BETTER CARBON SAVING: USING A GENETIC ALGORITHM TO OPTIMISE BUILDING CARBON REDUCTIONS

Christopher Pountney
AECOM UK
christopher.pountney@aecom.com

ABSTRACT

It has been common practice to develop a Marginal Abatement Cost (MAC) model to determine the cost minimal route to building-level CO_2 savings. This paper highlights several inherent limitations with this approach, not least the synergistic effects observed in applying measures sequentially. It demonstrates that a Genetic Algorithm (GA) can be used to provide a more accurate, detailed set of optimal solutions, whilst overcoming some of the shortcomings of MAC curve analysis. The SPEA2 GA is used in conjunction with SBEM and a set of capital costs to demonstrate the optimisation process using an office test case example.

INTRODUCTION

The recent financial crisis has brought into sharp focus the difficulties faced by the UK Government in setting bold carbon-reduction targets in new buildings. It is now clearer that environmental improvements cannot be achieved lightly; long-term paybacks are countered by immediate financial implications, intensifying the search for cost-minimal carbon-reductions.

At the macro-scale, it has been common to develop a Marginal Abatement Cost (MAC) model to determine the cost minimal route to aggregated carbon savings across the building stock. Typically, building carbon emission performance is measured in terms of annual CO_2 emissions. Given a baseline building emissions rate, the effects of varying the building technology specification can be plotted against the associated change in cost. A marginal abatement cost (MAC) curve orders these technology alternatives from lowest to highest cost.

As noted by Ellerman and Decaux (1998), MAC curves are often used heuristically in the context of international emissions trading. After constructing curves for each trading partner, the total cost for a given overall reduction can be minimised by equalising the individual marginal abatement costs (see Figure 1, where cost P gives a CO_2 reduction equal to Q1+Q2). This is an extension of a standard economic efficiency: An output maybe produced for the lowest overall cost by selecting the producers with the lowest marginal costs (Sheeran, 2006).

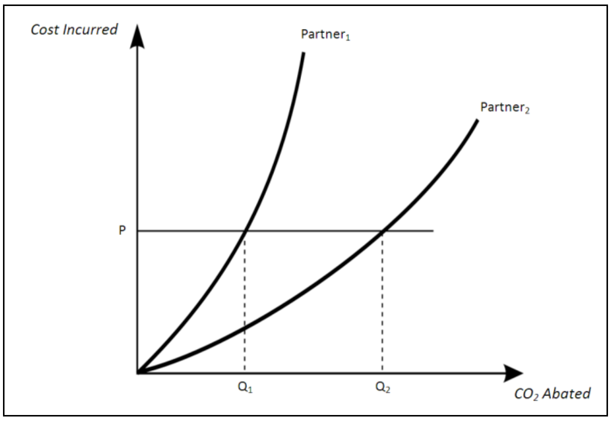

Figure 1 Equalising Marginal Abatement Costs (adapted from Ellerman and Decaux, 1998)

This approach translates directly to the national built environment scale, where the trading partners become individual weighted building types. Following carbon emissions modelling of typical building types, MAC curves can be developed where the building improvement measures are segments along the MAC curve. Characterising these curves, introducing the relevant weighting factors and equalising the MAC gives a straightforward relationship between total national cost and total national saving. As before, equalising the MAC sets different savings targets for each building type. The technologies required to achieve the saving in each building type (the 'specification') can be read directly from the MAC curve.

However the process of equalising MACs raises three particular concerns. Firstly, since it is difficult to model all the technology permutations across each building type, there is inevitably some uncertainty in the MAC data. Secondly, the MAC curve provides only a single route to cost minimal savings. Thirdly, adopting individual technologies results in discrete savings and may not provide sufficient resolution.

One recent study (DCLG, 2011) considered seventeen building types, incrementally varying up to twenty technologies across each type. Clearly, an exhaustive search of all the possible solutions is not preferable, yet the adopted parametric approach gives only a limited understanding of the synergistic effects of technology combinations. Fleiter, et al. (2009)

suggest that most technology combinations influence each other; these interactions result in synergisms affecting (often reducing) the quantity of saving potential. Thus the ordering of technologies used to create the MAC curves contains a degree of uncertainty, as does the precise quantities of CO_2 abated. Kesicki and Ekins (2011) go on to conclude that the failure to handle synergistic effects significantly impairs the usefulness of MAC curves.

As a solution to this problem, Stoft (1995) suggested recalculating the savings following the selection of a technology. Choosing the most cost-effective option at each stage removes the synergisms and finds the least-cost specification. This dramatically increases the computation required, yet even when it is possible, Fleiter, et al. (2009) recognised that the resulting solutions may not be practicable. This underlines the second issue with MAC curve analysis: In any real-world situation decision-makers prefer choice and the cost minimal solution provides no 'next-best' strategy.

Thirdly, since few technologies can be partially adopted, interpolation between points on the MAC curve has little physical meaning. This raises difficulties when needing to achieve specific CO_2 reductions, since there is no guarantee that the cost-minimal solution will give rise to a suitable specification.

In an attempt to address these issues, this paper proposes a Genetic Algorithm-based technique as an alternative to MAC curve analysis. As a subset of Evolutionary Algorithms, genetic algroithms (GAs) draw on principles developed in biological genetics to perform probabilistic searches of search spaces. This generate-and-test procedure has been shown to be particularly effective at optimising problems with poorly understood search spaces, and where approximate solutions are sufficient (Mitchell, 1996). Both these characteristics are true of the built environment MAC problem outlined above. In the first instance, the effect of varying technology specifications on energy performance is very difficult to predict on a building-by-building basis, depending on a wide-range of factors including building design, usage patterns and geographical location. Necessarily this results in a complex, undefined search space. Secondly, approximate accuracy matches the level of estimation involved in modelling the building-level carbon emissions.

This paper is structured as follows. First, there is a short overview of multi-objective optimisation, GAs and their application to similar problems. With this background, the proposed simulation methodology is outlined and details of the test case are given. The results are collated and compared to the equivalent MAC curves. The discussion and accompanying conclusions detail to what extent the GA methodology is able to address the three MAC curve limitations outlined above.

BACKGROUND

Most real-world problems involve multiple, competing design objectives; hence, in the field of computational optimisation these are referred to as multi-objective problems. Initially multiple objectives were aggregated into a single function, though these were demonstrated to be problem specific and difficult to set in advance (Richardson, et. al, 1989). Recognising that complementary objectives are forced to compete in the search for a single optimal solution, Schaffer and Grefenstette (1985) developed a vector measure of performance that quantified multi-objective behaviour. Although there were technical shortcomings to their approach, further development of multi-objective optimisation quickly followed. This culminated in the identification, by Goldberg (1989), of 'pareto-optimality' as a suitable method by which to rank solutions to multi-objective optimization problems.

Pareto-optimality states that no improvement in one objective can be achieved without a corresponding deterioration in at least one other objective. Given this, the primary aim of an optimisation routine is to find a set of solutions that approach the pareto-optimal set. Plotting the pareto-optimal set of solutions in objective space gives a surface known as the pareto front.

Yet there is also a secondary optimisation aim. Zitzler, Laumanns and Thiele (2002) concisely highlight the dual purpose, stating that "the distance to the optimal front is to be minimized and the diversity of the generated solutions is to be maximized." Not only should the pareto-optimal set be approximated as closely as possible, but the solutions should be spread out over the pareto front. A well-spread solution set is important as it gives a clearer picture of the trade-offs resulting from the selection of one solution ahead of another (Konak, Coit and Smith, 2006).

Genetic algorithms

One of the most common optimisation techniques invoked to seek the pareto-optimal set is the GA. With a GA, each solution to the optimisation problem is an individual and a set of individuals constitutes a population. Typically, GAs are 'generational'; a new population is developed by preserving the qualities of the best individuals from previous populations. Over time, the fittest individuals will tend to dominate future populations. To complete the metaphor, each individual is fully described by a chromosome. Each element of the chromosome is a gene, with each gene taken from a set of possible alternatives, the allele.

To begin the optimisation a population of individuals is instantiated; this is the first generation. This initial population will be either randomly created or seeded to provide an initial bias. Next, the fitness of each individual must be evaluated so that the better solutions can be identified. The individuals in the next generation are created by performing genetic operations on the fitter individuals in the current generation.

Various genetic operators have been developed; the most basic device is the ‘crossover’. In this case, two individuals (the parents) are selected and the genes taken in turn at random from one or other of the parents. This genetic data forms a new individual, the offspring. Since the parents are selected with a preference toward the fitter individuals, the offspring is expected to inherit the genes that will in turn lead it to be fitter also.

The second essential genetic operator is ‘mutation’. While crossovers maintain genetic fitness, mutation increases genetic variety; this is essential to avoid convergence at false optima. Typically, mutations will only need to be small modifications of specific genes. Nevertheless, depending on the application this can be sufficient to enable an individual to make a large jump through the objective space (Louis and Rawlings, 1993).

Having created the next population, the fitness of the new set of individuals will then be evaluated. In some cases, this generational search can continue until a set convergence criteria are met, otherwise the algorithm is stopped once a target number of generations is reached.

Related work

Various studies have incorporated GAs into building energy optimisation problems. Kusiak, Tang and Xu (2011) analysed the potential energy savings when controlling supply air temperature and pressure in an existing HVAC system. They were able to show significant savings, whilst maintaining indoor air quality above a predefined threshold. Wright, Loosemore and Farmani (2002) considered the relationship between thermal discomfort and energy costs in buildings with different fabric standards, demonstrating that GA-based optimisation appears well suited to building design problems.

More directly related to this paper is the work of Fialho, Hamadi and Schoenauer (2011). Using the EnergyPlus building simulation engine, they developed an optimisation methodology using an existing GA. Despite reporting on preliminary work, they concluded that the process had the potential to more quickly deliver energy-efficient building designs when compared with existing design practices. Furthermore, they explicitly recognised the value of the optimisation methodology to the decision maker. As previously discussed, the set of non-dominated solutions makes the trade-offs between alternate solutions clear, giving the decision maker the ability to make a more informed choice. As a basis for their work, only three decision variables were included in the optimisation problem: building orientation and two insulation materials. The authors anticipated increasing this number in future work to allow for a fuller building design optimisation.

Finally, of relevance is the work of Rabotyagov, Jha and Campbell (2010). Although in this instance the authors were concerned with improving water quality rather than building design, there are clear parallels with the focus of this paper to the extent that both studies seek the cost-efficient allocation of abatement activities.

The intricate relationship between human actions and environmental impacts inevitably forms a complex objective space. Noting that existing approaches were unable to handle the optimisation of both the spatial and hydrologic features, the authors identify a GA-based approach as being suitable for efficient solution seeking. Furthermore, they recognise that the pareto fronts found via the optimisation routine represent pollution abatement cost curves. They conclude: “In essence, the simulation-optimization framework employed allows us to make use of well-known environmental economics prescriptions of equalizing marginal abatement costs across pollution sources”. Having created the abatement curves the authors then select particular individuals from the pareto set to find real-world solutions that deliver a required pollution abatement.

The rest of this paper intends to extend this line of thinking to show that a GA-based optimisation approach is similarly appropriate to the cost-efficient allocation of carbon abatement efforts in building design. Moreover, not only is it suitable but it also more appropriate than MAC curve analysis.

SIMULATION

This section describes the approach taken to demonstrate the applicability of GAs to the building-level CO_2 optimisation problem. It begins by describing the selected GA and explaining on what basis the fitness is assigned. This is developed with a further discussion of SBEM, the adopted energy simulation engine. Finally, the complete optimisation routine is summarised and the test case is described.

SPEA2

To demonstrate the applicability of GAs to the built environment MAC problem, a model was constructed implementing the second Strength Pareto Evolutionary Algorithm (SPEA2) (Zitzler, Laumanns and Thiele, 2002). SPEA2 optimises a solution set using a fitness function that incorporates both pareto dominance and an assessment of the solution density at that point in the objective space.

To measure the pareto dominance each individual is assigned a strength corresponding to the number of other individuals it dominates. The raw fitness of an individual is the sum of the strengths of the individuals that dominate it, such that a non-dominated individual has a raw strength equal to zero.

This property alone however, is not able to indicate the spread of the solution set, nor can it provide direction to the search if many of the individuals are non-dominated. Zitzler, Laumanns and Thiele (2002) recognised that such a scenario is particularly likely to occur when the optimisation involves more than two objectives. For both of these reasons, they implemented a density estimator. The density measurement is a function of Euclidean distance in the objective space calculated using the *k*-th nearest neighbour method. The fitness function is a simple sum of the raw fitness and the density estimator.

More recently, there has been significant effort to develop GAs that improve on SPEA2. Much of the focus has been on the hypervolume indicator (Bader and Zitzler, 2011), a parameter that varies directly with pareto dominance. Indeed, Wagner, Beume and Naujoks (2007) identified that the quality of SPEA2 declined rapidly when increasing the number of objectives. Nonetheless, since the building-level CO_2 optimisation problem is a biobjective optimisation, SPEA2 has been used since it is both straightforward to implement and requires relatively low computational overheads. In this case, the position of an individual in objective space and hence its fitness, depends on the calculated annual carbon emissions and the corresponding capital cost.

Evaluating annual carbon emissions

To evaluate annual carbon emissions, the building energy demands are calculated using the Simplified Building Energy Model (SBEM) engine. The SBEM engine uses the standards listed in CEN Standard PG-37 to calculate monthly energy use. The SBEM input file contains the building description. This is updated for each individual to be tested and run through the SBEM engine. The resulting annual carbon density ($kgCO_2/m^2$) is recorded by interrogating the appropriate SBEM output file.

Although the analysis in SBEM is not as detailed as other simulation engines, SBEM is particularly well suited to GA-based simulation since the building models are described succinctly in a single, editable text file. This is advantageous as multiple parameters in the building definition need updating according to each individual’s chromosome. Furthermore, the energy calculations performed by SBEM are in accordance with the National Calculation Methodology (NCM). This readily allows carbon savings to be calculated in terms of reductions over UK Building Regulations. Lastly, the runtime of a typical SBEM model is in the order of less than a minute, so that many generations can be analysed relatively quickly.

At this stage, the only LZC technology included in the building description is PV. The quantity of PV is based on a percentage of roof area and the carbon savings calculated offline. A more complete building description would include a range of heat generating LZCs, some of which can be modelled directly in SBEM. However, this additional complexity was unnecessary at this stage to demonstrate the applicability of the GA-based approach to the CO_2 optimisation problem.

Evaluating total capital cost

The capital cost of an individual solution is determined by summing the costs associated with each gene in the chromosome. The capital costs are based on cost data gathered from a number of sources, primarily through direct contact with manufacturers. The costs of certain allele sets (i.e. lighting efficiency, fan coil unit efficiency, etc.) required various design assumptions to ensure that the costs could be derived as a function of the total floor area.

The allele set

As described above, each gene is taken from a corresponding allele. The group of alleles necessary to select an entire chromosome is the allele set. Following the recent DCLG analysis (2011), the allele set selected here describes the building fabric and HVAC efficiencies, as well as the size of the PV array. Table 1 gives the full list of alleles.

Table 1
The allele set

NAME	OPTIONS
Roof U-value (W/m^2K)	0.25, 0.20, 0.15, 0.10
Floor U-value (W/m^2K)	0.25, 0.20, 0.15, 0.10
Wall U-value (W/m^2K)	0.35, 0.25, 0.20, 0.15
Window U-value (W/m^2K)	2.0, 1.6,1.3, 0.9
Air Permeability (m^3/m^2hr)	7.5, 3.0, 1.0
Lighting Efficacy (lm/W)	55, 60, 65
Lighting Occupancy Control	OFF, ON
Lighting Daylight Control	OFF, ON
Boiler Efficiency (%)	84, 86, 88, 91
Heat Recovery Efficiency (%)	0, 40, 50, 70
Chiller Efficiency	2.5, 3.0, 3.5, 4.5
AHU Specific Fan Power (W/l/s)	2.2, 2.0, 1.8
FCU Specific Fan Power (W/l/s)	0.6, 0.3
PV Installation (% Roof Area)	13, 27, 40, 50

Outline of optimisation routine

The details of the GA implementation are described in pseudo-code format in Figure 2. An interface was written to create the SBEM models using the chromosomes produced by the GA. The interface then passed the collated costs back to the GA to generate the next generation.

It is important to emphasise the role of the archive in the algorithm. The archive is used to retain the best solutions from both the current generation and the preceding generations, eventually filling with non-dominated individuals. This is the record of the most optimal solutions discovered by the algorithm. It is from the archive that the future populations are derived. Therefore, if all the individuals in the next population are less fit than those in the current

Step 1: Create a random population (P_1) and an empty archive population (A_1).

[For each generation, n]

Step 2: *[For each chromosome, C_i, in P_n]* Assign fitness.

- Update the SBEM input file with the contents of C_i. Calculate the consequent costs and run SBEM to find the carbon density.
- Assign a strength value equal to the number of other solutions C_i pareto-dominates. Set the fitness equal to the sum of the strengths of the solutions that dominate C_i.
- Assign a weighting value based on the inverse of the objective space distance to the *k*-th nearest neighbour.

Step 3: Environmental selection.

- Combine P_n with the previous archive population (A_{n-1}) and order according to fitness, using the weighting value where necessary.
- Copy all non-dominated individuals from the combined set to a new archive population, A_n. Truncate A_n where necessary.

Step 4: If the stopping criteria are met, then stop and exit algorithm.

Step 5: Create new generation.

- Tournament selection is used to choose individuals from A_n. Select two individuals from A_n and compare the fitness values, putting the individual that wins the tournament forward for the genetic operation.
- Perform genetic operations on individuals selected from A_n and copy the new individuals to the new population P_{n+1}. Repeat this until P_{n+1} is complete. Return the two parent individuals to A_n.

[Return to Step 2]

Figure 2 Pseudo-code description of the GA

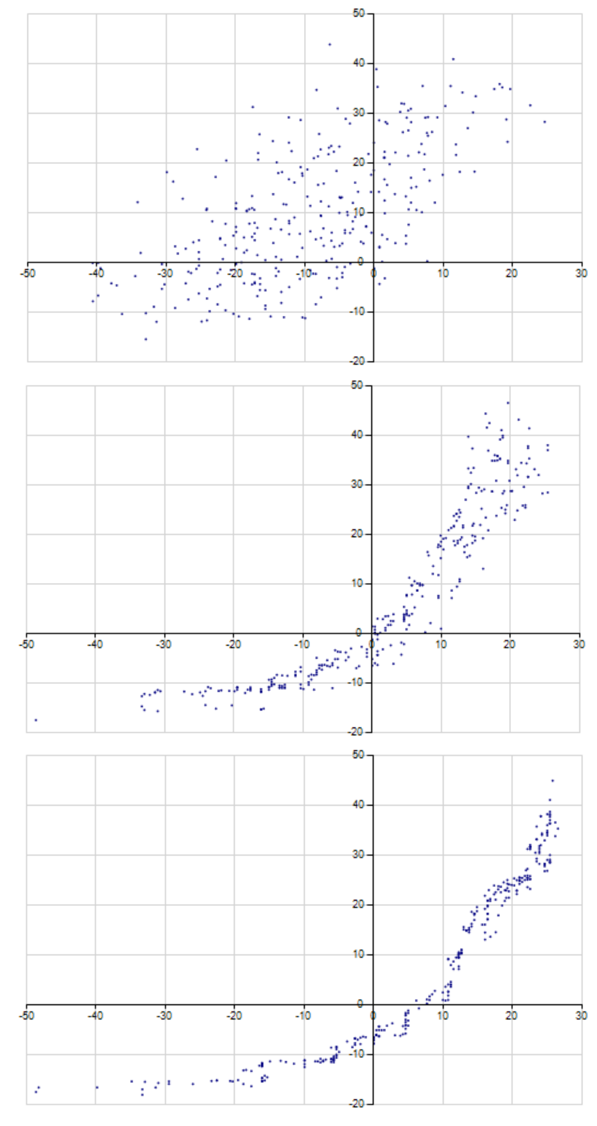

Figure 3 The archive after the 1st, 3rd and 5th generations

archive, the archive will remain unchanged and shall be used again to generate the next population.

Test Case: Office Building

The test case used to illustrate the GA-approach was a deep-plan, air-conditioned office building located in London. The model was created to represent a typical building of the type, with TUFA of 30,000 m^2 split over 10 floors. The facade assumed 40% glazing.

Following initial investigations, the GA was setup to include 300 individuals in each generation and run for 30 generations. In this instance, the two main genetic operators used were crossover (85% of the time) and mutation (10% of the time). A third operator, replacement (5% of the time), was included to introduce a new randomly generated individual. This aimed to maintain the genetic variation in the population. Each time the crossover or mutation operator was invoked, a random proportion of the parent individuals were altered.

Figure 3 illustrates the archive following the initial generations. The archive in the first generation is equivalent to the randomly generated first population and shows a good spread over the objective space. The following generations show the archive quickly moving toward the pareto front. For comparison a MAC curve was created using the same SBEM model and cost data. Assuming a nominal baseline specification, the impact on carbon density of each of the measures in the allele set is tested by varying them incrementally and running the associated SBEM model. Measures that show a carbon saving are plotted on the MAC curve.

For both the GA outputs and the MAC curves, the CO_2 savings are shown as percentage reductions over Part L 2010. Similarly, the costs are rebased on an assumed Part L 2010-compliance cost. This is derived from the 0%-reduction point on the MAC curve.

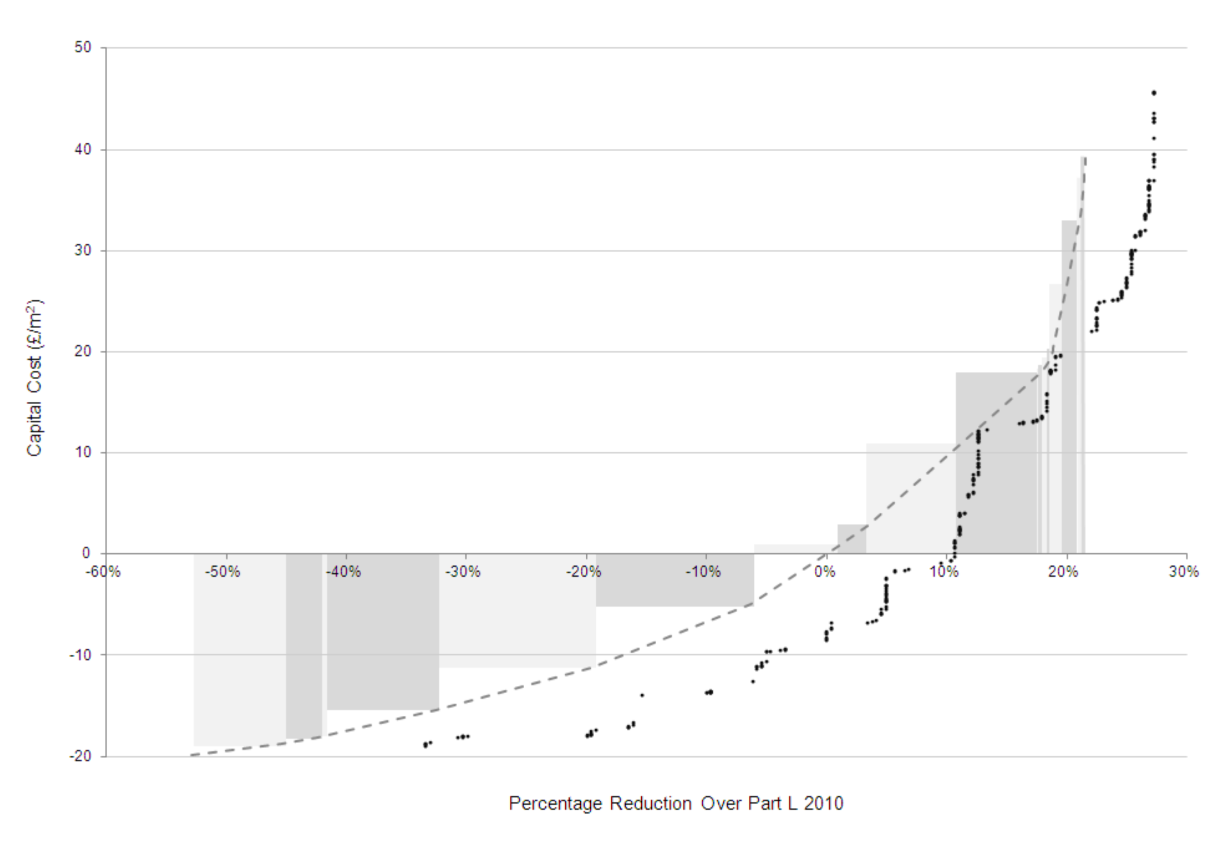

Figure 4 Comparing the MAC curve with the archive after the 30th generation

DISCUSSION

The MAC curve for the office test case is plotted in Figure 4. Each of the selected measures are represented by a grey rectangle, the corner of which is intersected by the MAC curve. Also shown in Figure 4 is the GA archive after the 30th generation. This section compares the results of the two approaches, noting several points of discussion.

Firstly, it is important to remember that each element in the archive represents a technology specification, that is, a physical building description. Thus, when the archive elements are to the right of the MAC curve there exist particular specifications that are more optimal than the MAC analysis suggests. As noted previously, the parametric approach used to construct the MAC curve has some inherent uncertainties, since the synergistic effects of adding technologies sequentially remain unresolved. The differences between the MAC curve and the GA archive in Figure 4 suggests that this synergistic effect is significant.

For example, consider the cost of Part L compliance. The GA identifies several solutions close to 0% that cost 8 £/m^2 less than the MAC curve prediction. Interestingly, there are several intervals along the horizontal axis between which a particular CO2 reduction can be achieved at considerably less cost. Pushing the savings to the MAC curve limit, at ~22% CO2 reduction there are GA solutions that save 16 £/m2.

To investigate further the discrepancy between the MAC curve and GA archive, the measures in the MAC curve were re-modelled to reduce the impact of the synergistic effects. Starting from the base case model, each measure was added in turn and the CO2 saving re-calculated. The resulting adjusted MAC curve is the solid line plotted in Figure 5. It will be noted that the smoothness of the curve is affected since adding the measures sequentially alters the cost-effectiveness in some cases.

This 2-stage modelling makes a clear improvement to the MAC curve and the maximum predicted CO2 saving corresponds directly with the GA archive. However, the remaining differences highlight a second issue with the MAC approach: Interpolating between the measures can lead to either over- or under-estimation of potential savings. This appears to apply to the measures shown on the MAC curve between 24-27% CO2 saving. Where the adjusted MAC curve predicts a cost of 18 £/m^2 for 24% CO2 reduction, the best archive solution is 7 £/m^2 more expensive. As noted previously, interpolating along a MAC curve may not result in realistic or feasible solutions. These results demonstrate that the predicted savings may not be achievable either.

Consequently, the GA approach has proven to give a greater number of routes to optimal CO2 savings. The design decision-maker is able to take a range of

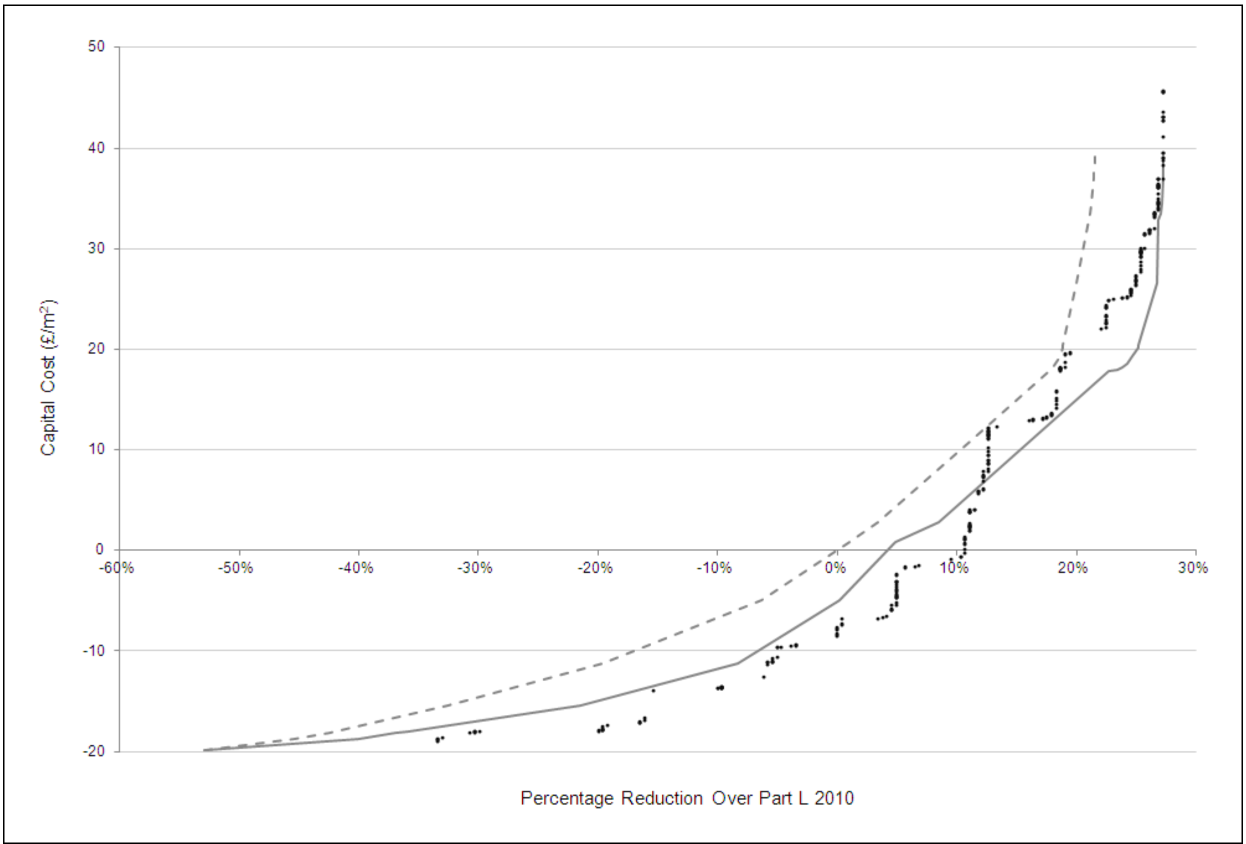

Figure 5 Comparing the MAC curve, the adjusted MAC curve and 8the Archive after the 30th generation

specifications forward to identify a preferred option.

Nonetheless, the GA solution set does not provide a complete picture. Although there is a good deal of spread in the archive, there are still regions where fewer solutions are observed. These appear in two forms. Firstly, there are flatter portions of the front, such as -29% along to -20% CO_2 saving. In these regions, there are no solutions in the archive; indeed, it is possible that such solutions do not exist. In the second place, however, there are intervals where the archive set is vertically separated, for example 19%-22% CO_2 saving. It is informally observed when plotting every tested solution, there appear to be solutions on the pareto front that are not included in the final archive. Further investigation into improving the quality and extent of the archive is required. This should seek to identify the importance of a larger archive size, weighing increased computational overhead with a more complete approximation to the pareto front. In this connection, development of the density estimator to further promote spread along the pareto front would also be beneficial.

There is clearly also further work required in developing test cases based on alternative building types. Although there are no apparent methodological differences in using the GA to model different building types, it is necessary to introduce this variation to ensure the conclusions of this study are more widely applicable.

Similarly, this investigation is unable to identify the importance of either more accurate cost data or a more detailed energy model. However, since the base data and the energy model are common to both the MAC and GA-based approaches, it is anticipated that the conclusions presented here are independent of changes to either. Again, further work is required to demonstrate this.

Finally, a brief comment on computational requirements is necessary. The MAC curve analysis required < 100 SBEM models, whereas the GA ran 9,000 SBEM models. Nonetheless, the time requirement of the GA was still manageable: Using a standalone quad-core laptop the simulation runtime was 18.5 hours. This has the potential to be much improved however, since the individual SBEM calculations can be efficiently distributed over a series of networked computers. This is currently under development and should make the GA analysis even more suitable.

CONCLUSION

This paper proposes that the three shortcomings inherent in MAC analysis are at least partially addressed through a GA-based approach. Firstly, the synergistic effects observed in MAC analysis resulting from implementing the selected technologies sequentially are avoided. The test case demonstrates that in many cases more optimal solutions existed than the MAC analysis predicted. This discrepancy is shown to be reduced by partially accounting for the synergistic effects through the 2-stage re-modelling.

However, comparing the adjusted MAC curve with the GA archive demonstrates the dangers of interpolating between MAC curve measures. In this test case, there are clear examples where CO_2 savings appear more costly in practice than when predicted by the adjusted MAC curve.

In the third place, the larger set of optimal solutions found by the GA analysis is a distinct advantage. Since interpolating along the MAC is often meaningless in reality, having a set of solutions at smaller cost intervals gives greater flexibility when identifying the most suitable specifications required to meet specific carbon reduction targets.

These conclusions provide a good basis for replacing MAC analysis with a GA-based approach when determining cost-minimal carbon emissions in new buildings.

ACKNOWLEDGEMENT

This paper is based on an essay originally prepared for the Interdisciplinary Design for the Built Environment masters course at the University of Cambridge. The supervisor of the essay was Dr Alan Blackwell.

REFERENCES

Bader, J. and Zitzler, E. 2011. HypE: An Algorithm for fast Hypervolume-Based Many-Objective Optimisation. Evolutionary Computation, 19(1), 45-76.

Department of Communities and Local Government (DCLG), 2011. Zero Carbon Non-domestic Buildings: Phase 3 Final Report, Department for Communities and Local Government.

Ellerman, A.D. and Decaux, A. 1998. Analysis of Post-Kyoto CO_2 Emissions Trading Using Marginal Abatement Curves. MIT Joint Program on the Science and Policy of Global Change, Report No. 40.

Fialho, A., Hamadi, Y. and Schoenauer, M. 2011. Optimizing Architectural and Structural Aspects of Buildings towards Higher Energy Efficiency. In Proc. GECCO ’11, 727-737.

Fleiter, T. W., Eichhammer, M., Hagemann, M., Wietschel, M., and Hirzel, S. 2009. Costs and Potential Energy Savings in European Industry - A Crtitical Assessment of the Concept of Conservation Supply Curves. ECEEE 2009 Summer Study, La Colle sur Loup, France.

Goldberg, D.E. 1989. Genetic Algorithms in Search, Optimization and Machine Learning. Boston, MA: Addison Wesley.

Kesicki, F. and Ekins, P. 2011. Marginal Abatement Cost Curves: A Call For Caution. Climate Policy, 12(2), 219-236.

Konak, A., Coit, D.W. and Smith, A.E. 2006. Multi-objective Optimization Using Genetic Algorithms: A Tutorial. Reliability Engineering & System Safety, 91, 992-1007.

Kusiak, A., Tang, F. and Xu, G. 2011. Multi-objective Optimization of HVAC System with an Evolutionary Computation Algorithm. Energy, 36, 2440-2449.

Louis, S.J. and Rawlins, G.J.E. 1993. Pareto Optimality, GA-easiness and Deception. In Proceedings of the Fifth International Conference on Genetic Algorithms.

Mitchell, M., 1996. An Introduction to Genetic Algorithms. MIT Press: Cambridge, MA.

Rabotyagov, S., Jha, M. and Campbell, T. 2010, Searching for Efficiency: Least Cost Nonpoint Source Pollution Control with Multiple Pollutants, Practices and Targets. Journal of Natural and Environment Sciences, 1(2), 75-90.

Richardson, J.T., Palmer, M.R., Liepins, G. and Hilliard, M. 1989. Some Guidelines For Genetic Algorithms With Penalty Functions. In Proceedings of the Third International Conference on Genetic Algorithms, Morgan Kaufmann: San Mateo, CA.

Schaffer, J.D. and Grefenstette, J.J. 1985. Multi Objective Learning Via Genetic Algorithms. In Proceedings of the Ninth International Joint Conference on Artificial Intelligence, 593-595.

Sheeran, K.A., 2006. Who Should Abate Carbon Emissions? A Note. Environmental & Resource Economics, 35, 89-98.

Stoft, S. 1995. The Economics of Conserved-Energy 'Supply' Curves. Energy, 16(4), 109-140.

Wagner, T., Beume, N. and Naujoks, B. 2007. Pareto-, Aggregation-, and Indicator-Based Methods in Many-Objective Optimization. In Lecture Notes in Computer Science, 4403, 742-756.

Wright, J.A., Loosemore, H.A. and Farmani, R. 2002. Optimization of Building Design and Thermal Control by Multi-Criterion Genetic Algorithm. Energy and Buildings, 34,959-972.

Zitzler, E., Laumanns, M. and Thiele, L. 2002. SPEA2: Improving the Strength Pareto Evolutionary Algorithm for Multiobjective Optimisation. In Evolutionary Methods for design, Optimisation and Control with Application to Industrial Problems (EUROGEN 2001), 95-100.