probability of mutation

The above charts show the effect of population size and probability of mutation on the optimisation process. The larger populations provide more potential for finding a high performing individual so they consequently take longer to converge to a small range. This can be observed in chart 3 showing the range between the most fit and the least fit member in the population. As a trend, a better "fitness" is obtained with a greater population size (chart 2) however the fitness obtained for a 200 member population is not better than that obtained for a 100 member population *after 20 generations*. If additional generations were created, the 200 member population would likely perform better than the 100 member population since "potential" still remains in the population (the range has not converged after the 20 generations).

The initial cost of the schedule as set by the station planners carried a cost of $ 2.1m. Optimisation reduced this cost to $ 1.3m using a population size of 100 members and a probability of mutation of 1:20.

6. Conclusion

The objective of establishing a method to cost and minimise the cost of a schedule of outages has been achieved. It has also been demonstrated that a 40 % saving can be made in the 98/99 year planned outage programme by employing such a optimiser. However, this seemingly remarkable result should be tempered by the following points:

1. Outage costs are partly seen as "funny money" as they are not directly observed in the operating revenue or costs of the company.
2. There is always some degree of uncertainty surrounding the cost of an outage as it is not possible to determine with certainty:
   - what a generator would have generated had it been available and
   - where the lost generation is replaced

7. Acknowledgements

Anping Wang for introducing the concept of a genetic algorithm to me.

8. References

[1] Dispatch Development Project Team, *Dispatch Development Closeout Report*, ECNZ Internal Report (Volume 2).

[2] C.R. Reeves (ed), *Modern Heuristic Techniques for Combinatorial Problems*, Halsted Press (1993).

[3] D.E. Goldberg, *Genetic Algorithms in Search Optimisation and Machine Learning*, Addison Wesley (1989).