OPTIMISATION FOR CHP AND CCHP DECISION-MAKING

Ralph Evins[1,2], Philip Pointer[1], Ravi Vaidyanathan[2]
[1]Buro Happold, 17 Newman Street, London, W1T 1PD, UK.
[2]University of Bristol, Tyndall Avenue, Bristol, BS8 1TH, UK.

ABSTRACT

We present a new analysis and optimisation procedure to aid decision-making regarding Combined Heat and Power (CHP) and Combined Cooling, Heat and Power (CCHP) installations. Our holistic model incorporates analysis of plant operation (including part-load performance) and provides guidance regarding applicability, sizing and phasing of plant.

A multi-objective genetic algorithm has been used to optimise a set of possible configurations. This produces a "trade-off front" of solutions. The outputs are reported for a case study. Additionally, a wide range of scenarios have been optimised and the outputs examined graphically to derive innovative design guidelines (a process known as "innovization").

INTRODUCTION

Combined Heat and Power

A Combined Heat and Power (CHP) system allows financial and carbon savings by making use of the heat produced when electricity is generated, which is usually wasted. The heat may be used to meet the thermal demands of a development, for example for space heating and domestic hot water, or used to run absorption chillers to provide cooling, known as a Combined Cooling Heat and Power (CCHP) system or tri-generation.

Great care must be taken in sizing a CHP system to match the demands of a development, and in particular the profiles of demand fluctuations. There is a minimum load, usually 50% of the maximum load, below which a CHP engine cannot run. Therefore the system must largely be used to supply a base load which is present for a large part of the time, perhaps 8 hours per day. A thermal store (TS), usually a hot water tank, can be used to buffer these fluctuations, but a large thermal store is expensive and takes up a lot of space. Also, CHP engines should not be started and stopped frequently; an average of one startup per day is recommended. If a system is too large, it will not operate often enough; if a system is too small, it will not be providing the full potential carbon and cost savings.

The general system under consideration consists of two CHP engines, a thermal store, an absorption chiller (if CCHP), a gas boiler used to meet any remaining heat demand, and grid electricity used for

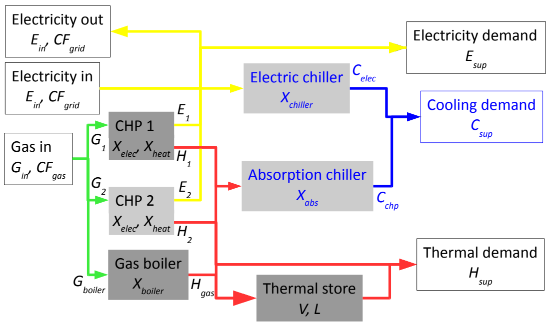

Figure 1: Schematic of the CHP system under consideration.

unmet electrical demand and unmet cooling demand via electric chillers (if CCHP); surplus electricity may be sold to the grid. This is illustrated in Figure 1.

Multi-objective Optimisation

Computational optimisation is a rapidly emerging discipline for aiding engineering design. Multi-objective optimisation is particularly useful as it involves the consideration of several objectives simultaneously, with no weightings or aggregations, allowing the robust resolution of complex trade-offs between conflicting objectives. This involves finding the *non-dominated-* or *Pareto-front*, a set of points in the objective space for which no point performs better in all objectives (see Figure 2).

Previous work

Ooka and Komamura (2008) developed a two-stage design process using genetic algorithms to simultaneously optimise plant capacities and operational details, applied over one day. Li et al. (2006) looked at the configuration of a CCHP system to maximise Net Present Value using a genetic algorithm. Tanaka et al. (2007) optimised plant configuration and operation using a genetic algorithm. Song et al. (1999) and Vasebi et al. (2007) investigated the CHP dispatch problem (see next section) using ant colony optimisation and harmony search respectively.

ANALYSIS OF CHP OPERATION

Predicted loads

Standard daily demand profiles have been used for heating, cooling and electricity demands (see Figure 3). Each profile consisted of a base component and a weather-dependent component which is scaled

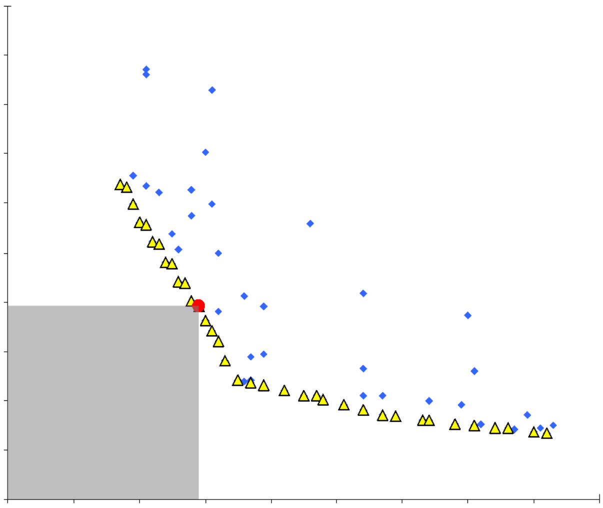

Figure 2: An example of a Pareto front for the minimisation of two objectives, one on each axis. Triangles are members of the Pareto front; dots are not. For the highlighted point there is no point within the shaded area, therefore it is non-dominated; in this work, this corresponds to there being no solution which has both lower emissions and lower costs.

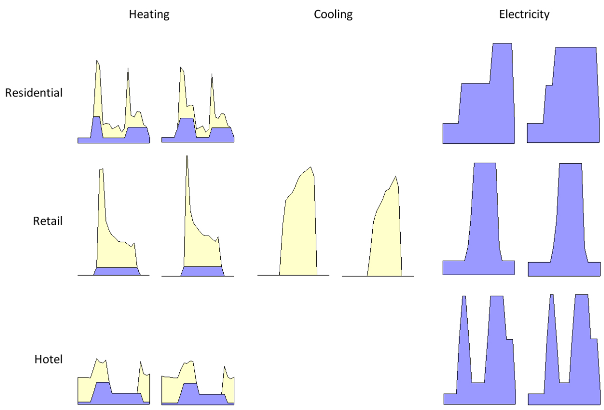

Figure 3: Daily load profiles for heating, cooling and electricity for three sector types. Blue areas are the base profile; yellow areas are weather dependent. The first graph of each pair is for weekdays, the second for weekends.

Table 1: Energy use benchmarks (KWh/m^2/year) and standard deviation σ used for profile diversity.

	Heating	Cooling	Electricity	σ
Residential	250	0	50	4
Retail	100	15	200	1
Hotel	220	0	78	2

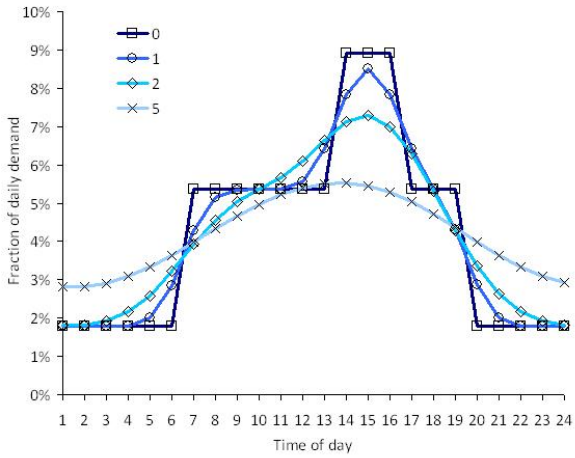

Figure 4: Example of diversity applied to a daily load profile. Each curve shows a different value of σ.

according to the temperature difference (taken from an annual series of daily averages) from a set point (below 15C for heating, above 18C for cooling). Different profiles have been used for each sector type (in the case study these are residential, retail and commercial) and for weekdays and weekends. Site-wide hourly demand profiles for heating, cooling and electricity were formed by summing the demands for each sector, scaled according to the area and benchmark values for annual demand (given in Table 1).

Diversity between different demands can have important implications in CHP design, as diverse demands will smooth peaks and troughs leading to lower maximum demands and a higher continuous baseline demand. Diversity has been introduced to the demand modelling described above by applying a normal distribution to all values in the daily demand profiles (see Figure 4). Each hourly value of the new profile x'_i is the sum of the original profile value x_i multiplied by the probability density function of the normal distribution with the mean at hour i and standard deviation σ (see Equation 1); three distributions offset by 24 hours are used to allow the profile to wrap around when bridging midnight. Different values for the standard deviation of the distribution have been used for each sector type (given in Table 1); retail has a low diversity as opening hours will be similar, whereas residential has high diversity as people get up and go to bed at very different times.

$$x'_i = \sum_{j=1}^{24} \left(x_i \frac{1}{\sqrt{2\pi\sigma^2}} \sum_{k=-24,0,24} \left(e^{-\frac{(j+k-i)^2}{2\sigma^2}} \right) \right) \quad (1)$$

Hourly simulation of plant performance

The site-wide hourly demand profiles for heating, cooling and electricity have been used as the input for the CHP operation algorithm, along with plant and

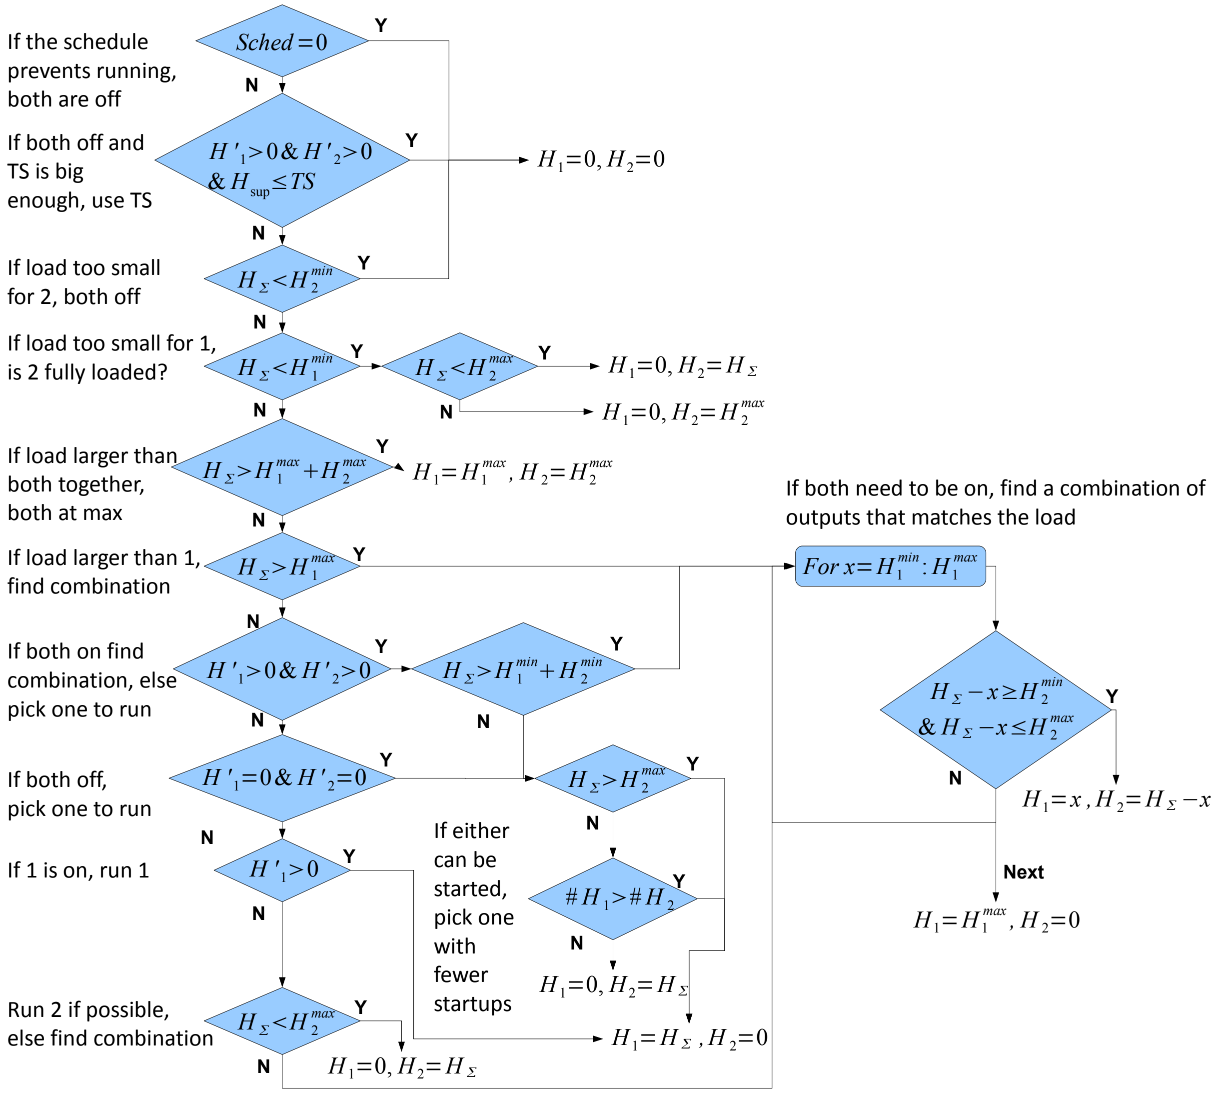

Figure 5: Control logic for CHP 1 and CHP 2. The algorithm determines heat outputs H_1 and H_2 for the two machines based on the availability schedule $Sched$, the heat supplied H_{sup}, the thermal store contents TS, the thermal demand including thermal store deficit $H_\Sigma = H_{sup} + (TS_{max} - TS)$, and the state of each machine in the previous hour H_1' and H_2', as well as the operating ranges of each machine H_1^{min}, H_1^{max}, H_2^{min} and H_2^{max}.

thermal store capacities. It is desirable to model the CHP system hourly (or better) since sharp peaks of short duration can have a great effect on performance (see Hawkes and Leach (2005)). It is desirable to model a full year of operation in order to assess performance over the whole range of expected demands (which vary based on the weather).

The main function of the CHP operation algorithm is to determine whether and at what load the generating plant will be operational. Rather than optimise the many parameters of an operational schedule (as for example Ooka and Komamura (2008)) or address the full CHP dispatch problem relating to balancing thermal and electrical demands with efficiency (as for example Vasebi et al. (2007)), a number of assumptions have been made to allow a fixed (though complicated) control logic to be used.

The first assumption is that surplus electricity can be sold to the grid for a reasonable price (and credit can be taken for the associated carbon savings) during peak hours (7am - midnight). This bypasses the dispatch demand problem as it will always be desirable to run at as high a load as the thermal demand permits. The second assumption is that efficiency differences between systems of different capacity and between full-load and part-load operation are small. This bypasses the efficiency drop-off problem, meaning it is always acceptable to run at part-load if this improves the contribution from CHP. The remaining requirement for the operation of CHP engines is the limit on the number of start-ups (taken as an average of once per day).

The control logic used (as shown in Figure 5) aims to maximise the CHP contribution to thermal demand whilst remaining within the machine operational limits and minimising the number of start-ups. Priorities of use differ depending on the previous state of the machines, for example if both are off, the load will be met by the thermal store if possible, whereas if any are on, they will be used to supply the load and charge the thermal store if possible. If it is necessary to start up one machine and it doesn't matter which one, the one with the fewest startups will be chosen. The availability schedule $Sched$ has been set to zero from midnight to 7am (when off-peak electricity prices make running uneconomic) and for two days per month and one week in the summer for maintenance (612 on-peak hours, ~10% of on-peak hours per year).

Environmental performance

Table 3: Financial and system inputs. Prices are given at year one of the project lifespan.

Price index: retail	4	%
Price index: electricity	5	%
Price index: heat	3	%
Price index: gas	3	%
Price index: construction	4	%
Price: electricity sold on-site	95	£/MWh
Price: electricity exported	80	£/MWh
Price: heat sold on-site	50	£/MWh
Cost: gas	37	£/MWh
Cost: grid electricity	95	£/MWh
Cost: CHP maintenance	10	£/MWh$_{elec}$
Network loss	5	%
Community boiler efficiency	85	%
Domestic boiler efficiency	88	%
Absorption chiller efficiency	100	%
Electric chiller efficiency	400	%
Carbon Factor (CF): gas	0.19	kgCO_2/kWh
CF: electricity from grid	0.43	kgCO_2/kWh
CF: electricity sold to grid	-0.52	kgCO_2/kWh

Using the control algorithm discussed above, an annual hourly series is constructed of thermal outputs from both CHP engines and gas boiler backup. The thermal outputs from CHP are converted into percentage loads, and the electrical outputs and fuel consumptions are found via interpolation between the values in Table 2. Grid import or export of electricity is calculated from the amount generated and the electrical demand. Fuel used and grid import and export are then summed for the whole year, and these totals are used to obtain the associated carbon emissions for the year using the appropriate carbon factors (see Table 3). Electricity exported to the grid is converted into a carbon emissions credit. The objective used for environmental performance was the carbon savings of the project over a baseline system (individual gas boilers and grid electricity), summed over the project lifespan.

Financial performance

The financial performance of the system has been evaluated by using a projected profit and loss approach over the project lifespan. Costs included capital expenditure, depreciation (calculated by dividing item cost by expected lifetime), maintenance, fuel and grid electricity; incomes included heat sales, electricity sales internally and electricity sales to the grid. All prices were subject to increase over time, governed by indices for retail costs, electricity, gas and construction costs (the percentage increase being applied cumulatively to the index). Table 3 gives details of the prices and index rates used. All costs and incomes were projected over the project lifespan, and a running balance kept. The objective used for financial performance was the cost saving of the project over the baseline system at the end of the project lifespan.

OPTIMISATION

There are many computational means of accomplishing multi-objective optimisation. One of the most widely applied is the genetic algorithm, and many other methods follow a similar approach. The Non-dominated Sorting Genetic Algorithm (NSGA-II) of Deb et al. (2002) used here is a very popular genetic algorithm for multi-objective optimisation. The algorithm maintains a population of possible solutions, each corresponding to a particular choice of input variables. Solutions may be changed to form new variations by crossover (combining features of two solutions) or mutation (randomly changing values). In this way a second population is formed, and solutions are selected to continue to the next generation from either population based firstly on non-domination rank[1] and secondly on crowding distance[2].

The two objective functions were the environmental objective and the financial objective as detailed above. Two constraints were imposed which limited the start-ups of each CHP to less than 365 per year. The variables used were CHP 1 unit (0 to 17) and CHP 2 unit (0 to 17) (see Table 2), Thermal Store (TS) size (15 to 150 m^3) and CHP or CCHP; there was an additional variable for the second study, CHP 2 construction year. The following NSGA-II parameter values were used: population size 20; number of generations 20; crossover probability 0.7; mutation probability 0.5.

RISK ANALYSIS

In order to better understand the CHP decision-making process and to provide an indication of risk for CHP

[1] Rank 1 solutions are the non-dominated front. These are removed and domination is recalculated to form a new front, which is given rank 2. This process continues until all solutions are ranked. This ensures that the algorithm progresses towards the true non-dominated front.

[2] A measure of the distance of a solution from its neighbors in the objective space. Solutions in less crowded regions are preferred, ensuring that the algorithm explores the whole front.

Table 2: CHP units available. Electrical output E_i, heat output H_i and gas used G_i are given for 100%, 75% and 50% load conditions.

Unit #	Capital cost, £	E_i, kW			H_i, kW			G_i, kW		
		100%	75%	50%	100%	75%	50%	100%	75%	50%
1	33,150	26	20	13	46	35	25	81	62	44
2	55,000	50	38	25	82	75	52	150	128	88
3	75,000	75	56	38	127	110	78	223	185	128
4	95,000	100	75	50	161	128	102	291	228	174
5	111,020	122	92	61	196	166	129	348	283	213
6	132,880	151	113	76	232	199	155	418	343	254
7	147,915	173	130	87	264	224	171	483	392	287
8	157,250	185	139	93	274	236	184	507	415	309
9	181,500	220	165	110	307	247	173	590	460	323
10	189,895	233	175	117	284	232	171	618	483	340
11	208,000	260	195	130	335	266	191	689	536	378
12	270,000	360	263	175	413	339	249	919	717	501
13	292,000	400	300	200	503	402	295	1055	822	581
14	353,000	500	375	250	608	490	367	1273	994	712
15	408,030	609	457	305	731	583	408	1559	1196	860
16	522,614	809	607	405	945	765	565	2057	1583	1090
17	622,000	1000	750	500	1216	980	734	2546	1988	1424

system decisions, six parameters of the system model have been altered and changes in the optimal solutions noted. The six parameters were project length, grid carbon factor, gas and heat price index, electricity price index, standard deviation of profile diversity, and demand for heat and power. Each parameter in turn was set to first 50% and then 150% of the original value. The performance of the *main solutions* found initially (Figures 8 and 13) was calculated for each of these scenarios, forming a normal sensitivity analysis. Additionally, a new optimisation run was conducted for each scenario, providing information on the *opportunity cost* of selecting each main solution in the context of each scenario. These results were analysed by finding the distance from each of the main solutions (using the peformance values for the relevant scenario) to the nearest optimal point for the new scenario. This goes beyond sensitivity analysis as it provides information on the performance of each solution relative to the optimal solutions for each scenario. However, the summary statistic - distance to the nearest optimal point - is only an indication of a single "better option"; for a more comprehensive answer, it would be necessary to visually compare the complete Pareto front obtained for each scenario with each of the main solutions.

CASE STUDIES

Study 1: Single-phase development

The first case study sought to optimise the CHP system for a development in which all buildings are constructed in a single phase at the start of the project. The development consisted of 30,000m^2 residential, 20,000m^2 retail and 10,000m^2 hotel. UK climate data was used.

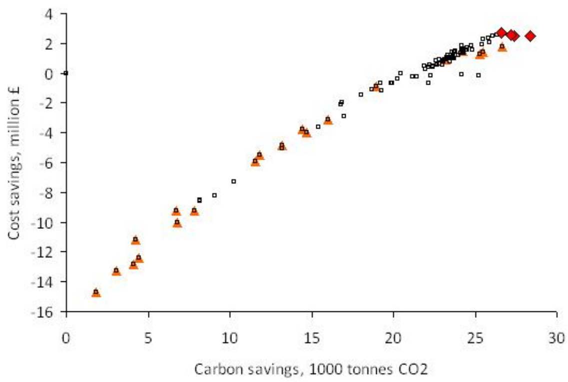

Figure 6: All solutions evaluated. Non-dominated solutions are shown in red. Initial random solutions are shown in yellow.

Figure 6 shows all solutions evaluated by the algorithm. There is a large degree of variability : carbon savings over the baseline range from 1 to 28 ktCO_2, and costs vary between £2.3m better than the baseline after 20 years to £14m worse.

There were only four non-dominated solutions: Figure 7 shows the objective values of these solutions in detail, and Figure 8 gives the variable values. The solution which performed best financially did not use CCHP; all others did. The thermal store was sized to the maximum permissible value in all but one case (solution 1). CHP 1 was sized at 809 kW$_{elec}$ in all but one case (solution 2 was sized at 609 kW$_{elec}$). CHP 2 was more variable in size: the total capacity increased gradually as the solutions progressed from low to high carbon savings.

Figure 9 gives the results of the sensitivity analy-

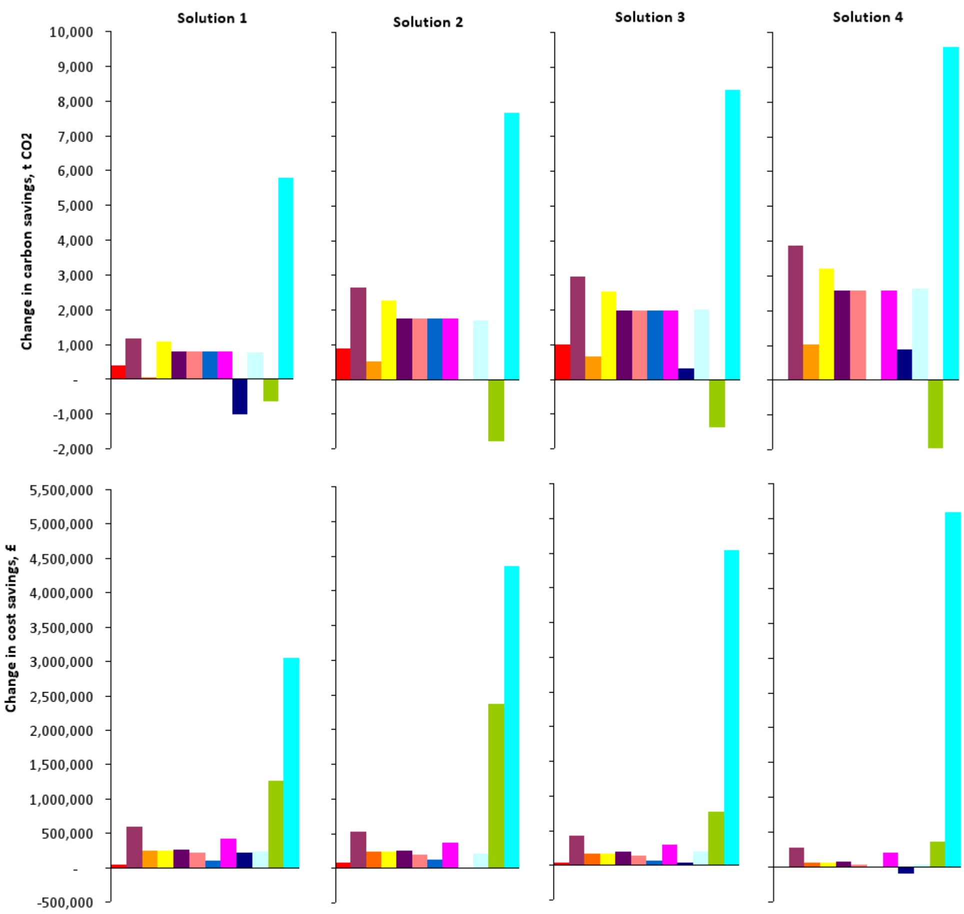

Figure 10: Variation under each scenario for study 1. Bars give the distance between the solution in question and the nearest optimal solution (negative values being an improvement). See Figure 9 for key.

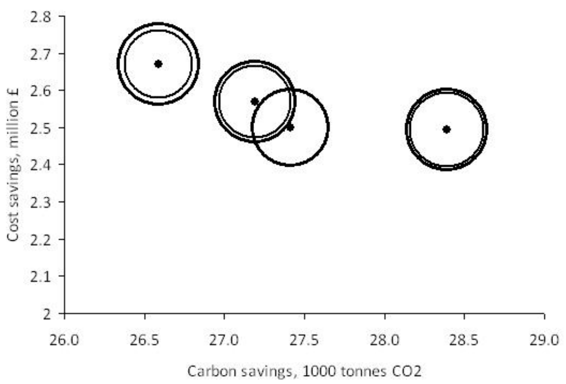

Figure 7: Non-dominated solutions. The radii of the two circles represent the CHP capacities.

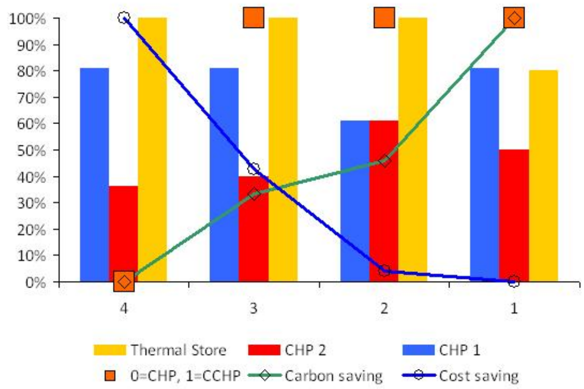

Figure 8: Variable and objective values for the four non-dominated solutions.

sis. This indicates which scenarios are beneficial, detrimental or neutral to each of the objectives. For example the greatest detrimental effect to the environmental objective came from a low carbon grid factor, whereas the most detrimental to the financial objective was a shorter project span.

Figure 10 gives the results of the broader risk analysis; solution numbers correspond to those in Figure 8. As was to be expected, the main solutions were almost always out-performed by the nearest optimal point (shown by positive values). Sometimes an improvement in one objective was balanced by poor performance in another (for example for the low demand scenario all main points were better environmentally but worse financially). For the low electricity price index scenario there was zero change for solutions 2 and 4, indicating that the main solution was a member of the optimal Pareto set.

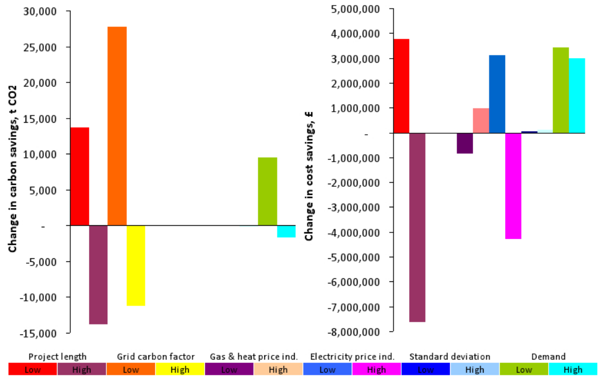

Figure 9: Average change in objective value under each scenario (negative values indicate an improvement over the main solution).

This information provides a valuable complement to the objective and variable values given above. For example solution 4 has the best financial performance, so may be chosen if the carbon savings are deemed to be sufficient. However, it has the greatest degree of risk regarding the carbon objective: 7 out of 12 of the scenarios would reduce the carbon savings by over 10%. Solution 1 provides a much lower level of risk for a relatively small financial penalty. Alternatively, if other aspects of a development indicate that high demand is unlikey (for example better fabric specification) then this source of risk may be discounted, making solution 4 onace more a plausible choice.

Study 2: Multi-phase development

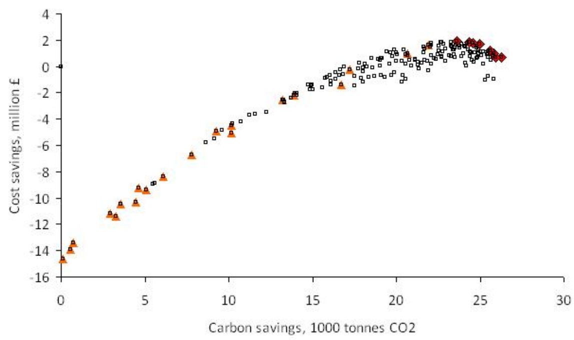

Figure 11: All solutions evaluated. Non-dominated solutions are shown in red. Initial random solutions are shown in yellow.

The second case study sought to optimise the same development, but taking place in five two-year phases: all phases consisted of 6,000m^2 residential and 4,000m^2 retail, with phase one having an additional 10,000m^2 hotel.

Figure 11 shows all solutions evaluated by the algorithm. There is a similarly large degree of variability. There were eight non-dominated solutions; Figure 12 shows the objective values of these solutions in detail, and Figure 13 gives the variable values. The thermal store was sized to the maximum permissible

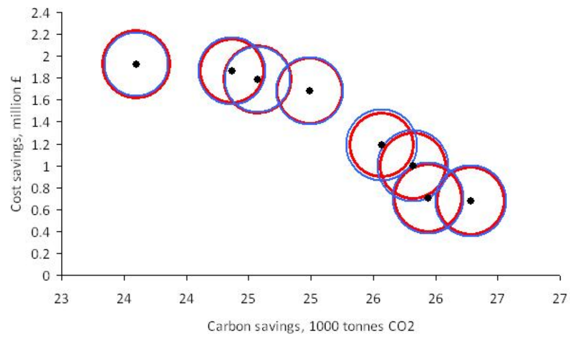

Figure 12: Non-dominated solutions. The radii of the two circles represent the CHP capacities.

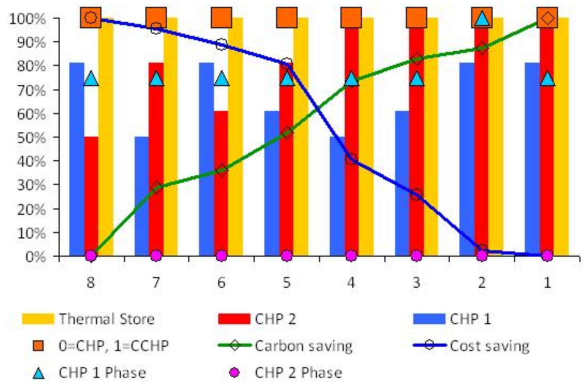

Figure 13: Variable and objective values for the eight non-dominated solutions.

value in all cases. The construction date of the first CHP unit was always the start of the project, and the second almost always six years subsequently, in phase four of five (the exception being solution 2, in which it was constructed in phase five). In most solutions (the exceptions being 6 and 8) the larger CHP unit was constructed first. All solutions used CCHP.

The sensitivity results for study 2 were very similar to study 1 (Figure 9). The main difference was a small improvement in the environmental objective for many scenarios. However, this is due to a wider spread of possible changes: the scenario may cause the solution to improve or to decline.

Figure 14 gives the results of the broader risk analysis; solution numbers correspond to those in Figure 13. Due to the large number of solutions, colour codes have been used to indicate changes. Again there is high variability between the solutions: solution 1 remains largely unchanged, whereas solutions 5 and 6 have many beneficial changes balanced by many detrimental ones.

CONCLUSIONS

This work has drawn together three key areas relating to CHP and CCHP decision-making: environmental benefits, financial performance, and the risks associated with model uncertainties. This has been achieved

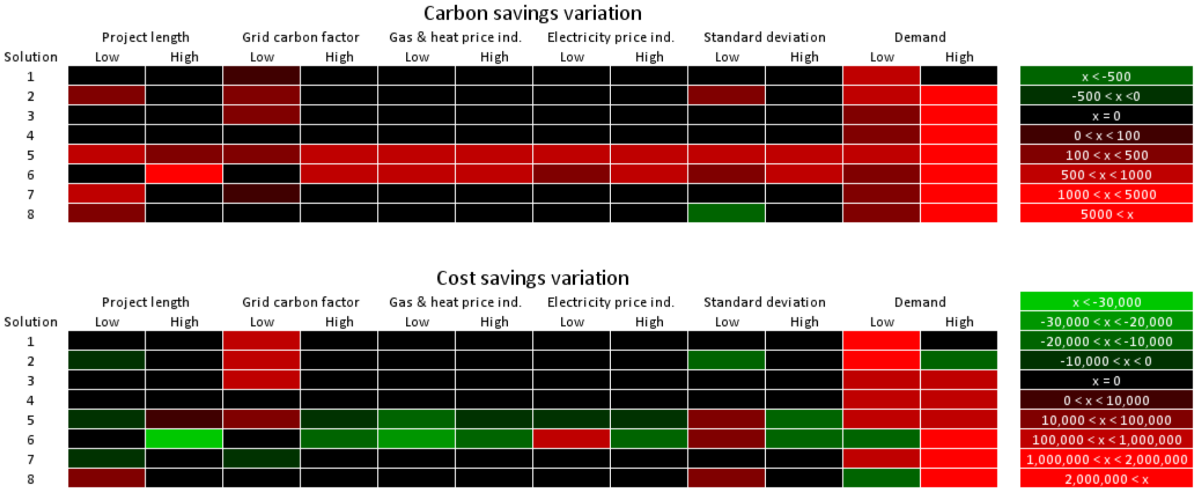

Figure 14: Variation under each scenario for study 2. Colours indicate the distance between the solution in question and the nearest optimal solution (negative values shown in green being an improvement).

through use of a holistic model, which combined plant control on an hourly basis, annual analysis based on weather-dependent loads, and financial evaluation over the project lifespan. This allowed optimisation of phased developments, adjusting plant capacities and construction dates.

The process has been applied to two case studies. Results were analysed visually to highlight variable trends amongst the optimal solutions (part of the process known as innovisation, coined by Deb and Srinivasan (2006)). The optimisation process was also conducted for a range of scenarios involving changes in the model parameters, and the affect on the optimal solutions analysed (Deb terms this higher-level innovisation).

The results show the large range of financial and economic performance of a CHP system: it is important to make correct decisions regarding sizing, as poorly-sized systems perform very badly. An interesting feature of the optimal solutions for both case studies was the asymmetric sizing of the CHP units: two smaller units appear to be always preferable to one large one, and it is very often good to have two units of different sizes. The analysis of many different scenarios has highlighted differences in resilience to changing external circumstances: some solutions remain near-optimal, whereas others are badly affected.

Future work in this area could investigate the links to other aspects of the development (balance of use-types, climate). A more holistic optimisation could examine CHP (and CCHP) as one of a number of supply technologies (biomass, fuel cells), or in combination with measures to lower demand.

ACKNOWLEDGEMENTS

Funding has been provided by EPSRC and Buro Happold Ltd. The lead author is a Research Engineer with the Industrial Doctorate Centre in Systems, Universities of Bristol and Bath, UK.

REFERENCES

Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *Evolutionary Computation, IEEE Transactions on*, 6(2):182–197.

Deb, K. and Srinivasan, A. 2006. Innovization: innovating design principles through optimization. In *Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 1629–1636, Seattle, Washington, USA. ACM.

Hawkes, A. and Leach, M. 2005. Impacts of temporal precision in optimisation modelling of micro-Combined heat and power. *Energy*, 30(10):1759–1779.

Li, H., Nalim, R., and Haldi, P. 2006. Thermal-economic optimization of a distributed multi-generation energy systemA case study of beijing. *Applied Thermal Engineering*, 26.

Ooka, R. and Komamura, K. 2008. Optimal design method for building energy systems using genetic algorithms. *Building and Environment*.

Song, Y. H., Chou, C. S., and Stonham, T. J. 1999. Combined heat and power economic dispatch by improved ant colony search algorithm. *Electric Power Systems Research*, 52(2):115–121.

Tanaka, Y., Umeda, Y., Hiroyasu, T., and Miki, M. 2007. Optimal design of combined heat and power system using a genetic algorithm. *Proceedings of International Symposium on EcoTopia Science 2007 (ISETS07)*.

Vasebi, A., Fesanghary, M., and Bathaee, S. 2007. Combined heat and power economic dispatch by harmony search algorithm. *International Journal of Electrical Power & Energy Systems*, 29(10):713–719.