A constraint-based intelligent controller for lighting systems

C. Ryan, K. N. Brown
Cork Constraint Computation Centre, Department of Computer Science, University College Cork, Ireland

A. Mady, M. Boubekeur, G. Provan
Cork Complex Systems Lab, Department of Computer Science, University College Cork, Ireland

ABSTRACT: A major factor in energy wastage in retail and public-office buildings is the inefficient control of lighting systems. This paper presents a constraint based optimiser for lighting control which minimises energy usage while maintaining occupant satisfaction. This controller is coupled with a reactive controller which we simulate using a hybrid/multi-agent platform. We apply this control strategy to a simulated open office scenario, and outline the simulation results and potential energy savings of the proposed strategy.

1 INTRODUCTION

A major source of energy inefficiency in buildings is suboptimal control. Lighting is one of the main energy consumers in buildings, responsible for 20% of total energy usage in US commercial buildings (EIA, 2003). Up to 58% of this may be wasted due to inefficient lighting control (Delaney, D.T. et al. 2009). The aim of our research is to define an intelligent control methodology which will optimise the light actuation levels to minimise energy wastage.

In this paper, we consider a two step approach, with a higher level controller periodically optimising initial settings, combined with a lower-level reactive controller actuating the light. There is existing work on both centralised optimisation (Singhvi, V. et al. 2005, Mahdavi, A. 2008) and combining decision support with existing control systems (Davidsson, P. et al. 2000, Boman, M. et al. 1999). We formulate the higher-level problem as constraint optimisation. We use hybrid systems models to model and simulate the building environment and reactive controller, to represent both the discrete and continuous components required in lighting control and presence sensing.

We assume a tracking system which monitors the location of individual occupants within the controlled area, and an interface which allows them to express preferences over the lighting. When significant changes occur in the environment, the constraint based optimiser computes initial actuator settings which minimise an objective function combining preferences and energy. The reactive controller then maintains the desired illuminance levels by actuating window blinds and internal lights to compensate for changes in the external light level.

The remainder of the paper is organised as follows: Section 2 introduces the reactive and optimising controllers and the process by which they share control. In Section 3 we specify the scenario and the reactive control strategy. Section 4 details the constraint based optimising controller. Section 5 briefly outlines the modeling methodology used in the simulated scenario and the results of that simulation. We conclude the paper in Section 6.

2 REACTIVE AND OPTIMISED CONTROL

To allow for both optimised and reactive control our approach uses two separate controllers which work together. In this section we describe their operation.

2.1 *Controller Algorithm*

The actuators in the simulated environment are controlled directly by the simulated reactive controller. This controller starts by actuating initial light settings provided by the optimising controller, the 'initial configuration'. The optimising controller provides both the optimal actuator settings and the predicted resulting sensed illuminance levels. It is assumed that any small changes to the environment, such as changes in daylight levels, will result in small changes to the configuration, and thus only small changes to the objective value. If, for instance, daylight increases gradually, the lighting nearest the window merely need to be reduced slightly. Thus by progressivly altering the initial configuration to maintain the predicted levels, the reactive controller alone can maintain an near-optimal configuration when small changes occur.

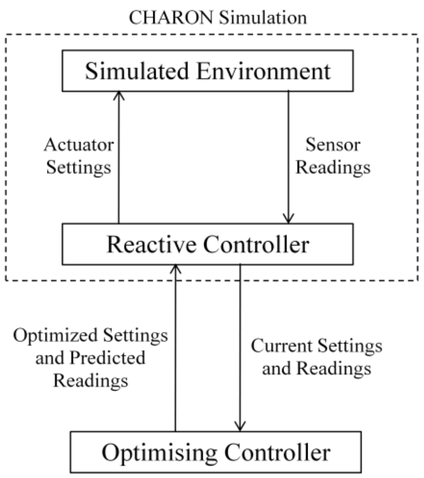

Figure 1. Control System

However, there may be circumstances in which the new optimal configuration is no longer in the neighbourhood of the current configuration, and cannot be found through reactive control alone. In our model two such circumstances exist: the movement of an occupant or significant drift from the initial configuration due to continued changes in the external light. In the case of an occupant moving the distribution of preferences across the environment will change, altering the preferred illuminance levels in a way the reactive controller cannot predict. In the case of configuration drift, we refer to small but cumulative changes to the external lighting of the environment causing the reactive controller’s configuration to become increasingly distant from the initial configuration. When the initial configuration is modified beyond a certain margin, it is assumed that the external lighting may have changed sufficiently to move the optimal configuration out of the neighbourhood of the initial configuration. In both of the above circumstances the reactive controller requests a new configuration from the optimising controller.

The reactive controller includes the current actuator settings, sensor readings and occupant locations in its request for a new configuration. The optimising controller uses these and the stored occupant preferences to compute the optimal configuration according to the objective function. It then returns this configuration, including both the optimal actuator settings and predicted sensed illuminance levels, to the reactive controller. This is the new ‘initial configuration’, which the reactive controller implements and controls around as before.

2.2 *Model Assumptions*

The simulation is an approximation of the scenario environment. We currently make some assumptions about the hardware and the propagation of light within the environment, as follows.

We assume that the propagation of light within the model is uniform and diffuse. The light received at any point from any source depends only on the distance between the point and the source. Similarly, blinding is assumed to uniformly decrease the amount of daylight entering the environment. With regard to the actuators and sensors, they function in an ideal fashion; the actuators emit exactly the amount of light that they are supposed to, and sensors always correctly report the exact illuminance level at their location. Furthermore, it is assumed that these devices are perfectly reliable, that they never suffer from control failures of any kind..

Future work will involve removing these assumptions from the model. Daylight entering the environment and the reduction of daylight via blinding is typically not uniform or diffuse. To account for this, the controller will need a model which can incorporate localised light and dark spots. The possibility of imprecise output from the actuators or inaccurate readings from the sensors adds uncertainty into the selection of the optimal configuration. The optimiser will have to account for the fact that its knowledge of the current configuration may be inaccurate, and that the suggested configuration may not have exactly the predicted effect, and choose the solution most likely to minimise the objective function. Control failure of actuators or sensors could change the options available to the controller, and so it will have to be able to detect such problems to avoid suggesting configurations which cannot actually be implemented.

3 INTELLIGENT LIGHTING SYSTEM SCENARIO

In this section we specify the scenario we are simulating and the reactive controller which controls the actuators.

3.1 *Scenario Specification*

We have adopted a typical architecture as shown in Figure 2. We focus on an open office area, which contains 6 controlled zones, where each zone contains one artificial light and one light sensor. One Radio-Frequency Identification (RFID) receiver is used to cover the whole area; there is one window/binding in the left border of the considered area and a fix number of predefined occupant positions.

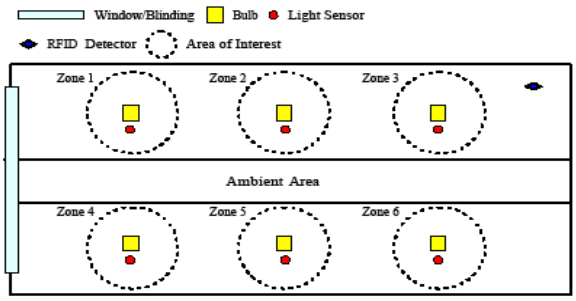

Figure 2. Scenario Model Specification

For the lighting model we integrate blinding and lighting controls. In order to enhance the efficiency of the resulting control model, an optimising controller has been implemented, as explained later.

As a summary, the lighting control scenario behaves as follows:

1 The user can switch on/off the automatic lighting system for several zones, or for all the system through a technician.
2 The occupants provide their light luminance preferences.
3 An occupant is tracked in each zone using RFID, and his preferences are considered wherever he is located.
4 The optimising controller receives the user preferences and sends back the optimal settings.
5 The reactive controller controls the artificial light and the blinding actuators in order to reach the optimal settings and to respond to the influence of changes in the daylight luminance.

3.2 *Reactive Control Strategy*

Figure 3 shows the agents of the control model and its interactions with the environment agents. The controller follows the following scenario in order to control the light intensity:

1 The optimising controller receives the user preferences for each person and its position, sends the optimal light setting and blinding position back to the reactive controller in order to refine the actuation values using a PI-Controller.
2 The reactive controller actuates the artificial light and the blinding position accordingly, then goes to 1 if the occupant preferences have been altered by the movement of an occupant or if a significant deviation from the optimised configuration has occurred.

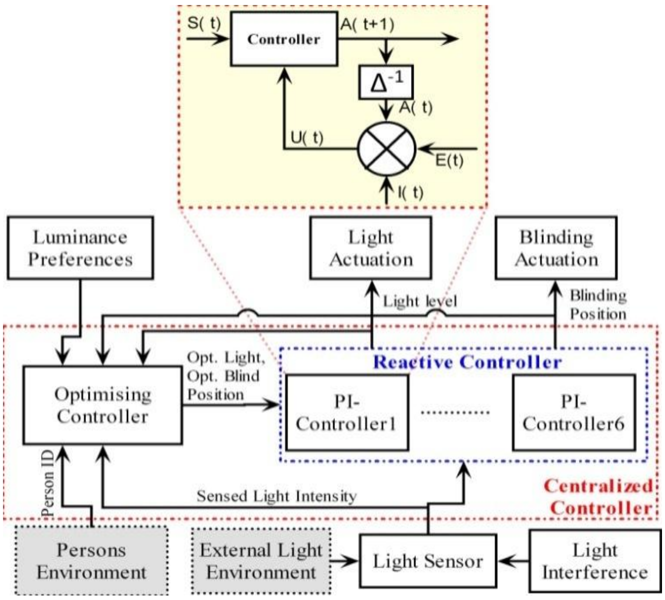

Figure 3. Control Model

The PI-Controller is used to predict the next actuation setting for the lighting dimming level in a close loop fashion (Kolokotsa, D. et al. 2009). The PI-Controller has two main status, first is unstable when the difference between the sensed light intensity and the optimal one is greater than 70 Lux (one light actuation level), and secondly, is stable, if the difference is less than or equal to 70 Lux.

4 CONSTRAINT BASED OPTIMISATION

The optimising controller is responsible for computing on request settings for the actuators which optimise both energy consumption and occupant satisfaction. To this end we model the energy cost and the effect on the building environment of the actuators, and the preferences of the occupants, as a Constraint Optimisation Problem (Dechter, R. 2003). A constraint problem consists of a set of variables, a domain of possible values for each variable, a set of constraints over the assignment of values to variables which restrict the values that may be assigned simultaneously, and an objective function over the assignments. A solution is an assignment of one value to each variable such that no constraint is violated. An optimal solution is one with the highest objective value. Solutions may be obtained by any suitable method, including backtracking and logical reasoning, mathematical programming, or local search.

4.1 *Preference Modeling*

To quantify occupant satisfaction with a configuration we must model their preferences over illuminance levels. We model each occupant's preferences as a preference curve, which associates a satisfaction level from 0 to 1 with each possible lux level. These satisfaction levels may be derived from direct occupant feedback or learned over time based on their actions with regard to the lighting. The preference curves in our simulated scenario peak in the range of 400-500 lux, falling off to 0 satisfaction within 300 lux above or below that range.

Modeling an occupant's preferences over the entire range of illuminance levels allows us to determine an acceptable range of illuminance levels for that occupant. By constraining the occupant's satisfaction to be above a certain threshold, we compute a range of illuminance values within which the occupant is simply considered satisfied, and thus within which the energy cost component of the objective function will be the deciding factor. Since we model the occupant's preferences across all values, this satisfactory range is determined by the occupant's own sensitivity to different lighting conditions.

To apply these preferences to the physical environment, we break it down into control zones, each of which contains a set of zero or more occupants. The zone is the smallest unit of the environment which the controller reasons over, so preference curves are grouped into these zones. By averaging the preference curves of the relevant set of occupants, we produce a zone preference curve, which describes the overall satisfaction of that set of occupants.

Conflicts in occupant preferences may make it impossible to completely satisfy every occupant in a set. Thus the controller first determines the maximum achievable satisfaction for the set of occupants, and the satisfaction threshold is offset from this value.

4.2 *Actuation Modeling*

We evaluate actuation in terms of the energy cost of the actuator configuration and occupant satisfaction with the resulting illuminance levels. The energy cost of each actuator is simply a direct function of that actuator's set point. However occupant satisfaction depends on the lighting conditions effected by the actuators, which is dependent on the combined effects of all the actuators. Thus, as with the preference model, we model the building environment as a set of zones, and model the effects the actuators have on that set of zones.

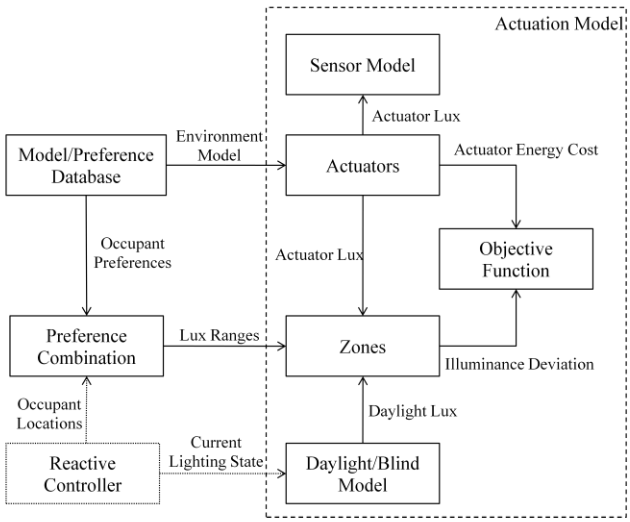

Figure 4. Optimising Controller Overview

We model each actuator as a set point variable and a set of constraints which describe their energy cost and lighting effect. These set point variables are the decision variables of the constraint problem, the variables to which we assign values when searching for a solution. In our simulated scenario there are six lights and one blind. The lights have 11 possible set points ranging from 0 to 10, while the blind has 5 set points, from 0 to 4. In this model both the energy cost and light output are proportional to the set point.

For each actuator a set of constraints relate the output of that actuator to the lux of each zone which it affects. The illuminance received by a zone from each actuator is proportional to the base output of that actuator and is modeled as a relational constraint. The total lux in each zone is the sum of the individual contributions. These predicted lux values are the dependant variables of the model, and are evaluated against the occupants preferences as part of the objective function.

The objective value of a solution is the weighted sum of the total energy cost and total deviation from occupants' preferred lux level ranges. The total energy cost is the sum of the energy costs of all actuators. The total deviation is the sum across all zones of the distance of the predicted lux level from the preferred lux level range.

The blinding actuator is a special case which integrates both the effect of daylight on the environment and the effect of blinding on the daylight. The daylight entering the environment is treated as the base output of the blinding actuator, and it is assumed to have no energy cost. The level of daylight is determined using the actuator settings and sensor readings supplied by the reactive controller when the optimising controller is invoked. Any light which is not accounted for by the settings of the internal lights is assumed to be daylight.

We model the sensors in the environment in the same manner that we model zones. When we have determined the optimal configuration, we compute the illuminance levels the sensors should report when that configuration is implemented. These illuminance values are passed to the reactive controller, allowing it to control around the target illuminance levels as well as simply implementing the actuator set points it is given.

Figure 4 shows the complete actuation model, including the external data sources. When the model is instantiated we search through the space of possible assignments to find the one which minimises the objective function. We do this using backtracking search interleaved with constraint propagation, using the min-dom and min-value heuristics. The model is implemented in and solved with CP-Inside (Feldman, J. & Freuder, E. 2009). We find the optimal solution in, on average, 250 milliseconds.

5 SIMULATION

The control system and its environment have been modelled and simulated using the Charon simulation tool-set (Charon). In this section we briefly describe the simulation models and discuss the results including power savings comparing to a typical control technique used in building automation.

5.1 *Modelling for Simulation*

In order to model the system, two types of agent have been used: Control- and Environment Agents. The behaviour of each agent is described using a hierarchy of modes.

Regarding the Control Agents, one main agent is used for the reactive controller, such that one sub-agent is used to refine the actuation values in each zone using a PI-Controller as depicted in Figure 5. Another agent is used to call the optimising controller. This agent is triggered whenever the user preferences change or a significant change in the actuator configuration occurs. Finally, the sensor agent is used to update the internal light value every sampling period, based on the actuation value, the light interference and the daylight light coming to the sensor. It considers an intensity attenuation factor of $1/r^2$, where r is the distance from the light source to the sensor.

For the Environment Agents, three Charon agents have been used to model Person Movements, Daylight Intensity and Window Blinding Occlusion.

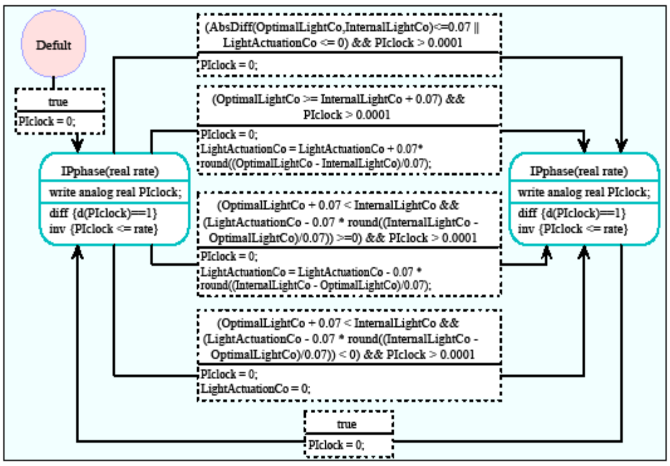

Figure 5. PI-Controller

Due to space limitation we only show one of the control modes depicted in Figure 5.

5.2 *Simulation Results*

To provide an overview of the control strategy and resulting lighting conditions in the office we show Figures 6a-6d. Figure 6a shows the locations of the occupants by zone, with zone 0 referring to occupants who are not present. Figure 6b shows the computed optimal light levels to be achieved by the reactive controller, while Figure 6c shows the actual sensed light levels. These levels differ due to changes made by the reactive controller to respond to the values in Figure 6d, the progression of daylight over the course of the simulated day.

To highlight the effects of the control strategy we consider zones 1, 3 and 6. Zone 1 is adjacent to the window and is thus heavily influenced by daylight. Due to this a pattern of decreases and increases in the internal light level of Zone 1 can be seen as the reactive controller corrects for the loss of daylight towards the end of the day. Shortly thereafter another occupant enters the zone, causing a small increase in the optimal light level, and a corresponding increase in the actual light level as the reactive controller implements the new optimised solution. Zone 6 is seen to have a more stable light level and more closely follow the optimal light level as it is further from the window. Zone 3 maintains a low but non-zero optimal light level despite being unoccupied most of the day. This is due to the optimised light levels being predicted sensor readings, which take into account the effects of all actuators.

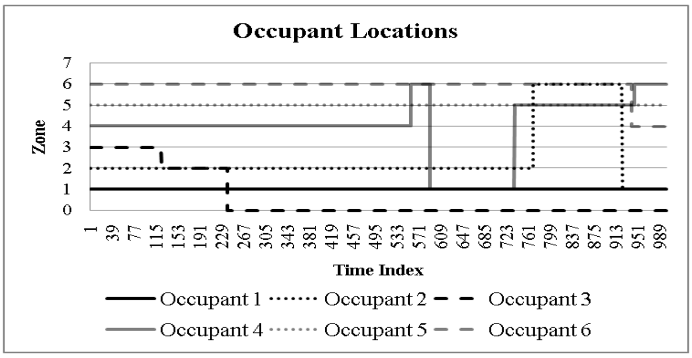

Figure 6a. Occupant Locations

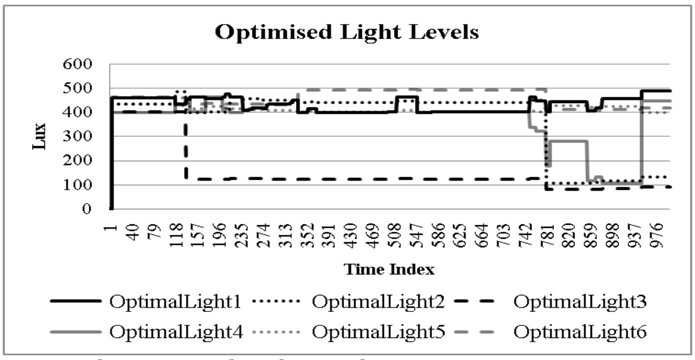

Figure 6b. Optimised Light Levels

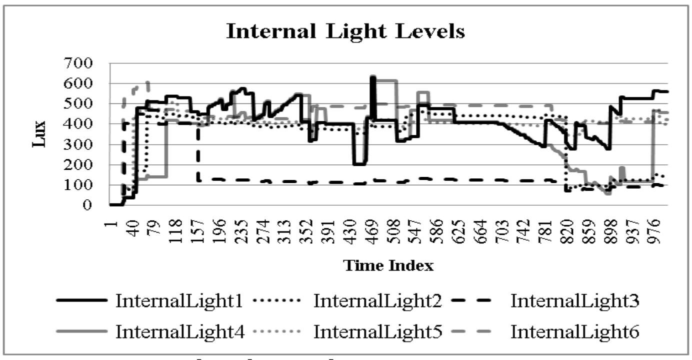

Figure 6c. Internal Light Levels

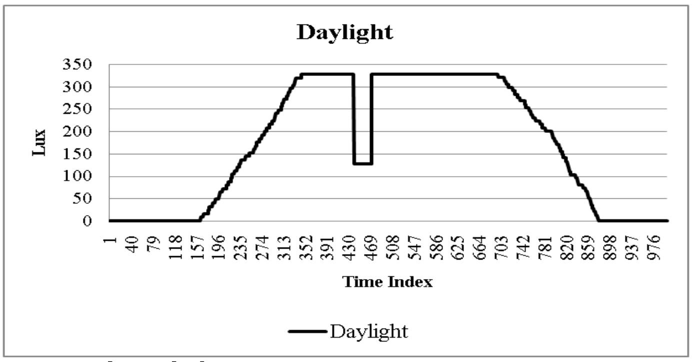

Figure 6d. Daylight

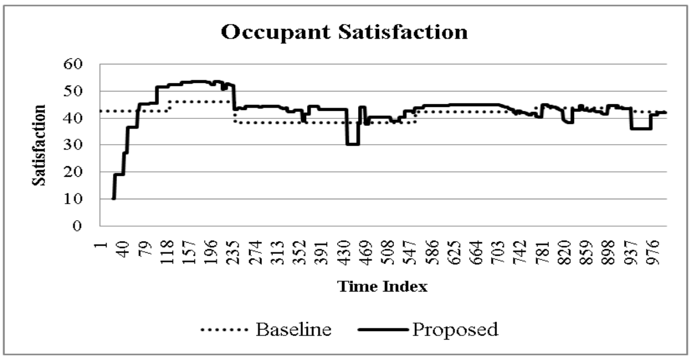

Figure 6e. Occupant Satisfaction

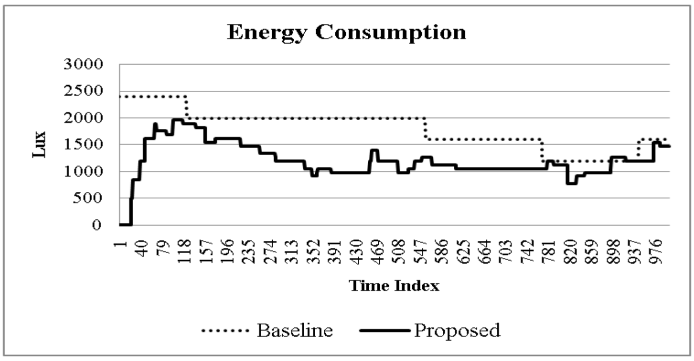

Figure 6f. Energy Consumption

To evaluate the potential energy savings and effects on occupant satisfaction of our proposed strategy, we compare the performance of our controller against a base line control system. This base line is a typical reactive control strategy used in building automation. We assume a passive infrared sensor for presence detection, with the light in each zone switching on to a predefined level whenever at least 1 occupant is in that zone. We consider 400 lux to be the optimal lux level for this strategy, as it has the capacity to satisfy the occupants under all conditions without undue energy use.

Figure 6e compares the actual occupant satisfaction level against that which would be achieved by the base line system. Sudden changes in actual satisfaction occur as there is a small delay between the movement of an occupant and the reaction of the controller. Note, though, that this effect is magnified in the results, since the simulation is running approx. 350 times faster than real time. In practice, the delay is never more than 1 second. These delays are also responsible for the unstable data prior to time index 118, which we disregard when evaluating energy performance below. These delays notwithstanding, our controller maintains occupant satisfaction at or above the base line.

Figure 6f shows energy consumption in terms of lux as energy consumption scales linearly with actuator output. The results show a 30% energy saving over the base line strategy.

6 CONCLUSION

We have proposed a constraint based optimiser and control strategy for an intelligent automated lighting system. The system uses supplied occupant preference data and a tracking system to maintain occupant satisfaction while optimising energy use. Simulation results show that the proposed strategy reduces energy consumption by around 30% compared to a standard base line approach while maintaining occupant satisfaction at high levels. Our current model makes a number of assumptions about the environment and hardware; future work will be to remove these assumptions and find solutions to the resulting challenges.

ACKNOWLEDGEMENTS

This work is funded by Science Foundation Ireland under grant No. 07.SRC.I1170, with the support of the collaborating companies HSG Zander, Cylon Controls, Arup, Vector FM, Spokesoft and Intel.

REFERENCES

Boman, M., Davidsson, P. & Younes, H.L., 1999, Artificial Decision Making Under Uncertainty in Intelligent Buildings, *UAI'99, Stockholm, 30 July-1 August 1999*.

Charon: Modular Specification of Hybrid Systems. http://rtg.cis.upenn.edu/mobies/charon/

Davidsson, P. & Boman, M., 2000, Saving Energy and Providing Value Added Services in Intelligent Buildings: A MAS Approach, In D. Kotz & F. Mattern (eds), *ASA/MA 2000, Zurich, 13-15 September 2000*.

Dechter, R., 2003, *Constraint Processing*, San Francisco: Morgan Kaufmann

Delaney, D.T., O'Hare, G.M.P. & Ruzzelli, A.G., 2009, Evaluation of Energy-Efficiency in Lighting Systems using Sensor Networks, *BuildSys 09*, Berkeley, 4-6 November 2009

EIA, 2003, Commercial Buildings Energy Consumption Survey, US Department of Energy, WashingtonFeldman, J. & Freuder, E., 2009, Cp-inside: Embedding constraint-based decision engines in business applications, *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*: 323-324

Kolokotsa, D., Pouliezos, A., Stavrakakis, G. & Lazos, C., 2009, Predictive control techniques for energy and indoor environmental quality management in buildings. In: *Building and Environment* 44(9): 1850-1863

Mahdavi, A., 2008, Predictive Simulation-Based Lighting and Shading Systems Control in Buildings, *Building Simulation* 1(1): 25-35

Singhvi, V., Krause, A., Guestrin, C., Garrett, J.H. & Matthews, H.S., 2005, Intelligent Light Control using Sensor Networks. In J. Redi, H. Balakrishnan & F. Zhao (eds), *Proceedings of the 3rd international conference on embedded networked sensor systems*: 218-229.