Reducing energy use and operational cost of air conditioning systems with multi-objective evolutionary algorithms

Cristian Perfumo, John K. Ward and Julio H. Braslavsky

***Abstract*— Air conditioning is responsible for around 60% of energy use in commercial buildings and is rapidly increasing in the residential sector. Although each system is individually small, the proliferation of air conditioning and the correlation of energy use with temperature is driving peak demand and the need for electricity distribution network upgrades. Energy retailers are now looking for ways to reduce this aggregate peak demand, leading to a tradeoff between peak demand, energy cost and the thermal comfort of building occupants. This paper presents a multi-objective evolutionary algorithm (MOEA) to quantify trade-offs amongst these three competing goals. We study a scenario with 8 air conditioners (ACs) and compare our findings against the case of having all ACs working independently, irrespective of global goals. The results show that, with statistically significant certainty, any run of the MOEA outperforms any scenario where the ACs function independently to keep a given level of comfort on a typical hot day.**

I. Introduction

Buildings are responsible for over a third of global energy related greenhouse gas emissions [1]. They have been identified by the Intergovernmental Panel on Climate Change (IPCC) as the largest and most cost effective sector for reducing greenhouse gas emissions with potential reductions equivalent to $3.2\mathrm{GtCO}_2$ by 2020 [2].

The most common means to manage consumer energy usage is through time-of-use (TOU) energy contracts (where electricity is charged at different rates depending on the time it is used) and through peak demand charges (where a component of the electricity bill is related to the maximum instantaneous energy usage (kW) during the billing cycle). Although successful in broadly shaping energy usage, these tariff structures do not provide any mechanism to deal with the aggregate peak demand that results from correlation between multiple customer loads.

To target these correlated loads, a number of Australian electricity retailers (e.g. ETSA, Country Energy and Integral Energy) have been trialling critical peak pricing and direct load control schemes where user loads (primarily air conditioning) can be controlled directly in response to periods of high aggregate power demand. These schemes have focused solely on reducing energy demand during peak times.

Cristian Perfumo (cristian.perfumo@csiro.au) is with the Commonwealth Scientific and Industrial Research Organisation (CSIRO) and the University of Newcastle.

Julio H. Braslavsky (julio.braslavsky@newcastle.edu.au) is with the ARC Centre of Excellence for Complex Dynamic System and Control, School of Electrical Engineering and Computer Science The University of Newcastle.

John K. Ward (john.k.ward@csiro.au) is with the Commonwealth Scientific and Industrial Research Organisation (CSIRO).

In this paper, we explore the application of MOEAs to optimise the operation of residential air conditioning (AC) systems and so realise the optimal tradeoff between peak demand reductions, energy cost, and occupants thermal comfort.

Although the primary goal of AC systems is to provide comfort to building occupants, it is necessary to balance its usage with other important factors such as greenhouse gas (GHG) emissions and running energy costs. Given that occupant comfort, GHG emissions, electricity network interactions and energy cost are neither completely orthogonal nor directly proportional to each other, it is generally impossible to simultaneously meet these competing goals optimally. Instead, they must be met in appropriately chosen tradeoff solutions. Such a scenario is a typical candidate for the application of multi-objective optimisation techniques.

Evolutionary algorithms (EA) [3] are widely applied to the optimisation of heating, ventilation and AC (HVAC) systems. Often they are used to find an optimal design for the topology of the HVAC system, the building shell, or some components of the HVAC system [4], [5], [6], [7]. EAs have also been applied to CO_2 concentration control [8].

In terms of the application of EAs to controlling HVAC systems, The authors of [9] use the non-dominated sorting genetic algorithm II (NSGA-II) to optimise thermal comfort (minimise the percentage of people dissatisfied) and energy consumption. They apply it to a model of a commercial building with a centralised control of about 70 zones. Their variables are the setpoints of each zone, as well as supply duct static pressure, supply air temperature and chilled water supply temperature. Every 15 minutes, the EA optimises the 73 variables for the next 15-minute period. The chief differences between [9] and the present work are *(a)* that we look at independent (as opposed to centralised) ACs; *(b)* we study a third objective: the demand during peak hours; and *(c)* our decision variables are the *ON/OFF* state of each AC. Also [10], [11] use EA for the same purposes, again, addressing centralised HVAC systems. In the case of [11], EAs are applied to develop fuzzy logic controllers for the HVAC system, taking into account energy performance and comfort.

Our research aims to apply EAs to study the interactions between different decentralised strategies for the AC systems, soft and hard constraints on the variables, and the resulting compromise solutions obtained. In the present work we explore the potential benefits of coordinating many residential-like ACs (e.g. a city block or different apartments in a building) instead of having them all work unaware of the

978-1-4244-8126-2/10/$26.00 ©2010 IEEE

existence of one another, as it is the case in reality at the moment. This initial study analyses a small set of identical ACs as a first step towards deriving general conclusions about larger problems.

The rest of the paper is organised as follows: Section II describes our building and AC model. The optimisation problem is defined in Section III. The technique (MOEAs) chosen to solve it and the details of the particular algorithm developed in this work are shown in Section IV. The results obtained are discussed in Section V. Finally, Section VI presents the conclusions as well as challenges for the future.

II. Description of the model

In this paper we consider the interaction of eight independent AC systems. These are considered to be operating under a basic on/off thermostat control, which is a situation representative of many of the existing installed residential AC systems.

An on/off thermostat works by monitoring room temperature, turning the AC compressor on when the temperature rises above a high threshold (i.e. $23C^o$) and off again once the temperature drops below a lower threshold (i.e. $21C^o$). The cycling of the compressor is often not noticed by the user since the air circulation fan usually operates continuously.

The term *on-off* refers to the operation mode of the ACs: each of them can be either working at its full cooling capacity (*on*) or having no cooling activity at all (*off*) with no intermediate state possible. Many modern systems use an inverter to allow the AC compressor to run at part load rather than just on/off. Although the approach used in this paper can be extended to include these devices, they are not considered for brevity, and we focus on the more common on/off thermostat control.

For each of the rooms modelled, the temperature of the air inside the room is given by the difference equation

$$\begin{aligned} T_{int}(k+1) = \quad & T_{int}(k) + P\tfrac{UA}{C}(T_{out}(k) - T_{int}(k)) \\ & +Q_{int}(k)\tfrac{P}{C} - Q_{ac}(k)\tfrac{P}{C} \end{aligned} \tag{1}$$

where

- $T_{int}(k)$ is the temperature inside the room at time interval k (C^o),
- $T_{out}(k)$ is the outside temperature at time interval k (C^o),
- P is the duration of an interval of time ($seconds$),
- U is the heat transfer coefficient of the walls, floor and ceiling of the room (W/m^2C^o),
- A is the total area of the of the walls, floor and ceiling of the room (m^2),
- C is the heat capacity of the room (J/C^o),
- $Q_{int}(k)$ is the heat transmission from internal load at time interval k (W),
- $Q_{ac}(k)$ is the heat transmission from the AC at time interval k (W).

The building parameters used for this work are listed in Table I.

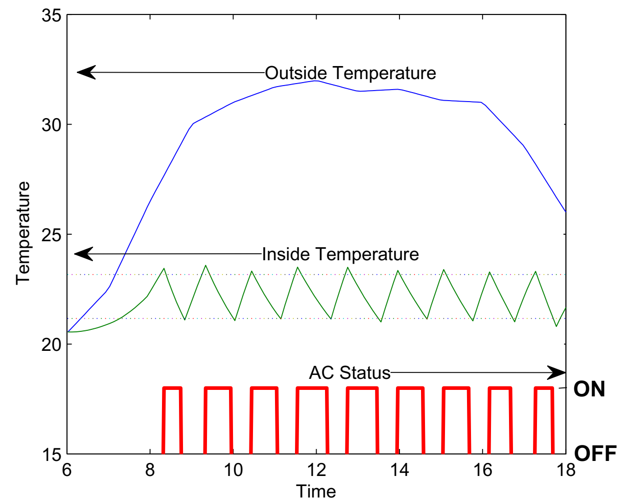

Fig. 1. External and internal temperature and AC activity for one room

TABLE I
Building parameters of our model in this work

Parameter	Value	Parameter	Value
P	360 $seconds$	U	1 W/m^2C^o
A	216 m^2	C	1600152 J/C^o
Q_{int}	525 W	Q_{ac}	1300 W
Setpoint	22 C^o	Hysteresis	2 C^o

For the sake of simplicity, in this study we assume that all of the rooms are identical, i.e. they have the same thermal load and mass and their ACs have the same cooling power. The regimes of operation of the ACs, however, will in general differ, depending on initial conditions and disturbances. We represent the ACs on and off periods with a binary array with a length equal to the number of time intervals modelled. Each position in the array indicates whether or not the AC is on in the corresponding time interval.

Figure 1 illustrates a typical sequence of on-off periods of the AC system corresponding to one room (lowest square waveform) to maintain the temperature of the room (middle saw-tooth waveform) within comfort limits in the face of the evolution of the external ambient temperature (top curve) on a typical hot summer day between 8 a.m. and 6 p.m. We can see how the room temperature decreases every time the AC is switched on, and increases when it is switched off, while the external ambient temperature is higher than the maximum set-point level of comfort. We also see that the AC on-periods are wider when the external temperature is higher in the day.

Figure 2 represents a possible scenario with eight AC systems operating in a completely decentralised and uncoordinated way. The figure shows on-off periods and temperature evolution plots, as in Figure 1, for each of the eight rooms. The plot in the lowest right hand side corner represents the evolution of the aggregate energy consumption

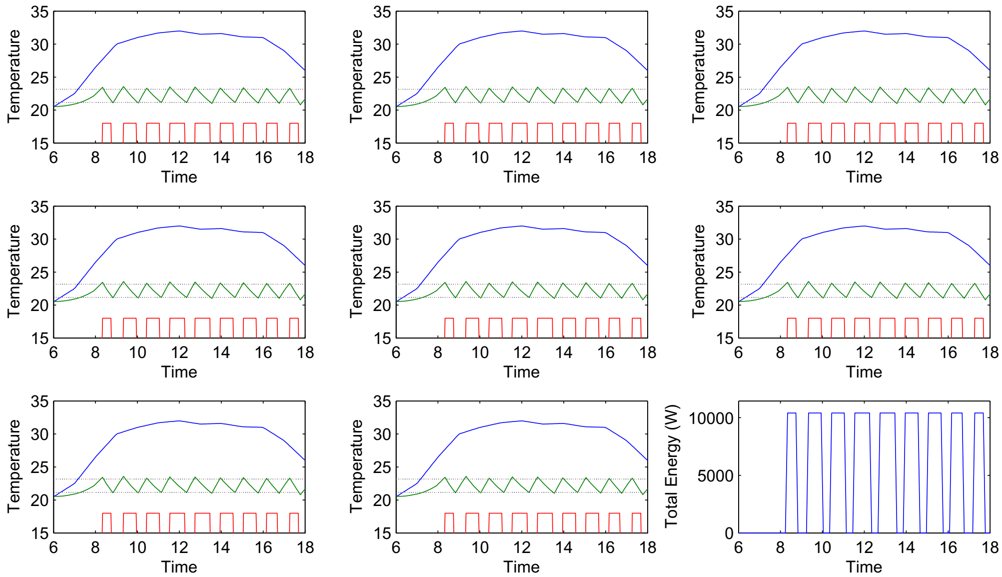

Fig. 2. One possible scenario, where all the ACs work in phase

of the eight AC systems. In this plot we see a potential worst-case scenario, in which the eight AC systems tend to be synchronised (in this case they are exactly in phase because of identical initial and functioning conditions of the rooms over time), yielding very undesirable large amplitude oscillations in the aggregate energy consumption.

Although the exact parameters and dynamics will vary between different systems, the AC model and parameters used in this paper, with dynamics as shown in Figure 2 are representative of the type of energy use and thermal performance characteristic typically seen in residential air conditioning systems. Regarding demand during peak hours, a full analysis is presented in Section V.

III. The optimisation problem

We are interested in finding sequences of on-off periods per room that minimise the following three competing objectives:

1) *Discomfort*: the occupants of the rooms should be as comfortable as possible.
2) *Energy cost*: the electricity bill should be as cheap as possible.
3) *Demand during on-peak hours*: the highest instantaneous energy consumption during the period of the day with highest electricity rates should be as low as possible.

It only makes sense to use setpoint and hysteresis as a comfort metric when the temperature profile has a uniform distribution like the one in Figure 2. When the profile can take any irregular shape, those parameters no longer can be mapped as comfort since the relationship between comfort and temperature is not linear. We use the widely-accepted Percentage of People Dissatisfied (PPD) as a discomfort metric. PPD is defined by the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) [12] and it estimates the percentage of people that would vote that they are uncomfortably cold or hot if they were surveyed.

TABLE II
Residential electricity rates in Australia as of 1 July 2009

Rate	Hours	Price (cents/kWh)
Off-peak	10 p.m. to 7 a.m. every day	7.4
Shoulder	7 a.m. to 2 p.m. and 8 p.m. to 10 p.m. weekdays - 7 a.m. to 10 p.m. on Sat, Sun and public holidays	12.8
Peak	2 p.m. to 8 p.m. on working weekdays	32.4

The PPD is calculated using the Predicted Median Vote (PMV), which in turn is a function of temperature, clothing, activity level of the occupants and air velocity [12]. Because the PPD is an instantaneous measurement, the arithmetic mean of the PPD in all the rooms at all time intervals is used as the objective to minimise.

The energy cost in dollars is calculated based on the residential rates stipulated by Energy Australia in New South Wales, Australia. Under this scheme, there are three different tariff time bands: *off peak*, *shoulder* and *peak*, each of them with a different price. Table II summarises these costs.

The maximum demand (in W) during the peak hours does not directly affect the end-user bill. However, the energy retailers do care about this instantaneous demand (which is the reason for higher prices during peak hours), since they have to plan their networks to cope with the highest load, no matter how infrequent it may be. In fact, the load control schemes mentioned in Section I are driven primarily by the demand during peak hours caused by the ACs.

The mathematical formulation of the objectives to minimise: discomfort, energy cost, and demand during peak hours is given by

$$Df = \frac{1}{R}\sum_{r=1}^{R}\left(\frac{1}{|I|}\sum_{\forall i \in I}\mathrm{PPD}_{ri}\right), \tag{2}$$

$$M = \sum_{r=1}^{R}\left(\sum_{\forall i \in I_{\text{off}}} W_{ri} * M_{\text{off}} + \sum_{\forall j \in I_{\text{sho}}} W_{rj} * M_{\text{sho}} + \sum_{\forall k \in I_{\text{peak}}} W_{r} * M_{\text{peak}}\right), \tag{3}$$

$$D = \max_{i}(W_{\text{AGG}}), \quad \text{where } i \in I_{\text{peak}}, \tag{4}$$

where

- Df, M and D are discomfort, energy cost and maximum demand respectively (%, \$ and W),
- PPD_{ri} is the Percentage of People Dissatisfied for the room r at the time interval i (%),
- I_{off}, I_{sho} and I_{peak} are the intervals of time falling into the off-peak, shoulder and peak hours respectively, $I = I_{\text{off}} \cup I_{\text{sho}} \cup I_{\text{peak}}$, $|I|$ is the cardinality of I,
- M_{off}, M_{sho} and M_{peak} are the costs of one interval of time at the off-peak, shoulder and peak hours respectively. (these prices are taken from Table II (\$))
- R is the total number of rooms, or ACs,
- W_{ri} is the power used by the AC for room r during the time interval i (W),
- W_{AGG_i} is the power used by all the ACs (aggregate) at the time interval i (W).

Because more than one conflicting objective is taken into account, there is no single optimal solution but a set of trade-off solutions, each of them giving different degrees of importance to comfort, cost and demand. Also, the search space grows exponentially as the number of rooms or time intervals increases, making the problem suitable for population-based techniches such as EAs. The way we applied EA to the problem described in this work is explained next.

IV. Our Multi-objective Evolutionary Algorithms

Multi-objective optimisation techniques aim to find a group of solutions that are as close as possible to the optimal set of non-dominated solutions known as *the Pareto front*. A solution X is said to *dominate* another solution Y if, at the same time: *(a)* X is not worse than Y in terms of any objective and *(b)* X is better than Y in terms of at least one objective [13].

In this paper we use the NSGA-II algorithm [14], whose main features are the usage of elitism, an explicit diversity preserving mechanism, with emphasis on non-dominated solutions.

Figure 3 shows the representation of solutions as well as the genetic operators (i.e. crossover and mutation):

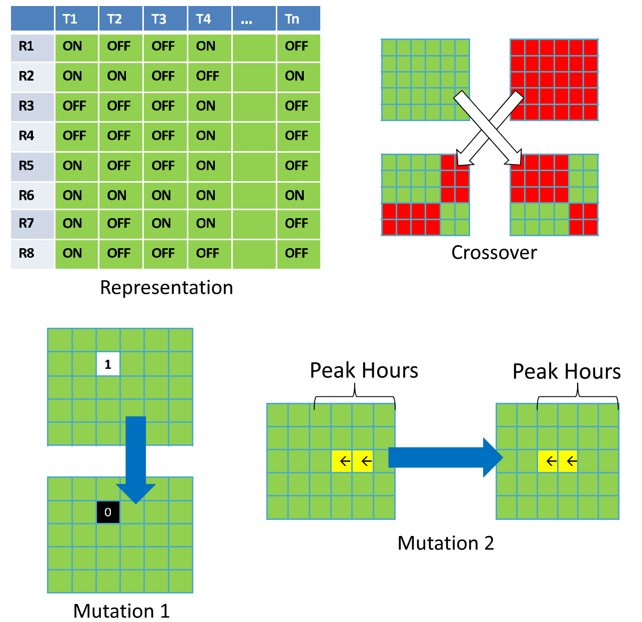

Fig. 3. Representation and genetic operators

- *Representation*: For simplicity, in this first approach each solution is coded as a binary matrix. Rows represent air conditioners and columns, intervals of time. So, if S is a solution, then $S(i,j) = 1$ *iff* the i-th air conditioner is *on* during the j-th period of time.
- *Crossover*: Two solutions (parents) are used to generate two new solutions (offspring) by randomly picking a row cross point and a column cross point. These cross points generate four sub-matrices for each parent. The top-left and bottom-right sub-matrices of one parent are combined with the top-right and bottom-left parts of the other, forming the first offspring. The unused sub-matrices create the second offspring. Either (but not both) cross points can be zero, leading to two sub-matrices per parent instead of four.
- *Mutation 1*: One random bit in the solution is flipped. This is the traditional mutation operator and its objective is to explore new solutions.
- *Mutation 2*: For one randomly-picked AC, a random number of consecutive bits within the peak hours is shifted left or right a random number of time intervals between 1 and 3, making sure to never exceed the boundaries of the peak-hours period. The idea behind this mutation is that by shifting the control bits, less spiky solutions (in terms of the electricity consumption chart) are encouraged. Preliminary results showed that the introduction of this operator improves the quality of the solution with respect to the peak demand objective.

One of the challenges faced during the development of this work was how to represent an evolution-aware, yet correct, demand metric. Semantically, the correct objective to minimise is the maximum demand hit during the period of interest. However, that plain number is not evolution-aware. We explain this concept by the following example: let us

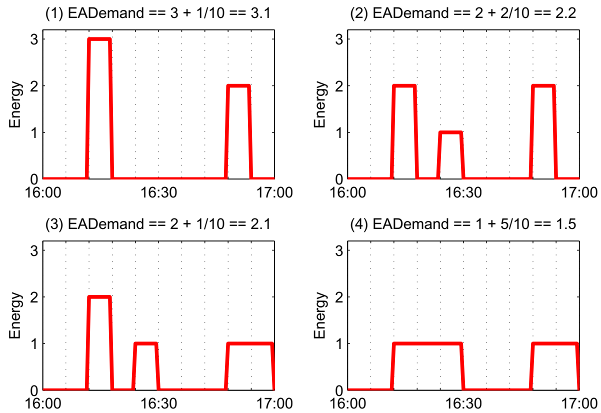

Fig. 4. Evolution-aware demand metric

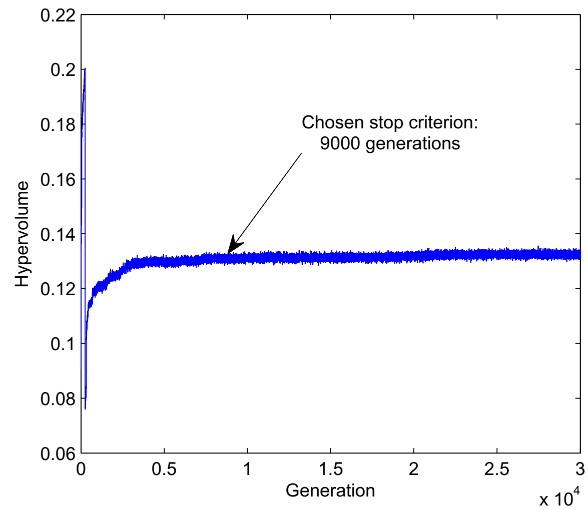

Fig. 5. Evolution of the hypervolume throughout generations. $r = (25, 30, 11700)$, where r_1 corresponds to comfort, r_2 to energy cost and r_3 to maximum demand during peak hours

say that one energy profile shows that the maximum, for instance 10 kW, is hit at four different times over the period of interest. If after applying a genetic operator (i.e. crossover or mutation), the number is reduced to three, the objective value will remain 10kW. It would be desirable that the EA "rewards" this new solution, because it is closer than the previous one to lower the actual maximum demand (once the other three maximums have been lowered).

This challenge was resolved by the simple, yet effective, *evolution-aware demand metric* that is calculated with the equation

$$D_{EA} = D + P\frac{N_D}{I} \tag{5}$$

where

- D_{EA} is the evolution-aware demand metric,
- D Is the actual maximum demand hit during the period of time,
- P is the power of one AC,[1]
- N_D is the number of times D is reached,
- I is the total number of intervals of time into which the period of interest is discretised.

Figure 4 shows the evolution-aware demand metric in action ($P = 1$, $I = 10$). As the first sub-plot (1) changes by shifting the load, the value of D_{EA} decreases. Note that should this scheme not be implemented, the solutions shown in sub-figures (2) and (3) would have the same objective value, i.e. 2.

V. Experimental results

Our MOEA was written in Matlab™ and the experiments run using Matlab 7.8.0.347 (R2009a), on Windows XP. We used commodity hardware since it was enough to cope with the computing power that the executions demanded. Specifically, the computer we ran them on features an Intel® Core 2 Duo™ P8600 CPU running at 2.4 GHz. The capacity of the main memory is 2 GB, with 3 MB of L2 cache.

[1] In the general case, where not all the ACs consume the same power, instead of counting how many times a maximum is hit, the Y-axis can be partitioned into several bands and then count the number of time intervals where the consumption reaches the highest (reached) band.

After evaluating preliminary results, the parameters chosen to run the EA are the ones summarised in Table III.

The number of generations were chosen based on the evolution of (an estimation of) the hypervolume [15]. The hypervolume of a set of solutions for an n-objective problem is the n-dimension volume enclosed by all the solutions in the set and one arbitrarily-chosen reference point r so that everything within the hypervolume is dominated by at least one of the solutions in the set. The hypervolume *measures the size of the portion of objective space that is dominated by those solutions collectively* [16].

Figure 5 shows how the set of trade-off solutions improve as the generations pass. Note that in very early generations the hypervolume is higher, because no solution satisfies the hard constraint ($\mathrm{mean}(\mathrm{PPD}) < 6$). Setting the number of generations to 9000 is found to be a reasonable tradeoff between computation speed and quality of the achieved solution, as the results do not improve significantly with further generations

In order to evaluate the variability of the results, Figure 6 shows the best, median and worst case for 31 independent runs of the MOEA (31 might seem a strange number, but in order to calculate the median, an odd number of samples has to be used). The best and worst case combined are plotted in the lower-right corner sub-figure, showing that the sets of solutions the MOEA is able to find are consistent throughout different runs. These curves were computed using a well-known technique for this purpose, called attainment surface [17].

Because the maximum demand directly translates to a discrete quantity (i.e., the number of ACs switched on simultaneously), each sub-figure in Figure 6 can show the "surface" of solutions in only two dimensions. The third dimension is represented by the shape of the points (diamonds, crosses, circles, etc). As a reference point, the solution in

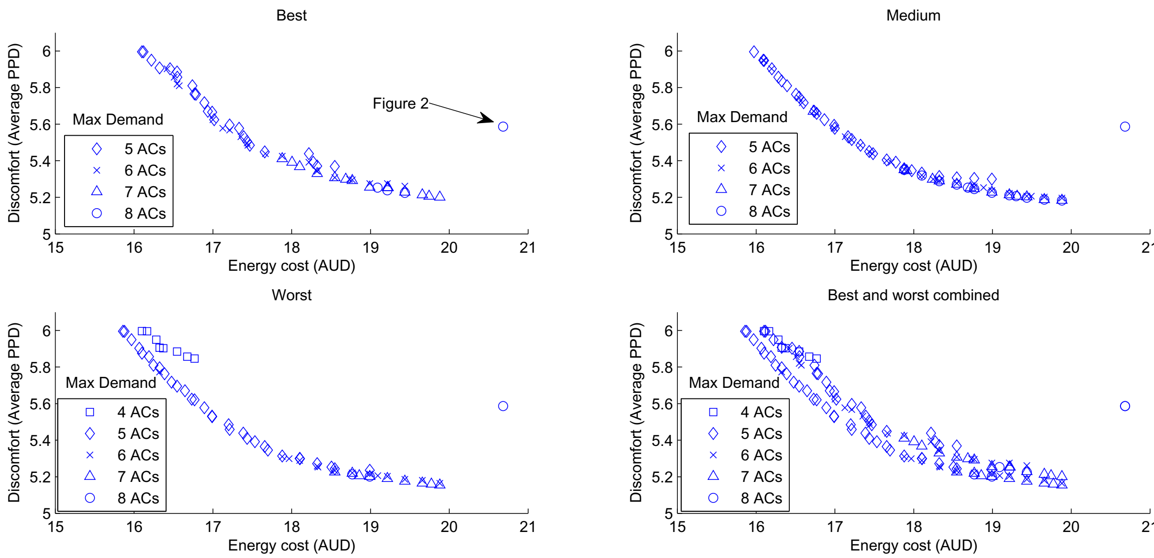

Fig. 6. Attainment surface showing the best, median, wost case and best and worst combined for 31 independent runs

Figure 2 is also shown at approximately (20.7; 5.6).

The lower-right plot combines the best and worst sets of points to make it easier to visualise that there is no considerable variability throughout several executions. This can be seen by looking at the shape of the curve (cost and discomfort) as well as the regions where certain kinds of points (demand) are found: only crosses, diamonds and squares on the left-upper part of the curve and circles and triangles appearing towards the lower-right end (easier to see in the independent plots than in the aggregate one).

Figure 7 shows, with the same representation of Figure 6, the set of non-dominated solutions that our algorithm was able to find for one particular run. It is interesting to look at the part of the curve that falls in the shaded area. All those solutions dominate the one presented in Figure 2. Amongst them, all sets of trade-offs can be found. A good example is the solution with a maximum of 7 ACs (triangle) and the one with 8 ACs (circle) extremely close to each other near the point $(19, 5.2)$. Only by sacrificing an imperceptible amount of money and comfort, the third objective can be reduced by 12.5%. It would be unlikely to spot this particular solution with a traditional, mono-objective technique. This demonstrates, in fact, one of the main advantages of the multi-objective optimisation approach.

On the other hand, the part of the curve outside the shaded area only contains solutions with a maximum of 5 ACs switched on at the same time. This makes perfect sense, as it is the part of the curve that corresponds with the cheaper and less comfortable solutions.

Figure 7 also highlights a trade-off solution selected from the non-dominated front for further investigation. Note that the immediate next solution to the right of the chosen one would have increased the third objective from 5 to 6 ACs. Once again, being able to pick amongst certain trade-offs is one of the chief features of multi-objective optimisation.

TABLE III
EA PARAMETERS

Population	50
Initial pop.	Random
Generations	9000
Crossover probability	0.8
Mutation 1 probability	0.1
Mutation 2 probability	0.1
Hard constraint	mean(PPD) < 6

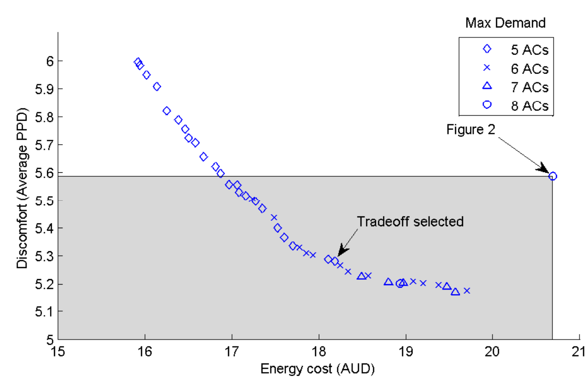

Fig. 7. The set of non-dominated solutions found by the algorithm

The selected solution is depicted in Figure 8 and compared to the solution of Figure 2 (also shown in Figure 7), it exhibits:

- 12.3% cheaper price
- 5.3% less PPD
- 37.5% less demand on peak hours
- 3.3% less energy consumption

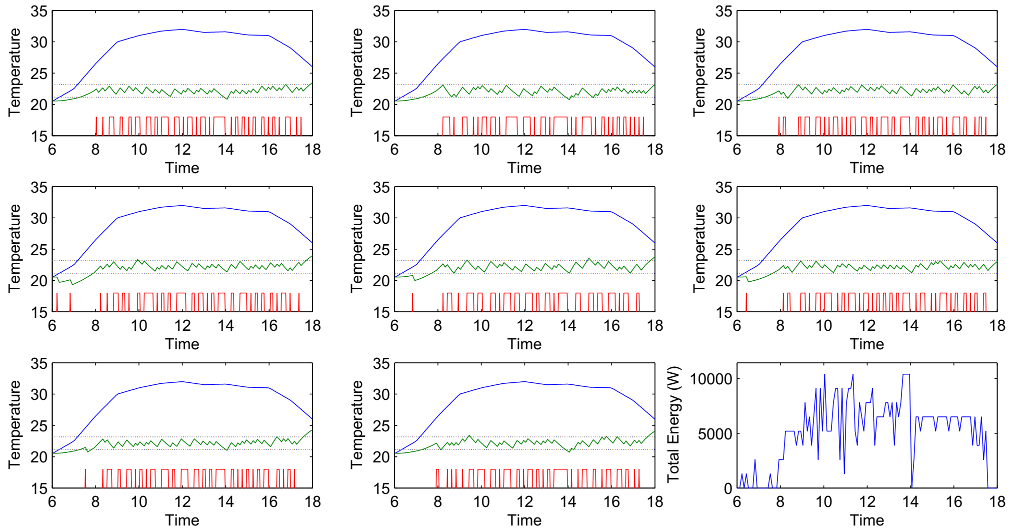

Fig. 8. One (arbitrarily-selected) solution found by the MOEA, representing an interesting trade-off

TABLE IV

COMPARISON OF PEAK DEMAND FOR THE MOEA, WMD, AMD AND BMD

Scenario	Absolute (# ACs *ON*)	MOEA improv.	Rounded up (# ACs *ON*)	MOEA improv.
WMD	8	37.5%	8	37.5%
AMD	7.62	34.3%	8	37.5%
BMD	5.2	3.8%	6	16.7%
MOEA	5	-	5	-

The first three items quantify by how much the selected trade-off dominates the one shown in Figure 2. Note that the fourth item above is not an objective that the MOEA takes into account for the optimisation. However, it is still worth looking at it, since it confirms that the improvements achieved are not at the expense of using more energy.

It is interesting to note that the solution found by the MOEA uses some degree of pre-cooling at two different times during the day. One of them is between 6:00 a.m. and 7:00 a.m., when some AC activity is observed. The reason for this is that the energy is cheaper before 7:00 a.m. (see Table II). The second, even more evident, is before 2:00 p.m. It can be seen in the energy consumption chart that all the ACs are switched on for about 15 minutes before 2:00 p.m.. Accordingly, the individual per-room temperature charts show an the internal temperature drop for all the rooms right before 2:00 p.m.. The reason for starting the "on peak" hours at a low temperature is to reduce the AC usage during this period, which has a two-fold motivation: after 2:00 p.m. the electricity is more expensive and from 2:00 p.m. on, the third objective is considered. In other words, switching an AC on after this time will increase the price and possibly the demand too.

As stated in Section II, it is fair to compare our results against Figure 2 in terms of comfort and energy cost, since it is a realistic situation. However, it is an unfair comparison in terms of the aggregate demand during peak hours, because the ACs are "in phase" (switching on and off simultaneously). On the other hand, the best case scenario for the same amount of energy consumed would be a rectangle comprising the peak hours with the same area as the one below the demand curve in Figure 2. For that particular example the rectangle would have a height of 6760 W. Because the actual heights should be a multiple of 1300, then the highest point would be 7800 W (6 ACs on at the same time). We define the *worst maximum demand* (WMD) and *best maximum demand* (BMD) as the maximum instantaneous energy demand during peak hours for the scenario in Figure 2 and the best-case (rectangular) energy profile respectively.

The solution found with the MOEA dominates the BMD, and therefore it outperforms in terms of the demand objective any other feasible scenarios using the same amount of energy as Figure 2. However, is the BMD the one expected to happen in reality? What about the WMD? The intuitive answer is that the most frequent, or *average maximum demand* (AMD), case should lie somewhere in between. In other words, it would be interesting to quantify by how much the AMD is outperformed with our algorithm.

In this work we estimate the AMD for our 8-room example by randomly generating control vectors that use the same amount of energy as the WMD (and BMD). After generating 5000 random vectors, the AMD during the peak hours is 9906 W (7.61 ACs turned on at the same time). The distribution is shown in Figure 9, showing that the WMD is also the most likely to be encountered.

Table IV shows the WMD, AMD, BMD as well as the maximum demand of the solution found by the MOEA. Because whether or not the numbers should be rounded up (does it make sense to state that 7.61 ACs are on at one given time?) is subject to interpretation, the absolute (not

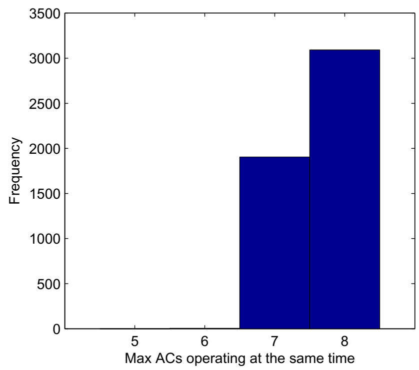

Fig. 9. Histogram of the maximum demand (in number of ACs) during peak hours for random AC binary control vectors

rounded) and rounded results are presented. Each result is accompanied by a percentage indicating by how much our MOEA solution outperforms the scenario.

VI. Conclusions and future work

In this paper we have shown that compared to the common-case scenario, it is indeed possible to save money and reduce demand during peak hours without compromising the comfort of the occupants of the building using the same, or even less, amount of energy.

The proposed EA is able to find several solutions that are strictly better than the scenario where the ACs are considered independently, which is representative of the incumbent operation of residential air conditioning. Moreover, even the results with best maximum demand are outperformed by the worst set of points that the EA is able to find in all the simulation runs.

The optimisation presented in this work assumes that the thermal loads of all rooms are known in advance. In practice it is not trivial to obtain this information, since people are not necessarily in their places all the time and also the weather forecast can be inaccurate. A simple approach for coping with this uncertainty is to regularly rerun the optimisation (as new information is available) instead of only once for the whole day.

Another challenge inherent to any multi-objective problem is how to decide on a convenient trade-off solution amongst the many that the algorithm is able to find. Even though the decision can be totally subjective, certain heuristics can usually be applied in order to automate the decision process. Although we have not addressed this issue here, considerable research has been conducted in this direction, such as [18], [19].

While further research will result in additional modelling and simulation improvements, the results shown in this paper already demonstrate that a MOEA approach to building energy optimisation has the potential to reduce energy usage and mitigate the need for electricity network augmentation without impacting on the comfort of building occupants.

Acknowledgment

The authors would like to thank Sergio Nesmachnow and Adam Berry for their invaluable comments on early drafts of this work.

References

[1] S. de Ia Rue du Can and L. Price, Sectoral trends in global energy use and greenhouse gas emissions, Energy Policy **36**, 1386 (2008).
[2] D. Ürge-Vorsatz and A. Novikova, Potentials and costs of carbon dioxide mitigation in the world's buildings, Energy Policy **36**, 642 (2008).
[3] D. E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning* (Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1989).
[4] J. Wight and Y. Zhang, An "ageing" operator and its use in the highly constrained topological optimization of HVAC system design, in *GECCO '05: Proceedings of the 2005 conference on Genetic and evolutionary computation*, pp. 2075–2082, New York, NY, USA, 2005, ACM.
[5] M. Palonen, A. Hasan, and K. Siren, A genetic algorithm for optimization of building envelope and HVAC system parameters, in *Eleventh International IBPSA Conference*, 2009.
[6] K. T. Ooi and H. Q. Lee, Multi-objective design optimization of a rotary compressor for household air-conditioning, Proceedings of the I MECH E Part E Journal of Process Mechanical Engineering **222**, 241 (2008).
[7] D. Daum and N. Morel, Assessing the saving potential of blind controller via multi-objective optimization, Building Simulation **2**, 175 (2009).
[8] V. Congradac and F. Kulic, HVAC system optimization with co2 concentration control using genetic algorithms, Energy and Buildings **41**, 571 (2009).
[9] N. Nassif, S. Kajl, and R. Sabourin, Evolutionary algorithms for multi-objective optimization in HVAC system control strategy, in *Annual Conference of the North American Fuzzy Information Processing Society - NAFIPS* Vol. 1, pp. 51–56, 2004.
[10] K. Fong, V. Hanby, and T. Chow, System optimization for HVAC energy management using the robust evolutionary algorithm, Applied Thermal Engineering **29**, 2327 (2009).
[11] R. Alcalá, J. M. Benítez, J. Casillas, O. Cordón, and R. Pérez, Fuzzy control of HVAC systems optimized by genetic algorithms, Applied Intelligence **18**, 155 (2003).
[12] R. Parsons, editor, *2005 ASHRAE Handbook: Fundamentals*, Si ed. (American Society of Heating, Refrigerating and Air-Conditioning Engineers, 2005).
[13] K. Deb, Multi-objective optimization, in *Search Methodologies: Introductory Tutorials in Optimization and Decision Support Techniques*, p. Chapter 10, Springer, 2005.
[14] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, A fast elitist multi-objective genetic algorithm: Nsga-ii, IEEE Transactions on Evolutionary Computation **6**, 182 (2000).
[15] E. Zitzler and L. Thiele, Multiobjective evolutionary algorithms: A comparative case study and the strength Pareto approach, IEEE Transactions on Evolutionary Computation , 257 (1999).
[16] L. Bradstreet, R. L. While, and L. Barone, A fast incremental hypervolume algorithm., IEEE Trans. Evolutionary Computation **12**, 714 (2008).
[17] C. M. Fonseca and P. J. Fleming, On the Performance Assessment and Comparison of Stochastic Multiobjective Optimizers, in *Parallel Problem Solving from Nature—PPSN IV*, edited by H.-M. Voigt, W. Ebeling, I. Rechenberg, and H.-P. Schwefel, Lecture Notes in Computer Science, pp. 584–593, Berlin, Germany, 1996, Springer-Verlag.
[18] I. Das, On characterizing the "knee" of the Pareto curve based on normal-boundary intersection, Structural and Multidisciplinary Optimization **18**, 107 (1999).
[19] C. A. Mattson, A. A. Mullur, and A. Messac, Minimal representation of multiobjective design space using a smart Pareto filter, in *AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, 2002.