

2010

Optimal Supervisory HVAC Control: Experiences in Australia

Glenn Platt
CSIRO Energy Technology

John K. Ward
CSIRO Energy Technology

Josh Wall
CSIRO Energy Technology

Follow this and additional works at: <http://docs.lib.purdue.edu/ihpbc>

Platt, Glenn; Ward, John K.; and Wall, Josh, "Optimal Supervisory HVAC Control: Experiences in Australia" (2010). *International High Performance Buildings Conference*. Paper 17.
<http://docs.lib.purdue.edu/ihpbc/17>

This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries. Please contact epubs@purdue.edu for additional information.

Complete proceedings may be acquired in print and on CD-ROM directly from the Ray W. Herrick Laboratories at <https://engineering.purdue.edu/Herrick/Events/orderlit.html>

Optimal Supervisory HVAC Control: Experiences in Australia

Glenn Platt*, John Ward, Joshua Wall

¹CSIRO Energy Technology,
Newcastle, NSW, Australia
E-mail: glenn.platt@csiro.au

* Corresponding Author

ABSTRACT

A number of heating, ventilation and air-conditioning (HVAC) system vendors and operators have started to implement demand management systems in commercial buildings, yet we contend that the impact of such systems on building performance and occupant comfort is poorly understood. This paper examines such issues, showing the results from demand response experiments in two large office buildings. For optimal performance, we believe the HVAC management system should be dynamic and intelligent, responding to changing events, and considering a variety of external factors such as occupancy, human comfort, electricity price, and weather forecast. An example of such a system is detailed.

1. INTRODUCTION

Recently, there has been a significant amount of media attention focussing on what is a perceived electricity supply problem in parts of Australia, and indeed many other parts of the world. In short, with aging network infrastructure, a growing peak electricity demand from loads such as air-conditioners, and ever increasing base-load energy consumption, electricity generation and distribution systems are being seen as unable to cope with the demand placed on them. Whilst the traditional solution to such issues has been to build more supply infrastructure, there is an increasing interest in solving these problems from the demand side. Demand side measures have many advantages- for example, consider the peak-load growth issue- addressing this through supply-side augmentation is an incredibly inefficient approach - whilst peak loads can be over double the average base load on a network, they often occur for only a few days per year. Add in issues such as carbon trading, energy efficiency programmes and increasingly volatile electricity trading markets, and there is a growing interest in energy saving and management opportunities that operate on the demand side of our electricity systems.

This interest in demand side measures extends to the heating, ventilation and air-conditioning (HVAC) industry, where there are a growing number of external pressures on the designers, owners and operators of HVAC equipment:

- Air-conditioning systems are being blamed in popular media as one of the causes of demand-related electricity network failures across parts of Australia
- Programmes are being introduced by electricity retailers that are specifically designed to control or limit the energy consumption of air-conditioning units at certain times of the day¹.
- Recognising issues such as those above, standardisation work has already started trying to standardise the interfaces used to control air-conditioning systems (Standards Australia, 2009).

¹ There are many examples of such programmes, however at this stage they can generally be grouped into two categories- programmes designed to control energy consumption by price signals (for example the ETSA Utilities “Beat the Peak” programme- http://www.etsautilities.com.au/media_release.jsp?xcid=1075) or programmes that use a physical control technology to directly manage air-conditioner consumption, away from the consumer (for example the Energy Australia Dynamic Peak Pricing study- <http://www.energy.com.au/energy/ea.nsf/Content/DPP+About+the+Study>)

Ultimately, HVAC systems participate in demand-side efforts through more optimal load control- changing when loads operate in order to shift consumption peaks around or reduce overall net energy consumption. A number of HVAC system vendors and operators have started to implement demand management systems in commercial buildings, yet we contend that the impact of such systems on building performance and occupant comfort is poorly understood. This paper examines such issues, showing the results of a number of demand-response experiments on large commercial office buildings.

The results of these experiments indicate that for optimal performance, the HVAC management system should be dynamic and intelligent, responsive to changing events, and considerate of a variety of external factors such as occupancy, human comfort, electricity price, and weather forecast. The latter half of this paper details our early efforts in developing such a system.

2. SUPERVISORY HVAC CONTROL

Demand or peak-reduction technologies are not entirely new to HVAC systems. If we consider the peak consumption issue discussed in the previous section, one widely discussed method for managing peak consumption periods with HVAC plant is to charge ice tanks in the building during off-peak times, and use the stored thermal energy of these tanks to cool the building during peak times (MacCracken, 1991). Whilst effective, these techniques rely on relatively expensive additional plant, and are not well-suited to responding to dynamic, rapidly occurring network contingencies.

A number of HVAC demand-side benefits can be realised without requiring new plant or significant infrastructure upgrades- these simply rely on more sophisticated control of the HVAC plant, altering running times and operating setpoints in order to realise a particular load profile. Such “supervisory control” techniques are relatively cheap to implement, and are suitable for retrofitting to existing buildings and HVAC systems. Such systems may be integrated in to a conventional building management system (BMS) suite, or may exist as a separate control system, that gives simple direct setpoints to the BMS.

Current peak management strategies in commercially available HVAC systems tend to be based on “load shedding” strategies, where the BMS is programmed with a prioritised table of loads to shed in the event of certain thresholds being crossed. Other more recent strategies are based on setpoint variations, where (for example) the indoor air temperature setpoints for a building may be varied to try and reduce energy consumption during a particular period.

Unfortunately, in our experience such relatively simple control strategies can result in fairly poor HVAC system performance, often realising unintended outcomes such as extreme temperature excursions in particular zones. These issues are explored in the following sections.

3. TEMPERATURE RESET STRATEGIES

As discussed in the previous section, one of the more recent approaches to dynamic control of HVAC systems is, rather than maintain a constant indoor air temperature, to vary the temperature setpoint in a building in a way that corresponds with external stimuli such as a load-shed request from an electricity utility, or increasing price signal. Our particular interest is varying the indoor temperature setpoint in a way that allows management of a building’s total and peak energy consumption. There are two different strategies commonly employed in such “temperature reset” or “smart thermostat” type approaches to HVAC management:

- **Seasonally raising the temperature setpoint** during summer (cooling mode) and lowering the setpoint during winter (heating mode), to achieve a reduction in overall building energy usage. Energy savings are achieved because (i) the temperature difference between indoor and outdoor conditions is reduced resulting in reduced heat load on the building and (ii) the HVAC system should operate more efficiently with a smaller differential between indoor and outdoor temperatures. The resulting energy savings should then lead to lower monthly energy usage bills and a reduction in greenhouse gas emissions.
- **Temporarily raising the temperature setpoint** during times of summer peak electricity demand to reduce the peak demand of a commercial building. This relies on the thermal inertia of the building to absorb heat for a period of time until the building has warmed up to its new setpoint temperature. In this way cooling

and associated electricity demand is eliminated for a period of time after the setpoint is raised- hence such techniques are often referred to as “demand response” mechanisms. Such techniques may be able to reduce maximum demand charges for a site or potentially attract an incentive payment from an electricity utility as a reward for this demand response.

Whilst the energy benefits of seasonal temperature variations are reasonably well understood, the wider implications of a dynamic and temporary setpoint variation in large buildings is less well known. In order to better examine such issues, we conducted a detailed dynamic setpoint reset experiment in a commercial office building in Melbourne, Australia.

3.1 Setpoint Reset Tests

A reasonably typical Melbourne office block was selected to obtain quantitative data on the cost and performance of using smart thermostat controls to produce a demand response. The building comprised 11 floors with a total floor area of around 10500m². The building contains two chiller units, however prior to commencement of the trials, one of the chiller compressors broke down, reducing the peak cooling available from the chiller system. Consequently chiller capacity was more constrained than might be expected in a typical commercial building. This could be seen as an advantage for these tests, because building operation under constrained chiller capacity conditions is an important feature of hot day power demand profiles.

To facilitate the trial, the BMS system was upgraded to allow a global temperature setpoint reset, providing easy adjustment of all zone temperatures from one position. Extensive instrumentation and data logging facilities were added to the building, logging core parameters such as zone temperatures, HVAC energy consumption, external temperature and so on.

3.2 Setpoint Reset Results

In this section, we investigate how a change in the global temperature setpoint for the building is able to bring about a demand response. A range of temperature setpoint control strategies were attempted:

- (i) Simple temperature setpoint increase
- (ii) Precooling followed by temperature setpoint increase
- (iii) Precooling followed by ramped temperature setpoint increase

Results from each of these control strategies are discussed below

Simple Temperature Setpoint Increase Control Strategy

The most straightforward method to implement a demand response in an HVAC system operating in cooling mode, without completely shutting off the system, is to increase all temperature setpoints. This has two effects:

- There is a substantial initial load reduction as the cooling demand for all zones is reduced while indoor temperatures slowly rise to the new setpoint. The greater the building thermal mass, the longer this will take and consequently the more energy that can be saved.
- Once the system has adjusted to the new setpoint, there will be lower power demand, because (1) the building internal temperature is closer to ambient conditions, which reduces the heat load; and (2) the HVAC system may be more efficient when operating over a lower temperature differential.

Note that once the load reduction event is over, the building energy demand may increase significantly (rebound) as indoor zone temperatures are reduced back to the original setpoint. Figure 1 shows an example of a direct load reduction experiment. In this case, the setpoint temperature was increased from 22.5°C to 24°C for 2.5 hours between 1:45pm and 4:15pm. We note that:

- This was a very hot day, peaking at 37.8°C. With the HVAC compressors running at maximum (constant) power, average zone temperatures prior to the trial were around 23.75°C (1.25°C above setpoint) and still rising. It is evident that there was insufficient cooling capacity to maintain the normal 22.5°C setpoint conditions.
- During the load reduction period, the average zone temperature appears to be converging to a steady state value of 24.5°C (0.5°C above setpoint), with the HVAC system better able to meet the new setpoint. Although the setpoint changed by 1.5°C, the effect on the zone temperatures was only half this and

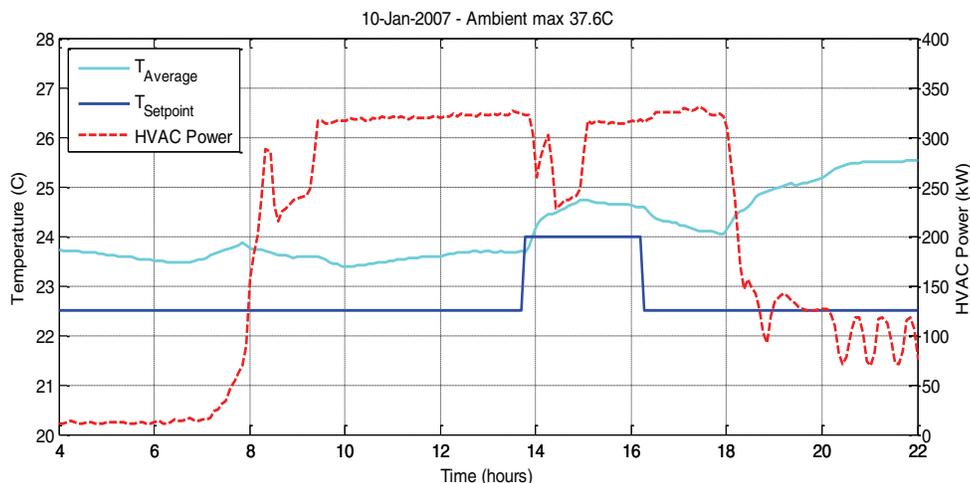


Figure 1: Direct load reduction through increased global temperature setpoint

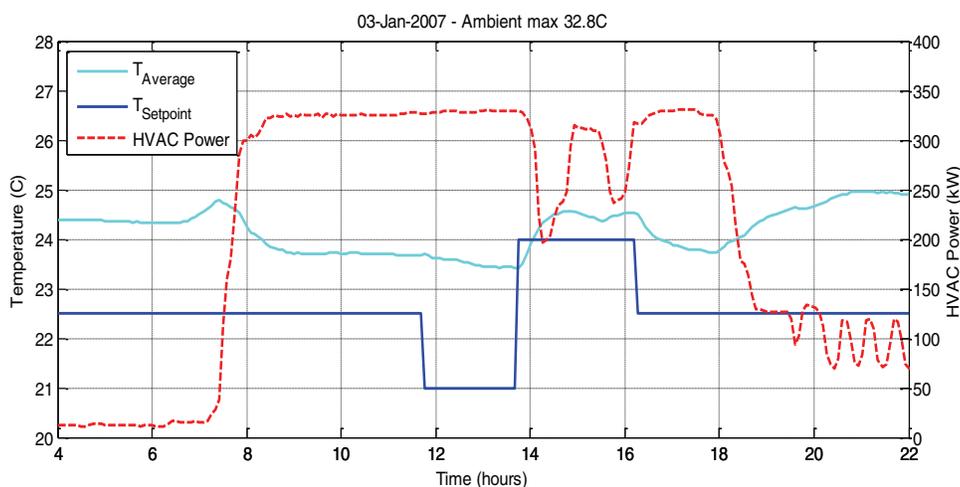


Figure 2: Load reduction with an initial pre-cool prior to increasing global temperature setpoint

consequently energy savings are not as great as might be imagined. The load reduction was only sustained for around one hour. The average zone temperature appears to be converging to a steady state value of 24.5°C (0.5°C above setpoint), with the HVAC system better able to meet the new setpoint.

- Although the setpoint changed by 1.5°C, the effect on the zone temperatures was only half this and consequently energy savings are not as great as might be imagined. The load reduction was only sustained for around one hour.
- In this trial around 65kWh (a 20% reduction) of electricity consumption was avoided over the period of one hour. This was achieved with a 0.75°C increase in average zone temperature.

Precooling Prior to Setpoint Increase Control Strategy

In the event that there is advanced warning that a load reduction is going to be required, a greater energy reduction can be achieved by first pre-cooling the building to build up a “thermal buffer” that can help extend the peak demand period with minimal loss of conditions. Figure 2 shows the results of an experiment examining this strategy. Salient features of this experiment include:

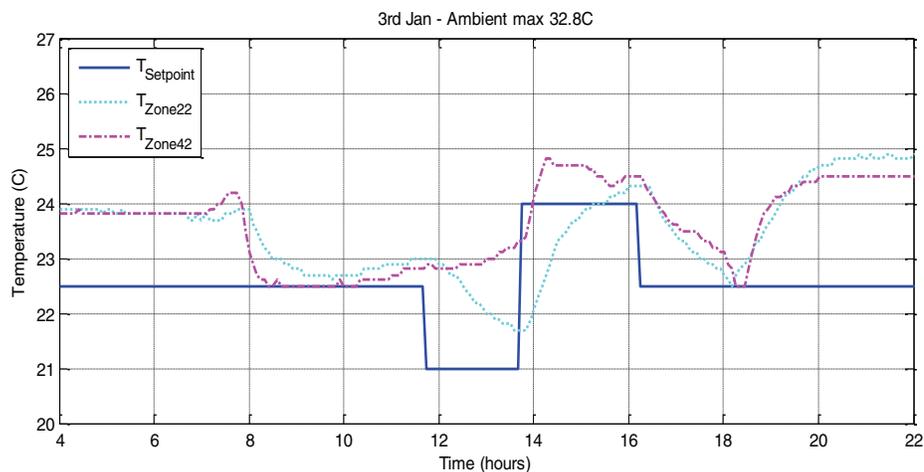


Figure 3: Two individual zone temperatures during the load reduction with pre-cool trial

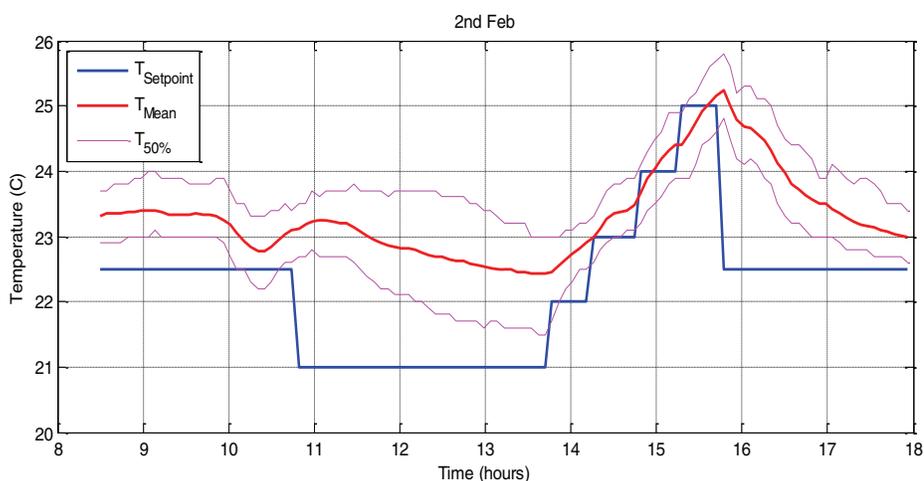


Figure 4: Zone temperature characteristics during a load shed event

- The test was performed on a somewhat cooler day (peaking at 32.8°C) than the test described in the previous section. However with the HVAC compressors running at full power, the average zone temperature prior to the trial was around 23.75°C (1.25°C above setpoint). It appears that cooling capacity was still constrained.
- During the load reduction, the average zone temperature increased to around 24.5°C and approximately 128kWh of energy consumption was avoided. In the first hour of the load reduction, around 77kWh was avoided (a 23% reduction). Once steady state conditions are established (with the compressor cycling on/off to meet demand) there is an average load reduction of around 43kW – this equates to a reduction in energy consumption of 12.7% achieved from a 0.75°C increase in average zone temperature.
- There was a delay of around 30 minutes from initiation of the load reduction to the peak energy reduction occurring. This may limit the immediacy of this strategy for providing a real-time response to sudden price spikes in the electricity market.
- With no additional HVAC capacity available for the pre-cool, any additional cooling supplied to some zones must be achieved through reduced cooling to other zones. That is, the pre-cool was only able to cause a redistribution of cooling between zones. Figure 3 shows this effect. When the building temperature setpoint was reduced, the temperature in zone 22 reduced as desired but the temperature in zone 42 increased.

This issue of HVAC capacity distribution is worth further exploring- figure 3 seems to indicate that a particular load shed event could unevenly affect the various zones in a building. Figure 4 examines this issue, with another pre-cool

and load shed event, this time investigating the spread of zone temperatures in the building- the 50% line shows samples 50% above and below the mean. As can be seen, during the pre-cool, the spread of zone temperatures increases markedly to the point at 12:30pm where the 50% region is $\pm 1^\circ\text{C}$. The reason for this is that when the setpoint is reduced at the start of the pre-cool, all zones are initially well above setpoint so all chilled water valves open to 100%, and cooling is reasonably evenly spread across all the zones. Consequently, the cooler zones (with less load) initially cool fastest and hotter zones can actually increase in temperature. It is not until the cooler zones have further cooled that there is sufficient capacity to start seeing temperature reductions in the hotter zones. This situation where the hotter zones increase in temperature during a pre-cool could be particularly problematic, since occupants of such zones may begin to experience a loss of thermal comfort even before the actual load reduction or demand shed begins.

4. INTELLIGENT DYNAMIC LOAD CONTROL

The results from the previous section show a significant issue around the application of setpoint reset strategies, even with pre-cooling efforts, to large multi-zone commercial buildings- whilst a particular demand reduction may be achieved, such a reduction can result in quite poor distribution of HVAC capacity across zones. This has a particularly significant affect on occupant comfort- in short, occupants could be uncomfortable *before* the actual demand reduction takes place. With these results in mind, we believe that for optimal demand response performance, an HVAC control system must go beyond global setpoint control, to intelligently predicting the effect of any demand response action on all zones of the building, on occupant comfort and energy performance, before initiating such an action. We have been working on such a system, described in the following sections.

4.1 Controller Details

As discussed, for optimal performance an HVAC system controller needs to find a balance between environmental conditions as well as resource usage in energy and financial costs. Finding an optimal balance requires an advanced building controller that operates in a fundamentally different way to a conventional system. Specifically, the controller will need an awareness of different energy sources and their consumption, and it should *plan* the operation of these resources in advance. As an example, many buildings in temperate climates will operate in heating mode (with gas heaters) in the morning, followed by cooling mode (using electrically driven chillers) later in the day. By taking into account anticipated weather and thermal loads later in the day, a *planning* HVAC controller could limit this behaviour, minimising any heating, thereby reducing both heating load and the subsequent cooling load.

Our new control system has been designed to optimise HVAC operation to balance (i) running costs, (ii) greenhouse gas emissions and (iii) occupant thermal comfort. User comfort is predicted via thermal comfort models based on Fanger's (1967) work that includes temperature, humidity and nominal values of other factors (airspeed, clothing and activity levels). The controller is based around two key components:

- An energy model that links HVAC control actions to energy consumption
- A human comfort model that links HVAC control actions to human comfort outcomes

With these two models, the controller runs a continuous optimisation process, planning the way it will run the HVAC plant throughout the next 12 hours. This plan is prepared for every individual zone of the building, and is assembled by determining the minimum HVAC operating parameters required to achieve a given occupant comfort target and maximum energy consumption. This plan is then updated at regular intervals throughout the day, thereby taking in to account changing conditions such as weather or energy price variations, and ensuring energy and comfort constraints continue to be met in a dynamic environment. With such a design, the demand response behaviour described in section 3 need not be explicitly managed- by simply offering a dynamic energy price to the HVAC system, the HVAC controller will re-plan the HVAC system's behaviour to minimise consumption at expensive times. Alternatively, rather than a price signal, the HVAC controller can take a simple demand cap signal, indicating that total consumption should not exceed a certain threshold at particular times of the day. In this instance, again, the controller will simply arrange the HVAC operation plan to respect the maximum demand constraint, whilst still achieving occupant comfort goals.

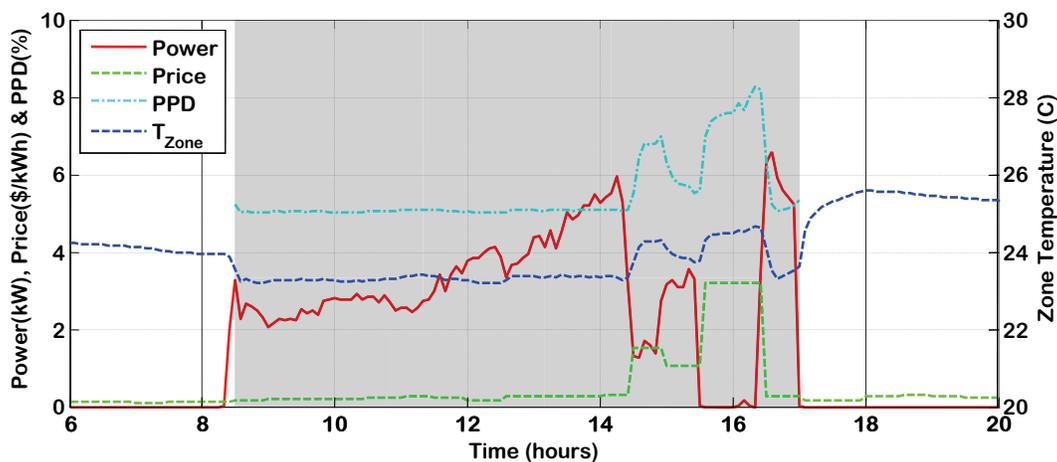


Figure 5: Intelligent HVAC controller results- price spike, zone 1

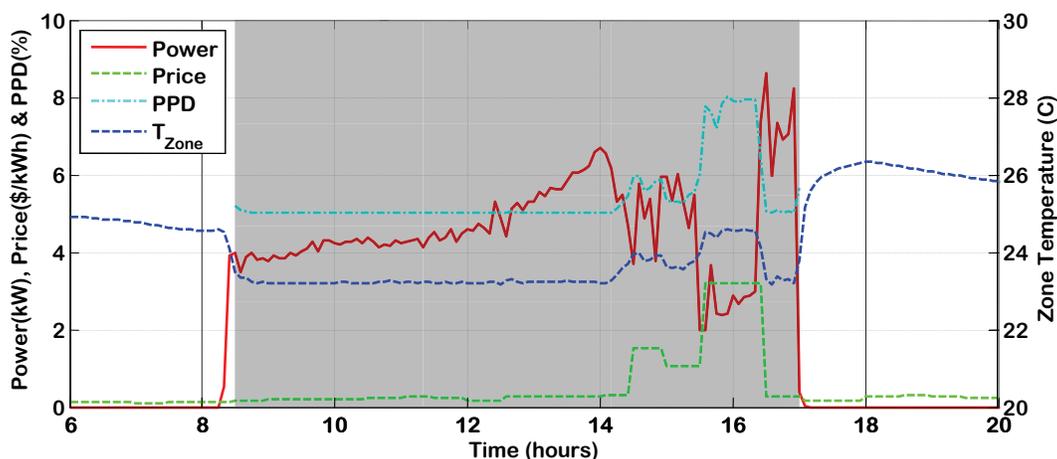


Figure 6: Intelligent HVAC controller results- price spike, zone 2

4.2 Intelligent Dynamic Load Control Tests

Having completed an early implementation of the HVAC control system described in the previous section, we have recently run a number of experiments examining the operation of this system in a real office building.

The system was deployed in a 3 story office building. In these early experiments the system was operated on only one floor of the building, approximate size 1000 m², with 5 separate HVAC zones on this floor. The building operates using a Siemens APOGEE building management system. This BMS was interfaced to our intelligent control system through an OPC interface, where our control system provided basic HVAC setpoints to the BMS.

Extensive instrumentation and data logging facilities were added to the zone under study, logging zone and ambient temperature and humidity, HVAC (chiller and fan) energy consumption and thermal energy flows through the chilled and hot water systems.

4.3 Intelligent Dynamic Load Control Results

The control system described in the previous sections has been operating in our test building for a number of weeks now, across a variety of external weather conditions. As part of our testing, we ran the building controller through a number of different response experiments. A typical zone data set from one of these experiments is shown in Figure 5. This figure shows HVAC power, electricity tariff (\$/kWh based on a real time pricing scenario), estimated

thermal comfort of building occupants (percentage of people dissatisfied, calculated as per the ASHRAE standard (2004)) and measured zone temperature. The operational hours of the building are from 8:30am to 5:00pm, as shown by the shaded region of the figure, outside of which thermal comfort does not need to be maintained.

As can be seen in the results, although there is no explicit startup/shutdown optimisation in the control scheme, the advanced controller has determined an operation schedule that has effectively achieved optimal startup/shutdown by seeking to minimise greenhouse gas emissions and energy costs. Once the system has started, for the first part of the working day, operation appears similar to a conventional HVAC system. However, when the price of energy starts to rise after 2.00pm, the system behaves quite differently to a conventional HVAC plant, as the system initiates a demand response to the changing energy price. At this time, faced with conflicting goals of minimising energy costs whilst maximising occupant comfort, the controller tries to balance these by pro-actively creating a different HVAC operation plan- the controller tries to minimise HVAC energy consumption (and thus operating cost) by allowing the zone temperature to rise during the first price increase at around 3.00pm, and then completely shutting the HVAC system down in this zone for the very severe price peak of over \$3/kWh occurring at around 4.00pm. The effect of these operations on predicted occupant comfort is also shown in the graph- for this zone, the percentage of people dissatisfied rises to approximately 8% during the price spike.

This trade-off and subsequent prioritisation of occupant comfort versus energy expenditure is one that a building's tenant or manager must make. Importantly, a global control strategy for an HVAC system would not allow such a decision to be accurately made- if we recall the results from section 3 of this paper, the global control strategy resulted in one particular zone of the building having quite large temperature excursions during the demand response period, likely resulting in very uneven distribution of occupant satisfaction across the building. Our new HVAC control system would avoid such an outcome, as it adapts a building's demand response on a zone-by-zone basis. This behaviour is shown in figure 6, where we look at another zone of the same building. In the zone shown in figure 6, we can see that during the same price spike period, HVAC energy consumption decreased, but the HVAC system in this zone did not shut down. This is because during this period this particular zone had different thermal properties to the other zone, and the HVAC controller predicted that to completely shut down the HVAC in this zone would have resulted in an excessive percentage of people dissatisfied figure. Thus, whilst the system strived to minimise energy usage and cost during the price spike, it balanced this by keeping the occupant comfort within similar boundaries to the zone shown in figure 5.

To date, results from this advanced HVAC controls trial are yielding energy savings of around 30% compared to the existing building controller. This has been achieved without compromising occupant thermal comfort levels.

5. CONCLUSIONS

There is a growing interest in demand-side participation of major electrical loads, including commercial HVAC systems, where the load will respond to external network requests such as reducing consumption during a particular time. We have shown that the participation of large multi-zone HVAC systems in such exercises is not straightforward- in particular, a global approach may yield undesirable consequences in a particular zone of a building. Addressing these issues, we have implemented an adaptive HVAC control system that plans its behavior for the day ahead, re-planning in response to changing environments or network requests. The controller operates at the zone level, planning a zone's activity based on considerations such as energy consumption, financial expenditure and occupant comfort. The system is designed to be retrofitted to existing buildings, operating via a simple interface to the incumbent BMS.

REFERENCES

- Standards Australia, 2009, Australian Standard AS4755 "Demand Response of Air-Conditioners"
- ASHRAE, 2004, ANSI/ASHRAE Standard 55-2004, Thermal Environmental Conditions for Human Occupancy. American Society of Heating, Refrigerating and Air-Conditioning Engineers. Atlanta, USA. www.ashrae.org
- Fanger, P.O., 1967, Calculation of thermal comfort: introduction of a basic comfort equation. ASHRAE Transactions 73(2):III.4.1.
- MacCracken, C. D., 1991, Off-peak Air Conditioning: A Major Energy Saver, ASHRAE Journal, pp. 12-23.