Optimisation of a Dairy Farm Model - Comparison of Simulated Annealing, Simulated Quenching and Genetic Algorithms

D.G. Mayer, J.A. Belward*, K. Burrage* and M.A. Stuart

Queensland Department of Primary Industries, Animal Research Institute, LMB 4, Moorooka QLD 4105;

*Centre for Industrial and Applied Mathematics and Parallel Computing, University of Queensland, QLD 4072.

Abstract Available optimisation techniques vary widely in terms of derivation, application, and efficiency. A complex dairy farm model was used to benchmark those which have been used previously in the model optimisation field. The more traditional methods, including random search, hill-climbing and direct search, were notably inferior in identifying the economic optimum of this agricultural system. Genetic algorithms proved quite efficient, but overall results were marginally down on those from the simulated annealing methods. Initially, these proved to be quite slow, but a retuned simulated annealing algorithm was found to be more efficient, thorough and safe. It's extension to simulated quenching proved best for this problem, safely identifying the optimum at a good rate of convergence. As this program is freely available and relatively easy to use, it is strongly recommended. Also, initial investigations with the tabu search strategy are reported, which show it to have potential.

1. INTRODUCTION

The management of a dairy farm to achieve the economic optimum is a complex operation. Sufficient feed must be provided for the herd to maintain a high level of milk production, and this must be balanced against the relative costs of each. The overall management strategy covers critical decisions such as stocking rate, calving pattern, supplements (type, level and timing), fertiliser and irrigation inputs, planting and grazing strategies, and the balancing of pasture and forage areas and species. The interacting effects of these options can only realistically be investigated by a simulation model of the system. An available dairy farm model was used as a test case for the optimisation methods, with the definable management options forming a 16-dimensional optimisation problem. Even at a course level of resolution, there are of the order of 10^{13} possible combinations of management options, which is an imposing number to explore and evaluate. Previous studies [Mayer et. al. 1989] showed this problem to be of a difficult form for optimisation algorithms, due to non-smooth surfaces and multiple optima.

The purely random search method has been included in a number of optimisation studies [Corona et. al. 1987, Bramlette & Bouchard 1991, Davidor 1991, Syswerda 1991], more as a poorly-performing benchmark against which to compare the more targeted methods. It cannot either theoretically or practically be justified as a serious optimisation technique. The more traditional optimisation algorithms include hill-climbing (conjugate-gradient or quasi-Newton techniques) and direct search methods (the Nelder-Mead simplex, and complex algorithms). In some cases [Karr 1991, Mayer et. al. 1991] these have been shown to perform well, but in general they are inferior to the more recently-developed optimisation techniques. The first of these, genetic algorithms, has been demonstrated as superior to hill-climbing methods, in terms of both speed of convergence and values of reported optima [Bramlette & Bouchard 1991, Davidor 1991].

The second more recent technique covers the field of simulated annealing and it's extension of simulated quenching. On test functions, Corona et. al. [1987] showed simulated annealing to be more efficient than the Simplex algorithm. Bramlette & Bouchard [1991] demonstrated it as clearly superiour to hill-climbing, and marginally better than a genetic algorithm, in optimising a model of combat aircraft design. Goffe et. al. [1994] reported that simulated annealing found far better solutions to econometric parameterisation problems than did either hill-climbing or the Simplex algorithm, but that this improvement was at the expense of much greater computation time. A notable exception to this trend of simulated annealing being the best is reported in Styblinski & Tang [1990], where hill-climbing proved faster and more accurate on a set of smooth, unimodal test functions. This is not surprising, as these are the types of problems which conform exactly to the underlying assumptions of the hill-climbing methods. With problems which are more 'real-world' in nature, hill-climbing is generally found to be inferior to the more robust methods.

In a comprehensive test between the more modern techniques, Ingber & Rosen [1992] compared genetic algorithms against simulated annealing over a range of multi-dimensional functions. Overall, simulated annealing proved to be superiour, especially in terms of the values of the identified optima. There was some variation in these results, however, with the genetic algorithm notably performing better than simulated annealing on the higher-dimensional problem.

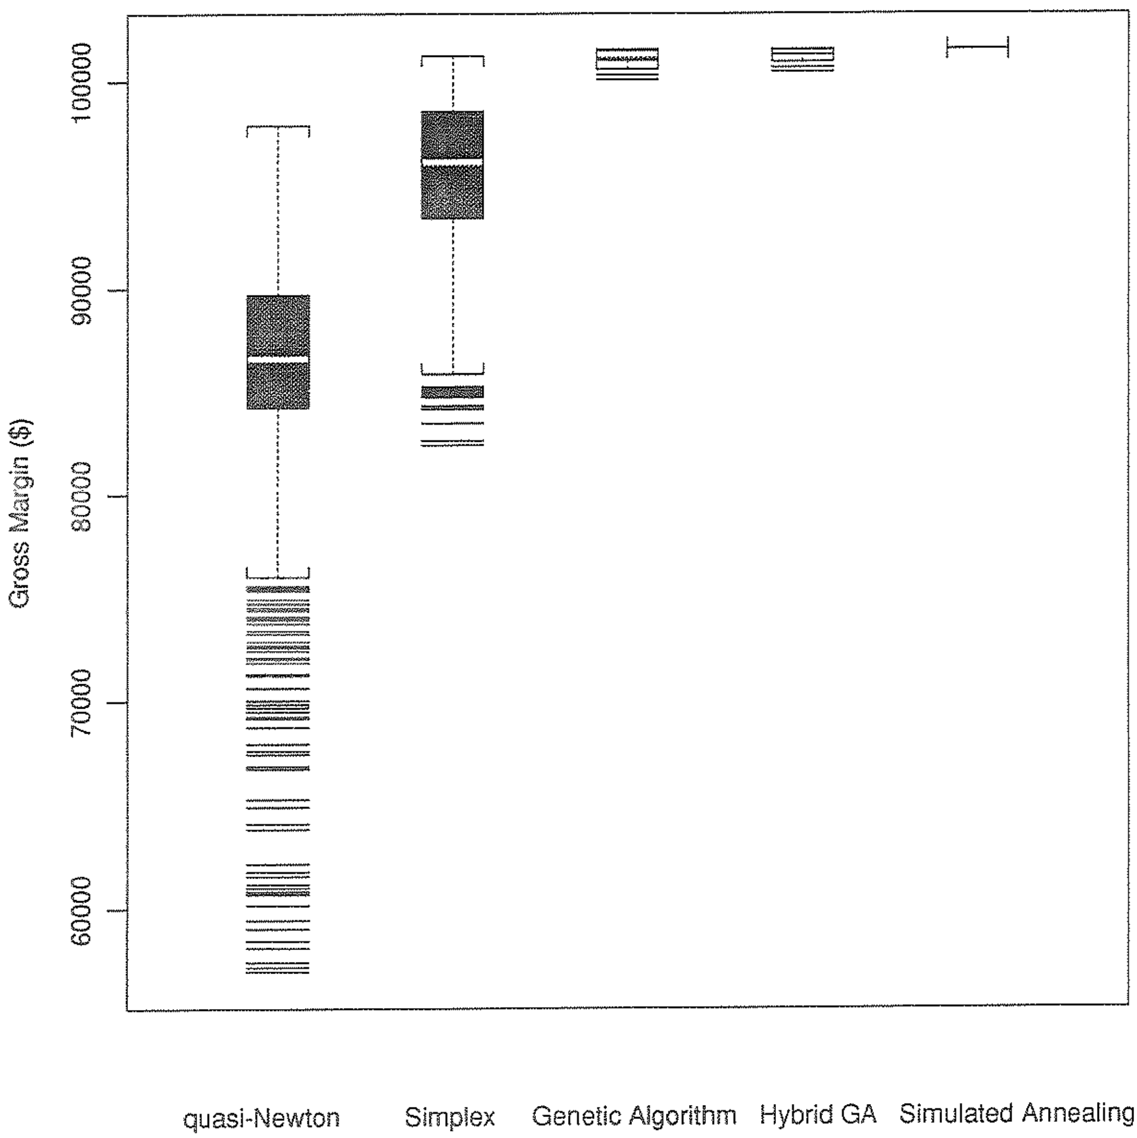

Figure 1. Boxplot distributions of identified optima of the dairy farm model, by optimisation method.

These newer optimisation methods have been previously tested against the more traditional algorithms, using the dairy model as described [Mayer et. al. 1995]. Each method was replicated so that the total number of model runs was in the order of half a million, and the resultant distributions of the reported optima (non-discounted farm gross margin, in dollars per annum) are shown in Figure 1.

The 'hybrid GA' is a hybrid of a genetic algorith which has been fine-tuned by incorporating a hill-climbing routine as it's final process, to allow this discrete method to converge to the actual optima in the continuous hyperspace. This extension showed little improvement over the basic genetic algorithm, so was not warranted. From this figure, it is clear that the hill-climbing (quasi-Newton) and direct search (Simplex) methods are inferior for this particular problem. As these results are consistent with others found in the literature, these methods were discarded from further consideration. This paper covers a more thorough investigation into the properties and performance of genetic algorithms and simulated annealing, including an evaluation of the latter's extension of simulated quenching.

2. GENETIC ALGORITHM (GA)

Based on the natural selection theories of John Holland, GAs mimic the operation of sexual reproduction between individuals [Davis 1991]. The operational parameters are coded onto 'genes', and over modelled generations the 'survival of the fittest' rules tend to produce successful combinations of options, to eventually

Source	Population Size	Crossover Rate	Mutation Probabiity	Generation Gap	Elitism	Rank Based
Davis [1991]	50	0.65	0.008			Yes
Bramlette and Bouchard [1991]	4-100	0.01-1	0.01-1		No	
South et. al. [1993]						
– standard	50	0.60	0.001	1	Yes	
– quick	30	0.95	0.01	1	Yes	
– comprehensive	80	0.45	0.01	1	Yes	
GENESIS defaults	50	0.60	0.001	1	No	No
Dairy model						
– ranges	40-80	0.45-0.95	0.001-0.01	0.9-1	Yes	No, Yes
– optimal	40	0.60	0.001	1	Yes	No

Table 1. Operational parameters of the genetic algorithm (blank entries indicate not specified).

arrive at the optimum. This cross-mixing and searching strategy is very targeted and efficient, and has allowed the solution of more higher-dimensional problems which could not previously be attempted [Radcliffe & Wilson 1990]. The use of GAs in real-world applications is becoming more widespread [South et. al. 1993]. Annevelink [1992] used a GA to solve a horticultural design problem which had proved intractable with a mathematical programming approach.

The performance of GAs is largely controlled by a number of operational parameters. Table 1 lists the more important of these, along with ranges suggested in the literature as worthy of consideration. Also given in Table 1 are the values trialed with the dairy model, using the Genetic Search Implementation System (GENESIS, Version 5.0). For this investigation, elitism was used throughout, as this guards against losing an identified optima, and grey coding was used for the integer options to protect against the effect of Hamming cliffs. Factorial analyses of variance of these results revealed the optimal combinations of the control parameters (as listed in Table 1), with population size, mutation rate and the score-based selection method proving to be the most important.

For the optimal values of these GA parameters, all 8 replicates converged (from a practical viewpoint) to the global optimum of the dairy model. The rate of convergence for the best and worst of these replicates is shown in Figure 2.

3. SIMULATED ANNEALING (SA)

SA covers an evolving family of methods which mimic the cooling (annealing) process in metallurgy [Kirkpatrick et. al. 1983]. Originally based on the Boltzmann distribution [Metropolis et. al. 1983], SA probabilistically accepts less successful (or backwards) steps, allowing this method to escape areas of local optima. This probability is controlled by the temperature schedule, and is quite high in the initial 'searching' stages, eventually becoming negligible in the final stages of convergence. Whilst there is no guarantee of finding the global optimum, SA's slow, thorough nature usually ensures this [Ingber 1993]. Practical applications show it to be far superior to hill-climbing methods, in studies on the harvest scheduling of forestry stands [Lockwood & Moore 1993], and in optimising groundwater remediation strategies [Kuo et. al. 1993].

Previous applications of SA to this dairying model using 'very fast simulated reannealing' showed it to be successful in always finding the global optimum, but taking on average 10^5 runs to achieve this [Mayer et. al. 1995]. An improved SA algorithm, ASA (Adaptive Simulated Annealing Ver. 5.10 ; available through e-mail from asa-request@alumni.caltech.edu) was obtained and implemented. This uses an exponential temperature schedule, making it doubly-exponentially faster than the original (Boltzmann) SA. Other improvements in ASA include the ability to define true integers, having control over their annealing process, and easier control of operational parameters through an options file rather than recompilation of the C-code each time.

In ASA, many of the operational parameters are claimed to be critical. Initial investigations of these on a one-at-a-time basis were confounded by the random nature of the search path, with some apparent successes merely due to chance. A more rigorous method is to replicate comparisons, and this is most efficiently achieved with factorial designs. A number of these factorials were used to investigate ranges of the more important parameters. Somewhat surprisingly, shifts of up to a factor of 10 in either direction from the default values had little or no effect for some, and a detrimental effect for others, so for these the default values were adopted. These variables (using parameter names as defined in the on-line ASA documentation) included the initial_parameter_temperature, controlling the starting point temperature for annealing; the temperature_annealing_scale, which assists in the annealing schedule; the cost_parameter_ratio_scale, which relates the cost annealing schedule to that of the parameters; and delta_x, controlling the calculation of tangents for reannealing. Also, investigations forcing proportionately more acceptances, tending towards the 'rejectionless annealing' approach of Ingber [1993], performed poorly, so this strategy was abandoned.

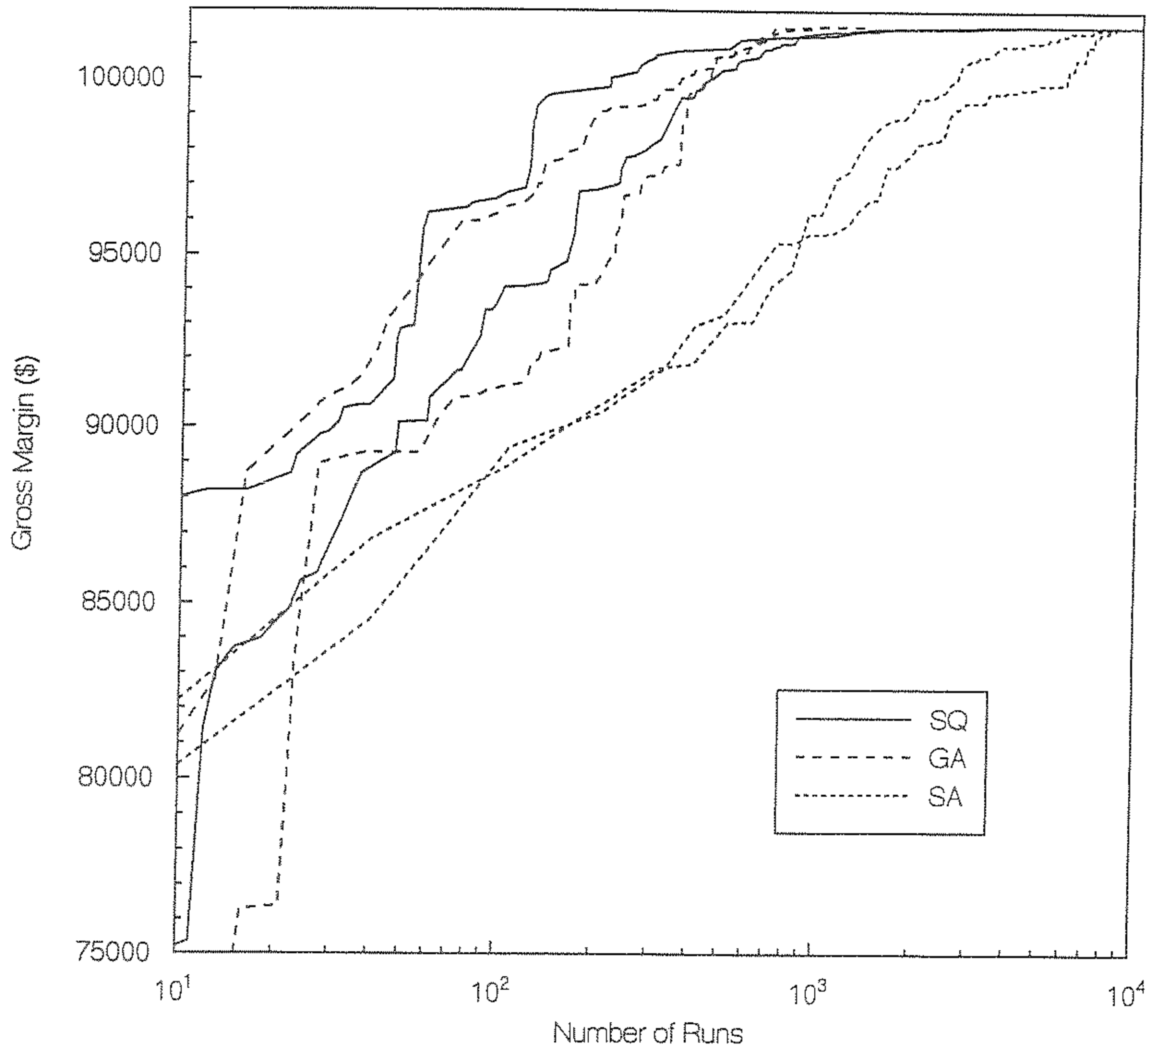

Figure 2. Rate of convergence to optimum for two replicates of the simulated quenching (SQ), genetic algorithm (GA) and simulated annealing (SA) methods.

The single parameter which consistently and significantly (P<0.01) affected performance was the temperature_ratio_scale, which controls the rate of temperature decline. With a default of 1E-5, higher values (1E-4) slowed the search too much, resulting in excessive computation required to find the optimum. Smaller values (down to around 1E-10) improved the efficiency, but values towards the bottom of this range or lower proved inconsistent, occasionally failing. The 'safe' default value was thus adopted for the SA test runs, with the rate of convergence for two random replicates plotted in Figure 2. It is notable that the combined effect of these improvements resulted in convergence ccuring in under 10^4 runs, which is a factor of 10 better than earlier SA optimisations of this problem [Mayer et. al. 1995].

4. SIMULATED QUENCHING (SQ)

SQ, also termed simulated tempering, is an extension of SA. Quenching implies a more rapid cooling schedule, which allows it to be a faster or 'more greedy' process. However, this also introduces dangers - if too fast, the probability of arriving at a suboptimal solution becomes more than negligible [Ingber 1993]. This speeding up is determined by the quenching factor (Q), which by recommendation is approximately the number of dimensions in the problem [Ingber 1993]. Fractional values can also be used to slow down the cooling process, making the basic SA process even more thorough.

A geometric progression of Q was investigated, namely values of 0.5, 1, 2, 4, 8, 16, 32 and 64. These were crossed with four values of the temperature_ratio_scale (1E-5, 1E-6, 1E-8 and 1E-10) and two random starts (replicates) in a complete factorial design, and results subjected to analysis of variance. This showed a significant (P<0.01) interaction between the design variables, as graphed in Figure 3. Here, the best Q across the range of temperature scales is 8, although results for 16 (the number of dimensions, and expected to be optimal) were not significantly higher. Whilst the lowest temperature scale (1E-10) appeared the best, its results tended to be erratic and more variable, and values in this range proved unreliable when used for SA. Hence, the

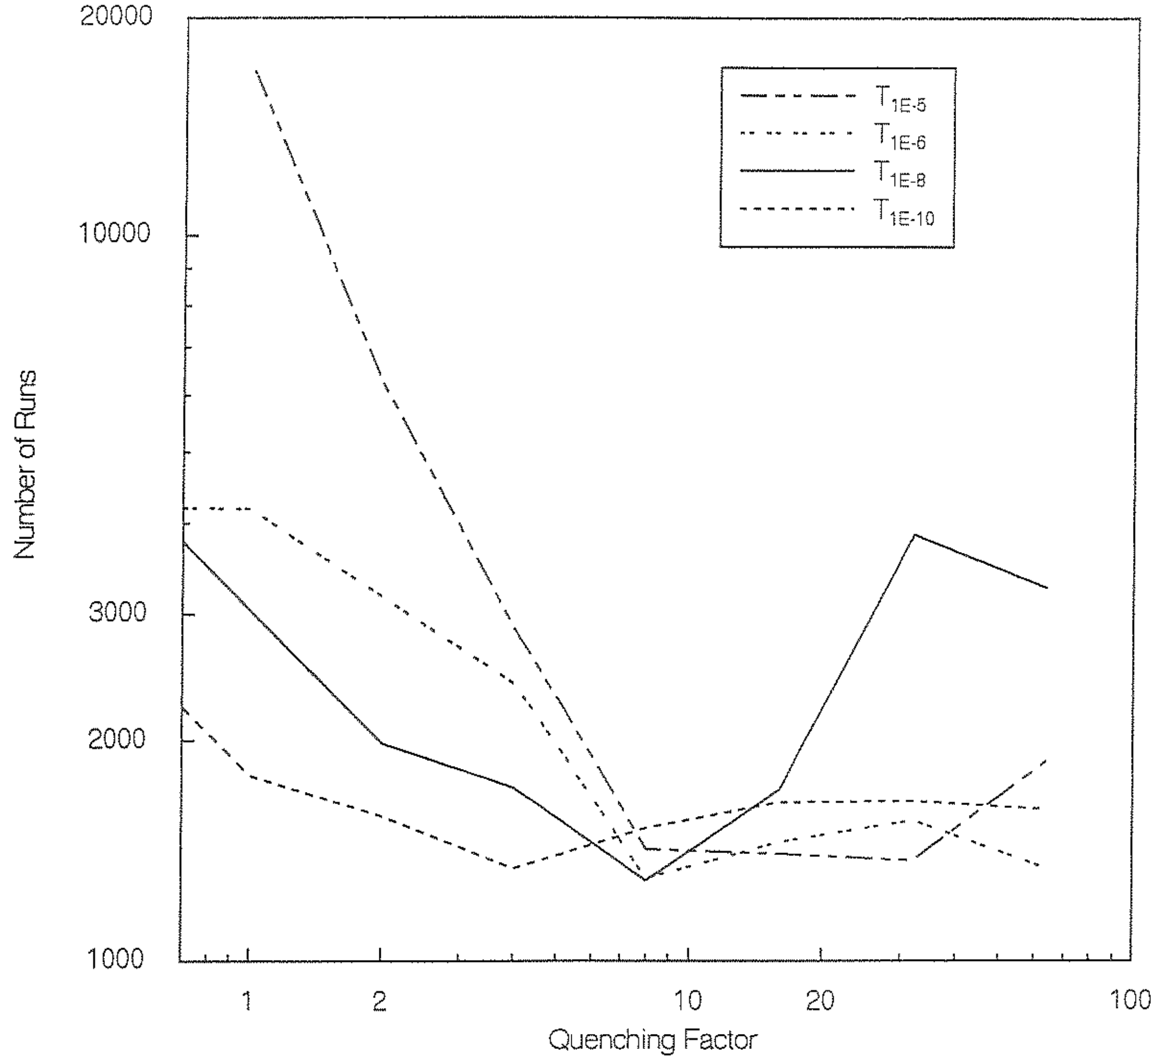

Figure 3. SQ runs to achieve optimal value, for a range of quenching factors and temperature ratio scales (T).

actual best combination (Q of 8 and temperature scale of 1E-8) was taken as approximately optimal for this problem, and the rate of convergence for the two replicates is graphed in Figure 2. Here, SQ's rate of convergence is about the same as that of the genetic algorithm. Unlike the GA, however, SQ always converged to the global optimum of the system, and is thus more reliable.

These data, along with further trials, also revealed a few failures with 'super-quenching' - of a total of 22 trial runs with Q in the range 32 to 100, 3 failed to identify the global optimum. As the rates of convergence of these 'super-quenching' runs were no better than those of the more moderate Q values, the latter appear appropriate.

5. TABU (OR TABOO) SEARCH STRATEGY

Tabu search is a metastrategy which can be applied to other optimisation methods [Glover 1990]. It primarily involves maintaining a list of recently-visited options, which are regarded as 'tabu'. This ban both prevents redundant re-evaluations, and also forces the search away from recently-investigated areas of the hyperspace, allowing it to escape from local minima. It has most commonly been used in conjunction with hill-climbing algorithms, and one such version was independently developed as the 'steepest ascent mildest descent method' [Hansen & Jaumard 1990]. A tuned tabu/hill-climbing method was shown to be marginally superior to SA on low-dimensional multiple-optima test functions [Cvijovic & Klinowski 1995].

Our investigations of this method on the dairy model are ongoing, and include use of the Reactive Tabu Search Code from the University of Trento, Italy (http://rtm.science.unitn.it). Initially, we attempted adding a tabu list into the SQ algorithm. It offered only marginal improvement (eg, on average 10 % fewer runs to get to within 99.9 % of the optimum), and encountered slight problems regarding convergence.

It appeared that the discrete nature of obtaining a 'tabu match' interfered with the final fine-tuning of SQ, and these runs terminated on average 0.01 % short of the optimum. Given also the practical requirements of extra code and program interfaces, and the necessary decisions over the required level of discretisation and tabu list length and period, it appears that its disadvantages outweigh the marginal improvement offered, and its addition to SQ is not merited. This could be because SQ works so well on this particular model, and results may well vary with different problems or indeed optimisation methods. In line with literature studies, it is expected that a tabu/hill climbing algorithm will perform well, and this method is currently under investigation.

6. DISCUSSION AND CONCLUSIONS

Simulated annealing appears a thorough and reliable method for optimising simulation models. It is useful both in its own right, and via its more efficient adaptation of simulated quenching, provided the level of quenching is kept within reasonable levels. Apart from this obvious factor, the rate of temperature decline is the most critical control parameter. When this method is used in a practical sense, a range of values should be tried for both these parameters.

The genetic algorithm offers a good rate of convergence, but is not quite as reliable in terms of the final values found, although certain parameter settings inspire more confidence. With higher-dimensional problems it may well turn out to be the best optimisation method, as the noted inefficiencies of the simulated annealing approaches may become overwhelming.

On real-world problems, with non-smooth surfaces and multiple optima, either genetic algorithms or simulated annealing (including simulated quenching) should outperform the more traditional optimisation methods such as hill-climbing or direct search algorithms. Hence, either of the former methods, with perhaps the addition of the Tabu search strategy, are worthy of consideration for use in practical applications.

7. ACKNOWLEDGEMENTS

We are grateful to Lester Ingber and John Grefenstette for making their optimisation algorithms publically and freely available, and wish also to thank Jan Neale for assistance with the figures, and Amanda Lloyd for word processing.

8. REFERENCES

Annevelink, E., Operational planning in horticulture : Optimal space allocation in pot-plant nurseries using heuristic techniques, *J. Agric. Eng. Res.*, *51*, 167-177, 1992.

Bramlette, M.F. and Bouchard, E.E., Genetic algorithms in parametric design of aircraft, in *Handbook of Genetic Algorithms*, edited by L. Davis, pp. 109-123, Reinhold, New York, 1991.

Corona, A., Marchesi, M., Martini, C. and Ridella, S., Minimizing multimodal functions of continuous variables with the "simulated annealing" algorithm, *ACM Trans. Math. Soft.*, *13*, 262-280, 1987.

Cvijovic, D. and Klinowski, J., Taboo search : An approach to the multiple minima problem, *Science*, *267*, 664-666, 1995.

Davidor, Y., A genetic algorithm applied to robot trajectory generation, in *Handbook of Genetic Algorithms*, edited by L. Davis, Reinhold, New York, 1991.

Davis, L., *Handbook of Genetic Algorithms*, Reinhold, New York, 1991.

Glover, F., Tabu search : A tutorial, *Interfaces*, *20*, 74-94, 1990.

Goffe, W.L., Ferrier, G.D. and Rogers, J., Global optimization of statistical functions with simulated annealing, *J. Econometrics*, *60*, 65-99, 1994.

Hansen, P. and Jaumard, B., Algorithms for the maximum satisfiability problem, *Computing*, *44*, 279-303, 1990.

Ingber, L., Simulated annealing : Practice versus theory, *J. Mathl. Comput. Modelling*, *18*, 29-57, 1993.

Ingber, L. and Rosen, B., Genetic algorithms and very fast simulated reannealing : A comparison, *J. Mathl. Comput. Modelling*, *16*, 87-100, 1992.

Karr, C. L., Air-injected hydrocyclone optimization via geneitc algorithm, in *Handbook of Genetic Algorithms*, edited by L. Davis, Reinhold, New York, 1991.

Kirkpatrick, S., Gelatt, C.D. and Vecchi, M.P., Optimization by simulated annealing, *Science*, *220*, 671-680, 1983.

Kuo, C., Michel, A.N. and Gray, W. G., Design of optimal pump-and-treat strategies for contaminated groundwater remediation using the simulated annealing algorithm, *Adv. Water Res.*, *15*, 92-105, 1992.

Lockwood, C. and Moore, T., Harvest scheduling with spatial constraints : A simulated annealing approach, *Can. J. For. Res.*, *23*, 468-478, 1993.

Mayer, D.G., Belward, J.A. and Burrage, K., Use of advanced techniques to optimize a multi-dimensional dairy model, *Agric. Syst.* (in press), 1995.

Mayer, D.G., Kelly, A.M. and Butler, D.G., A review of optimisation techniques for non-linear simulation models, Proceedings 8th Biennial Conference, Simulation Society of Australia, Canberra, 229-234, 1989.

Mayer, D.G., Schoorl, D., Butler, D.G. and Kelly, A.M., Efficiency and fractal behaviour of optimisation methods on multiple-optima surfaces, *Agric. Syst.*, *36*, 315-328, 1991.

Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, A. and Teller, E., Equation of state calculations by fast computing machines, *J. Chem. Phys.*, *21*, 1087-1090, 1953.

Radcliffe, N. and Wilson, G., Natural solutions give their best, *New Scientist*, *126*, 35-38, 1990.

South, M.C., Wetherill, G.B. and Tham, M.T., Hitch-hiker's guide to genetic algorithms, *J. Appl. Stats.*, *20*, 153-175, 1993.

Styblinski, M. A. and Tang, T. S., Experiments in nonconvex optimization : Stochastic approximation with function smoothing and simulated annealing, *Neural Networks*, *3*, 467-483.

Syswerda, G., Schedule optimization using genetic algorithms, in *Handbook of Genetic Algorithms*, edited by L. Davis, Reinhold, New York, 1991.