AN ADVANCED STRATEGY FOR ENERGY EFFICIENT LIGHTING INCORPORATING DISTRIBUTED SENSING AND TAILORED CONTROLS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF CIVIL AND ENVIRONMENTAL ENGINEERING AND THE COMMITTEE ON GRADUATE STUDIES OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Amanda Askin

October 2011

Amanda Askin

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Kincho Law
Principal Advisor

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Anne Kiremidjian

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Jack Baker

Approved for the Stanford University Committee on Graduate Studies.

ABSTRACT

Improving energy efficiency in buildings is an increasing national and international priority in response to concerns regarding environmental impacts and global climate change as well as energy security and the stability and condition of the distribution infrastructure. The building sector accounts for 41 percent of total energy use in the United States with 46 percent of this energy consumed by the commercial sector. With 78 percent of commercial building energy use coming from electricity and heavy use times coinciding with daily peak electricity demand, the commercial sector has a disproportionately high impact on the aging electricity grid system. The daily peak window also coincides with maximum daylight availability driving the development and installation of daylight compensation dimming systems. Currently available systems have been demonstrated to achieve significant energy savings, but due to their dependence on distributed control with sparse sensing and actuation points, they lack consideration of the individual needs of occupants and miss potential energy savings.

This research designs a new, tiered energy allocation system for commercial building lighting that effects daylight compensation through the integration of distributed sensing, tailored lighting scenes, and individualized preferences. At the base level of the system are densely distributed sensors and lamp actuators. The sensors collect light level, light level preference, and occupancy information. The actuators set dimming levels on individual lamps to provide maximum flexibility in the lighting scene. The middle level of the system is composed of zone managers which coordinate the sensors and actuators within their zones and use the provided information to create a zone energy use utility curve. The top level of the system is

the building server which allocates building-wide energy resources in accordance with the utility curves from all zones. The zone level is capable of acting quickly in response to changing local conditions while the centralized building level enables energy use and performance tracking and facilitates energy use curtailment, either in response to a demand response request or to minimize use during peak pricing, by allocating limited energy resources to parts of the building that can best put them to use. The advantages of this tiered resource allocation system center around defining the explicit relationship between energy use and the service level provided to occupants. Basing resource allocation decisions on the relative energy cost of quantified performance improvements across the building enables the system to focus on maintaining the maximum achievable performance level for the occupants while using minimal resources. This service-based approach enables the explicit designation of acceptable service level standards and the inclusion of relative importance weighting for individuals and areas throughout the building.

A prototype wireless hardware system is designed to show the implementability of the system and a building simulation tool is developed to show the system performance. The prototype hardware system is designed on a wireless platform to demonstrate the applicability of the resource allocation system to retrofit projects as well as new construction and to emphasize the adaptability of the system for use in reconfigurable spaces. The evaluation of the tiered resource allocation system shows both decreased energy use and improved occupant performance as compared to conventional dimming systems.

ACKNOWLEDGMENTS

I would like to thank my dissertation advisor, Professor Kincho Law, for his guidance and encouragement throughout my time at Stanford. I would also like to express my gratitude to Professor Anne Kiremidjian, Professor Jack Baker, Professor Ram Rajagopal, and Professor Elizabeth Sattely for serving on my thesis committee and providing me with insightful feedback on the presentation of my work.

This work would not have been possible without the guidance of past members of the Engineering Informatics Group, Professor Jerry Lynch, Professor Yang Wang, Professor Jack Cheng, and Dr. Julie Ekstrom, and I am grateful to them for sharing their technical expertise and research experiences with me.

I am thankful for the wonderful support of the staff and all of the students in the John A. Blume Earthquake Engineering Center, with special thanks to Andy Myers, Marc Ramirez, Wil Srubar, Madeleine Flint, Victor Victorsson, Ting Lin, Chris Smith, Allen Cheung, and Racquel Hagen. The students in the Blume Center have been a great source of encouragement and research guidance throughout my time at Stanford.

I am truly grateful for all of the support I have received from my family. Their constant support has always helped to push me in the right direction. I am also thankful for all of my extraordinary friends, especially Lindsey, Kristen, Joan, Dan, Helen, Heather, Valerie, Colin, Warren, Jeff, Elliot, Scot, and Alex.

This research has been sponsored by the National Defense Science and Engineering Graduate Fellowship, the Stanford University Leavell Fellowship, and the John A. Blume Earthquake Engineering Center.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1.

INTRODUCTION

1.1 BACKGROUND

Providing safe, affordable, and reliable energy for now as well as for the future is one of the most compelling challenges facing society. As the world population increases and standard of living improves, energy demand also rises. Worldwide primary energy consumption increased an average of 2.6 percent per year from 1998 to 2008 (EIA, 2009), a trend that is predicted to continue. As demand increases, resource supply and delivery ability are limited. Over the same period, world electricity consumption increased at a rate of 3.4 percent per year (EIA, 2009), requiring enormous investments in generation, transmission, and distribution infrastructure. Despite billions of dollars invested per year in electricity transmission lines in the United States, the increase in demand outpaces the increased supply and the transmission infrastructure continues to be outdated and in need of upgrading (ASCE, 2005).

Beyond transmission issues are limitations in electricity generation, political stability concerns, and environmental considerations. While peak oil timeline estimates vary, some researchers believe the peak has already occurred (Aleklett et al., 2010). Coal resources are considered more abundantly available, but the ability to make use of these resources is directly affected by greenhouse gas emissions policy. While renewable resources are becoming more widely available and new ways of integrating them into the electricity grid are being developed, renewable energy including solar, wind, hydroelectric, biomass, and geothermal currently totals only 8 percent of national energy use in the United States (EIA, 2011b). With the dependence on foreign sources of oil to meet electricity demand, the United States is highly vulnerable to internationally controlled markets and to the stability of the relations with other countries. Wherever the energy source is located and whichever generation method is used, environmental consequences are a consistent concern. Even among domestic, renewable energy resources, environmental concerns persist regarding the manufacturing of photovoltaics, the installation and operation of wind and water turbines and geothermal generators, and the harvesting of biomass. Decreasing energy use directly reduces the associated supply, distribution, security, and environmental problems. Improving the efficiency with which energy is used to provide for the needs of society is a necessary component of solving the continuing energy challenge. This thesis proposes a new strategy for energy efficiency in commercial building lighting which comprises a significant portion of energy use.

1.1.1 Energy Use in Commercial Buildings

In total, buildings including both commercial and residential facilities account for 41 percent of total energy use in the United States, larger than both the transportation and industrial sectors. Therefore the efficiency of energy use nationwide is significantly impacted by the efficiency of these buildings. In the United States, the commercial sector accounts for over 19 percent of total energy use with 78 percent of this energy use coming from electricity use and electrical system losses thus comprising nearly 40 percent of total electricity use (EIA, 2010). The electricity grid is of particular significance as the aging grid infrastructure experiences decreasing distribution efficiency and increasing reliability concerns. For the United States, it is predicted that increases in electricity demand for electronics will more than compensate for potential energy savings from implementation of currently available efficiency technologies (EIA, 2011a). Ever more advanced efficiency technologies are necessary to counteract this increasing trend.

The commercial sector is defined to include all non-residential spaces that are not used in manufacturing or transportation. This category includes such spaces as shopping centers, schools, hospitals, restaurants, and office buildings. Together, office buildings and retail spaces comprise over 50 percent of commercial building energy use in Europe and in the United States (Perez-Lombard et al., 2008), thus comprising a sizeable portion of the total overall energy consumption.

In addition to the large contribution to total energy use, commercial buildings have a disproportionately large effect on the peak electricity demands. For example, in the state of California, commercial buildings are projected to account for 36 percent of coincident peak electricity demand in 2011 (Kavalec & Gorin, 2009). Because commercial buildings typically reach their highest occupancy rates at the hottest hours of the day, the ventilation, cooling, and lighting demands for these buildings contribute significantly to the afternoon electricity demand peak effect. Ultimately high peak demands are very costly to utility companies and this cost is passed on to consumers. Entire energy plants are designed for use only in peak demand times which occur very infrequently but drive up the overall system cost and redirect resources from maintaining and expanding distribution and transmission infrastructure toward adding capacity. Because these facilities are offline most of the time, they are enormously underutilized and expensive assets. If the peak energy use can be decreased for the short duration of normal occurrence, this investment in additional generation capacity would be unnecessary.

From a building owner standpoint, the ability to predict and minimize operating costs is crucial to the economic viability of the real estate investment. In 2000 to 2001, California suffered from a severe wholesale electricity price fluctuation due to deregulation, an increase in fuel costs, and increased demand (Sueyoshi & Gopalakrishna, 2008). This drastic increase in electricity cost hurt commercial business consumers who were unable to predict the 500 percent price increase and did not have the means to restrict energy use in an intelligent way to combat the increase in operating costs.

1.1.2 LIGHTING AS A COMPONENT OF COMMERCIAL BUILDING ENERGY USE

A 2007 study by McKinsey & Company identified the use of more efficient lighting in commercial buildings as a potential source of significant energy and cost savings (Creyts et al., 2007). Commercial building control systems were also identified as an opportunity for combined emission and cost savings. This considerable potential savings arises from artificial lighting accounting for 25 percent of total fuel consumption and 39 percent of electricity consumption in commercial office buildings (EIA, 2008).

Commercial building lighting also contributes significantly to the peak load problem. Because commercial facilities are generally heavily used in the afternoon hours, the demand for lighting is also highest during these peak utility usage hours. Commercial lighting comprises 30 percent of summer peak demand in California (Rubinstein & Kiliccote, 2007). During hot days, the use of lighting can further increase cooling demand due to waste heat emitted by lighting fixtures. The impact on total energy use is tempered by the increased heat necessary for cold days, however because space heating is delivered via a mix of gas and electricity, the peak electricity grid concerns are more significant on the hot weather cooling days (Sezgen, 2000). These peak usage hours coincide with periods of high daylight availability, which can reduce the energy demand for the commercial sector when the natural light is properly harvested and compensated for by reducing reliance on artificial lighting systems. As utility companies develop and refine demand pricing schedules, it becomes more advantageous to commercial building owners and tenants to reduce usage during these peak times. Reducing the use of artificial light during high electricity demand times enables building owners to save money while reducing environmental impacts and system loads.

The McKinsey report references to energy efficient lighting point to the installation of more efficient lighting types such as light emitting diodes (LEDs) and compact fluorescent bulbs (CFLs). CFLs are slated to replace incandescent bulbs where they remain in use in the commercial industry and LED technology is progressing toward use as a primary light source for building installation. However, currently over 92 percent of commercial building floor space and 98 percent of commercial office floor space is at least partially illuminated by standard fluorescent lights (EIA, 2008). Therefore any lighting control system aimed at reducing the impact of commercial lighting energy use for existing as well as new buildings must be capable of interfacing with fluorescent lighting systems. Looking toward the future, a control system designed with applicability to a wide variety of lighting types is advantageous as new technologies become available.

1.1.3 Current State of Energy Efficiency Technologies

Although building designers have started to implement available energy saving technology and energy efficient lighting such as tubular fluorescents with electronic ballasts have been widely adopted, the use of sensor-linked control systems remains below 10 percent based on floor space and the percentage of daylit space is near 14 percent (Andrews & Krogmann, 2009). The adoption of new technologies is driven by large, high-end, owner-occupied buildings where upfront investment is seen as an offset to long term operating cost reductions with older buildings being slow to adopt new technologies (Andrews & Krogmann, 2009). However, as building standards become more stringent and the cost savings available through use of new technology is demonstrated, technologies designed for easy integration with existing buildings will be able to capture the potential savings that exist throughout the building stock.

1.1.3.1 Available Components: Occupancy Sensors, Lamp Dimmers, and Ballast Control Systems

A wide variety of components is currently available for building installation. One of the most commonly used energy efficiency components is the occupancy sensor. Occupancy sensors use a variety of technologies to detect the presence or absence of people in the space and turn off lights when the space is empty. Occupancy sensors use sound responsive technology, passive infrared technology, ultrasonic technology, or a combination of these to provide enhanced performance. Sound responsive sensors essentially work like a microphone and pick up local noise. Some sophisticated sensors are capable of learning background sounds over time and filtering them out to determine sounds most precisely correlated to occupancy.

Passive infrared sensors are essentially heat sensors that detect the presence of warm bodies but require line-of-sight to the occupant. Ultrasonic sensors detect changes in the return of emitted high frequency waves which signal movement. Dual-technology sensors incorporate passive infrared and a second technology to minimize errors. Typically both sensors must be triggered to turn on the light but only one must be triggered to keep the lights on. Examinations of installations in real buildings have shown that the use of occupancy sensors can save up to 26 percent of lighting-related energy in private office spaces (Jennings et al., 2000).

Dimmers and light sensors can be very effective lighting energy reducing tools. Typically a photodiode sensor is installed and calibrated to maintain a set point where the incident light energy at the work surface is equivalent to the required desktop illuminance. In an open loop system, the sensor is angled toward the daylight source and the room lights are dimmed in accordance with the measured daylight contribution. In a closed loop system, the sensor is positioned to determine the total light level on the work plane and adjusts the artificial light sources as appropriate to maintain the desired light level. Both of these types of systems can be difficult to install and properly calibrate because each sensor can be responsible for dictating the lamp settings for a wide area and changing room conditions can affect the system performance. The dimming effect can be achieved either through the use of a continuous dimming system where all lights are capable of dimming through a wide range of fractional output or bi-level dimming where a dimming effect is achieved by switching off individual lamps within a fixture to vary the output. Studies show significant savings are available from the implementation of daylight-linked dimming systems. Study of a large federal office building in San Francisco found a 27 percent light energy use savings from automated dimming controls installed in private offices with external windows but noted that the savings would have been higher if the target light level had been adjusted to a lower, task-tuned level (Jennings et al., 2000). By contrast, only a 9 percent energy savings was recorded for the use of non-daylight linked, manual dimming. Another study in the same building found up to 41 percent light energy savings in the outermost row of fluorescent lights in a shared office space and up to 22 percent savings in the second outermost row (Rubinstein et al., 1998).

Measurements of high frequency dimming controls in a Hong Kong open plan office found a 33 percent energy savings due to the use of dimming controls (Li et al., 2006). Another study conducted by the same group found up to 50 percent energy savings in electric lighting for perimeter offices (Li & Lam, 2001). The variation in the findings of these studies are related to the differences in location, floor plan, control systems, and building types of the study targets, but all results indicate significant potential for savings.

A variant on the daylight-linked dimming system is the daylight-linked on/off system which does not have dimming capabilities but turns the lights off when enough daylight is available in the space such that artificial light is not necessary. These systems have also been shown to produce savings, though less significant savings than are achievable with their dimming-capable counterparts. By contrast, a study of lighting controls in atrium spaces found that use of a continuous dimming system saved 68 percent of lighting energy while an automatic on/off system saved only 31.5 percent of lighting energy during the time of main occupancy for the facility (Atif & Galasiu, 2003).

Typically daylight compensation systems are structured such that the light settings in a large room or set of smaller rooms are determined by the light level measured at a single centrally located, ceiling mounted sensor. Possible performance issues related to this setup have been noted (Love, 1995; Littlefair & Motin, 2001), such as the inability of the system to guarantee sufficient light levels for the use of the space due to variable shading conditions and use of blinds and difficulties in responding to changing sky conditions.

Achieving the energy-saving continuous dimming in a fluorescent ballast is more complex than for an incandescent bulb. For incandescent lights, the applied current signal can be clipped around the zero crossover points in the alternating signal to achieve the reduced power consumption and reduce the output without negative impact on the lamp system. However, in a fluorescent tube, maintaining a minimum quantity of current flow is necessary for the health of the bulbs. Special dimming ballasts have been developed to provide the reduced level of light output in accordance with a received control signal. These ballasts are controlled with a 0-10 Volt analog signal, a modulated power signal, a digital message according to a specialized protocol, or are interfaced with wireless communications. The advantage of the power signal and wireless communications methods are that they require no extra wiring and can be installed easily in a retrofit project. The digital protocol and wireless communication ballasts allow for individualized settings to be selected for all ballasts. The tradeoff for this tailored light setting ability is an increased cost and limited array of available products. Wireless ballasts are an emerging technology that is not currently widely available or typically installed in buildings. The move toward wireless technologies is driven by the need for reduced installation and commissioning costs and to facilitate the use of denser sensing and control networks as conventional building control system installation can account for up to 80 percent of the total system cost (Kintner-Meyer & Conant, 2005). Investigations on the use of wireless technologies for the building controls sector indicate the focus of wireless system development must be on scalability, reliability, customizability to a variety of interfaces, and the ability to integrate both local and building-wide controls. Zigbee 802.15.4 protocol is widely suggested as the appropriate wireless protocol for this application due to its minimal power consumption, mesh networking capabilities and the ability to auto-detect new nodes as changes to the system are made (Kastner et al., 2005; Kintner-Meyer & Conant, 2005; Menzel et al., 2008). As the technologies advance and as more sophisticated control systems are developed to make use of the additional features of the enhanced ballasts, the market for these components is expected to increase driving down costs in the longer term.

1.1.3.2 System Needs: Responsiveness to Individual Preferences

Compensating for the influx of daylight by reducing artificial light output both saves energy and brings provided light levels closer to light provision standards and preferences. A survey of workers in a St. Louis, Missouri office building found that they considered lighting to be one of the most important environmental factors (Ne'eman et al., 1984). The study showed over 90 percent of the workers ranked quality and quantity of light for completing work-related tasks to be very important. The importance of maintaining office worker comfort has been studied and links between comfort, health, satisfaction and productivity have been established (Wyon, 2000; Roulet, 2006; Newsham et al., 2009).

In a study looking at office lighting preferences of secretarial workers, a wide variation in preferences was found. While the findings showed differing preferences between workers, consistency over time for individual preferences was found, indicating that the experimentally derived preferences were significant to the individuals (Tregenza et al., 1974). The nature of office work has changed since this study was conducted; however, the indication of variation in human preference remains applicable. A wide range of individual lighting preferences is reinforced by a later study that again showed variability in desktop illuminance preference, ranging from 100 to 800 lux with over 60 percent of occupants preferring light levels below IESNA recommendations (Veitch & Newsham, 2000). Due to the variability in preferences, ensuring everyone is comfortable is not a simple task, and thus as Wyon states, "The best solution is a workplace where you can change something". An experiment in comparing occupant satisfaction in a laboratory room mock up under fixed lighting conditions and under occupant-controlled conditions found improved mood and satisfaction and decrease in discomfort when the occupants were allowed to control their own environment (Newsham et al., 2004). The results were found to be attributable not only to the provision of individual control, but to the actual implementation of the preferred settings. The desire for less light than that recommended by the prevailing standard points to potential energy savings from meeting the actual preference levels instead of provision of light in accordance with the standard. A study of the utilization of individual dimming controls in 58 private office spaces found that nearly three quarters of occupants chose to have their lights set to less than full output with several occupants choosing to work with the lights off entirely (Maniccia et al., 1999) further indicating that potential savings exist for allowing individuals to set their preferred light levels.

In a study examining the manually selected diming settings of office workers in a daylit space, occupants selected dimming levels that offset the availability of daylight indicating a preference for daylight compensation measures and desktop illuminance level was found to be the best predictor of dimming level selection (Newsham et al., 2008). The variability in light level preference, the occupants' desire for individual control and ability to implement the preference, and the relationship of measured desktop illuminance level to control actions indicate a need for a lighting control system that provides individualized dimming control relative to measured desktop illuminance levels.

1.1.3.3 Demand Response Systems

Demand response systems fall under the umbrella of peak load management. Unlike demand limiting systems which simply cap the peak usage during a time of anticipated peak usage or demand shifting systems which strive to shift energy intensive systems to use during non-peak hours, demand response (DR) systems are event driven. DR systems respond to excessive demand on the electricity distribution infrastructure or to inflated, demand-responsive pricing by curtailing the use of electricity.

Currently installed demand responsive systems for lighting are characterized as either absolute reduction systems, meaning the building lights are switched to a preset conservation scene, or relative reduction systems, where all lamps are dimmed by a certain fraction from their current state. In the absolute reduction case, the system is unable to guarantee any savings if the light settings already are at levels equivalent to the conservation scenario. The relative reduction case requires more advanced control systems and knowledge of the current light level settings (Rubinstein & Kiliccote, 2007). Neither system is capable of providing information about the actual comfort cost to the occupants nor do they necessarily carry out the energy use reduction in an optimal manner. Field studies conducted to assess the performance of DR lighting systems demonstrate that these systems can achieve anticipated energy savings levels; however, the studies rely on the presence or absence of complaints to the building manager and actions taken by occupants to change the light settings after the event to determine the acceptability of the response implementation, which unfortunately are imprecise measures of overall occupant satisfaction (Newsham & Birt, 2010).

One of the major barriers to implementation of DR systems for lighting is the lack of market penetration of lighting controls. Only 7 percent of commercial floor space has lighting controls, but this number is anticipated to increase for at least California buildings with the implementation of California Title 24 which increases building performance standards (Rubinstein & Kiliccote, 2007). Reasons for the limited use of centralized dimming systems for DR include increased initial costs, the need for utility companies to have verifiable, readily accessible load reduction sources, and the concern among building owners about maintaining contractually obligated performance for occupants (Newsham & Birt, 2010).

The installation of DR systems has implications not only for limiting peak demand but also for improving distribution system reliability. Investigations have been conducted into the use of DR-type reserves as spinning reserves to offset problems in the grid. Spinning reserves are the extra generation capacity that the utility operator has available to quickly respond to fluctuations in demand. Traditionally this reserve is provided by extra generators, but with significant demand side response abilities, the total available reserves could be increased, thereby allowing operators more flexibility in handling high-demand situations and avoiding rolling blackouts (PIER, 2007).

1.1.3.4 Control and Optimization Systems

Based on the potential for energy savings and improved occupant satisfaction, new types of sophisticated building system controls have been proposed for both lighting and HVAC systems. These control systems are typically either fully decentralized or centralized systems. Decentralized systems enable quick local response but are unable to provide centralized information regarding the full state of the building. Fully centralized systems can be slow to respond to changes in conditions and require large stores of data. The systems are typically designed around the optimization of a single parameter or a single function inherently balancing multiple parameters. Where a single parameter is optimized, usually either comfort or energy, the impact the system has on multiple performance parameters is overlooked. In optimizing a single function, the tradeoff decisions must be predetermined, no information is provided about the level of performance for any individual parameter, and there is no ability for systems to interact unless they are all tied in to a very complex objective function which further obscures the performance metrics. Typically the control systems are feedback based but do not incorporate spatially distributed sensing and control on a small scale. Generally one sensor is used to assess the conditions in a wide area, and this sensor measurement is used to universally apply a new setting across the assigned area based on standard code-based set points. This method does not allow for flexibility in addressing variable conditions across a space nor does it allow for the incorporation of variable preferences.

The most commonly installed lighting control systems are straightforward feedback loop systems. A similar approach has been derived for an integrated artificial lighting and blinds controller with blind position and light dimming fraction set directly from sensor measurements (Roche, 2002). Results of this system in a summertime installation show that it maintains good visual comfort performance with minimal computational complexity. One major drawback to the blinds controller is that it requires knowledge of the geometry of the space which would require individualized programming at installation. Additionally, the ceiling-mounted sensor used for inference of the desktop illumination also must be calibrated on site and requires manual recalibration after changes to the arrangement of items in the space.

A group from Carnegie Mellon University has proposed a distributed sensor network and control strategy for optimal light control that uses a single optimization function to maximize their definition of utility. Total utility, in this case, is defined as a combination of predefined user utility functions and a function representing building manager utility, essentially a representation of resource usage (Singhvi et al., 2005). This control system can be modified to incorporate the actual preferences of the occupants and uses a spatially distributed sensing and actuation system, but does not convey information beyond the light setting output of the optimization program. The system falls in the category of decentralized systems as it is implemented on a zone-by-zone basis throughout the building, making it unsuitable for DR applications and for providing performance information to the building manager.

Others have proposed the use of neural networks to enable the building system to adapt to the comfort preferences of the users (Sierra et al., 2006). These systems generally require storage of very large quantities of past data for constructing and adapting the network. Typically the focus of these systems is exclusively on providing quality environmental conditions for occupants without specific regard to energy use.

Fuzzy controllers have also been suggested. One type fuzzy controller for building control has been designed with a pre-determined cost function that incorporates the major factors of interest to the control designer with relative importance weightings (Sierra et al., 2007). Simulations performed by the researchers have indicated that this fuzzy control system can be effective at tracking set points under variable external loads, but it does not incorporate individualized preferences or enable tailored controls. While the proposed system attempts to incorporate both HVAC and lighting systems, it is limited by its fully decentralized implementation. An adaptive neuro-fuzzy controller for lighting has also been suggested for an artificial lighting and automatic window blinds integrated controller that considers impacts to both lighting and HVAC systems though evaluation of both daylight and solar heat gain (Kurian et al., 2008). This system relies heavily on extensive training data sets, and although it is capable of considering both visual and thermal comfort in conjunction with energy savings, the tradeoffs between comfort and savings and relative weightings are not explicit.

Genetic algorithms (GA) have also been suggested for in situ parameter updating for control decisions (Guillemin & Morel, 2001). The window blinds control scheme suggested by Guillemin et al. utilizes GA for nightly parameter updates. These updates are based on knowledge of manual overrides to the automated system, i.e. when window blind angle was set by the automated system but changed by the occupant. In this way the system adapts to the preferences of the occupants. This approach is useful for blind control because of the generally poor characterization of objective criteria for determining the preferred settings of blinds in daylit spaces with respect to known information including sun position and irradiance. This technique may have implications for other systems such as ventilation and temperature which have complex interactions with respect to comfort; however, with environmental conditions such as desktop illuminance where occupants are capable of explicitly specifying their preference, more reliable information can be gained by providing them with direct selection.

As the algorithms for building control become more sophisticated, industry and research entities continue to develop the hardware and communication systems that implement these algorithms. Web-based monitoring systems are of interest because they make use of the already existing computer hardware and networking abilities and can allow building managers to monitor buildings remotely and building occupants to view their environment and make changes to their preferences conveniently at their desktop (Kensek et al., 2000). A proposed trend is for the integration of wireless networks with web-based systems to take advantage of the flexible distributed sensing capabilities of a wireless network and the convenience of internet-based retrieval of data (Jang et al., 2008).

None of the control methods discussed in this section provides the ability for a building manager to assess the tradeoffs between energy use and quality of service provided to occupants. Decision parameters are set at system design and commissioning stages and these values are held constant for the lifetime of the system yielding no ability to dynamically assess building conditions or identify potential energy tradeoffs. If the systems have a centralized component to their installation, during a demand response scenario they can choose a universal dimming fraction for all lamps or attempt to reallocate resources based on previously indicated, static priorities in order to lower the electrical load due to lighting, but they cannot provide information on where the true energy versus comfort tradeoff points are or what the decrease in comfort will be due to the energy curtailment.

1.2 Research Objectives

This research project is designed to investigate the potential for a new type of energy efficient building system controls. This new type of control incorporates distributed sensing, occupant preference and location, tailored control capability, two levels of decision making, and wireless communications to form a building-wide system that minimizes energy use while providing a high quality occupant experience.

Contemporary building system controls use relatively simplistic, typically wired, feedback control methods leaving many questions to be answered in moving forward toward a more sophisticated system. These research questions span the design, implementation, and performance and evaluation of the system.

Questions with respect to the system design include:

- How should a tiered resource allocation scheme be structured to enable local, real-time response in conjunction with global system standards and resource limitations?
- How can decision analysis concepts be applied to efficient resource allocation? How would a utility function be defined for a collection of individuals for use in this resource allocation?
- How can the potentially opposing goals of minimizing energy use and providing quality occupant performance be effectively incorporated?

Questions with respect to implementation include:

- What is the most effective way to structure the wireless communication?
- How are wireless units assigned to groups to perform tasks?
- What are the necessary considerations in the design of distributed sensing and control capabilities?
- Can the necessary decision-making and optimization capabilities be embedded on a microcontroller for distributed processing?
- What is the energy overhead of installing this type of system? To what extent does it impact the energy performance of the building?

Questions with respect to performance and evaluation include:

- How effective is this type of system effective at conserving energy use?
- How effective is this type of system at meeting occupant performance requests?
- Are there additional benefits to implementing this type of system?

- How should a model be designed to assess the performance of the system at full building scale?

This research is designed to find answers to the above questions and to demonstrate the value of an enhanced control system. The goal is not to create a single, commercial grade product, but instead to propose a new way of thinking about energy efficiency systems and to bring in concepts from decision analysis and mathematical programming to develop a better type of resource allocation system than those that are currently available.

1.3 Thesis Outline

This thesis is organized to provide a comprehensive picture of a new type of commercial building energy efficiency system. The research incorporates both physical implementation and computer simulation to investigate the anticipated behavior of the proposed energy control framework in a full scale building system. This document is structured to first give an overview of the structure and design of the proposed, novel system, followed by a demonstration of physical implementation, and lastly to give evidence of the potential benefits of the system through larger scale simulation. The specific outline of this thesis is below.

- Chapter 2 of this thesis addresses the design and structuring questions with an overview of the system components. It begins with an introduction to the tiered structure of the system and a discussion of the roles of actors at each level.
- Chapter 3 is tasked with addressing initial questions with regard to physical implementability by introducing a laboratory prototype system designed as a small-scale demonstration. This chapter details the design of the wireless sensing, computation, and control unit and the software backbone of the system.
- Chapter 4 details the optimization and resource allocation algorithms used in the decision-making scheme including both the design of the initial laboratory

system as well as alternative formulations considered in the design of the simulation.

- Chapter 5 is tasked with addressing further questions related to realistic physical implementation through a discussion of the results of a small-scale laboratory demonstration.
- Chapter 6 investigates the full building scale performance and evaluation questions with a discussion of the development of a building model suitable for testing the algorithm and the results gleaned from this simulation. The simulation discussed in this chapter is used to evaluate energy performance as well as occupant preference tracking.
- Chapter 7 provides a summary of the project and offers overall conclusions and suggestions for extensions. The summary discusses the implications of the research project and the potential for integration of the sensing, actuation, and resource allocation system into a real building environment.

Chapter 2.

TIERED RESOURCE ALLOCATION SYSTEM DESIGN

The proposed resource allocation system is designed to control the light settings throughout a building in accordance with the following goals: minimize energy consumption, provide quality occupant performance, facilitate efficient and minimally invasive demand response, and provide meaningful decision making information to building owners and managers. To accomplish these goals, all energy use in the system is characterized in terms of the degree to which it contributes to the overall performance. Defining this relationship between energy use and performance is critical to changing the perspective that increased energy use is necessary for increased performance and to providing the decision makers the tools required to evaluate opportunities for increased efficiency. The conventional perspective on lighting energy use is to consider energy use in terms of light output at the source; however, the system described in this thesis shifts this perspective to consideration of energy use in relation to the level of light service at the points of interest. Shifting to a service-based approach allows for the potential integration of other building systems to share energy resources with the ultimate goal of providing a high level of total service to the occupants with minimal resources.

2.1 SYSTEM REQUIREMENTS

Providing quality service necessitates the acquisition of information about the continuously evolving state of the building. Collecting this information requires a distributed array of sensors capable of gathering current information about the occupants' presence, preferences, and experiences. Timely and tailored response to this evolving state is crucial to maintaining a high performance standard and minimizing energy use. Localized decision making is necessary in order to facilitate a timely response, and tailored settings require distributed points of control. Centralized information and decision making capabilities are required in order to provide system-wide performance tracking and energy use information and to implement building-wide standards and usage limits. Supplying the central repository with raw data from all sensors and settings for all individual control points is taxing to the communications and central processing infrastructure of a building system and does not provide actionable, decision-related information. Therefore an efficient system must process its local information into a readily digestible form for the central actor such that communication requirements are reduced while relevant, meaningful information is shared. Tasking the local, lower level processors to consolidate the information and process it into a characterization of their respective needs for energy resources, and the corresponding level of service they are able to provide with those resources, allows the central actor to directly use this report from all building subsections to allocate energy units across the building and to be continually informed of the occupant service level provided across the building.

2.2 Tiered System Design

A tiered resource allocation system is developed to combine the benefits of decentralized and centralized systems. The decentralized, distributed control level of the system is able to respond quickly to changes and update lighting scenes as necessary while the centralized level is able to monitor system performance, set total energy use limits, initialize demand response, and manage the quality of the occupant experience. The distributed processing at the decentralized level of resource allocation allows multiple microprocessors to solve small scale problems in parallel which minimizes computational and information storage and transfer requirements at the building-wide level. This section details the design of the system and the decisions made at each level and the tiered structure of the system is diagramed in Figure 2-1.

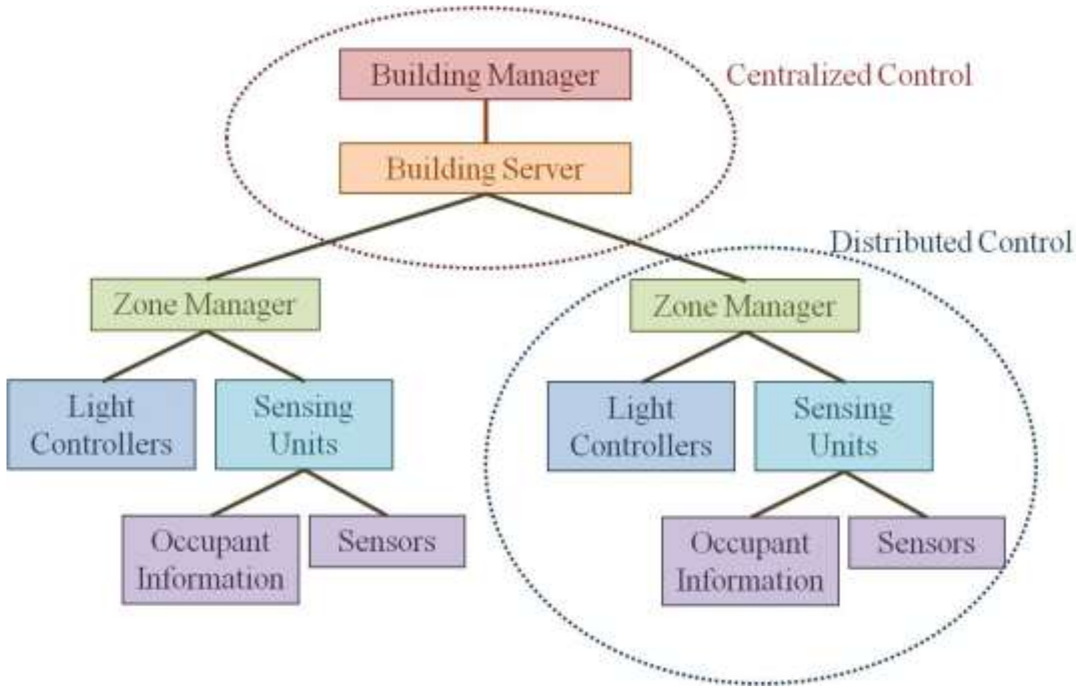

Figure 2-1: Control and Allocation Hierarchy

At the top level sits the building manager who is responsible for setting all relevant system parameters and has the ability to view pertinent information regarding system performance as necessary through an interface with the building server. The building manager and the building server comprise the centralized control level. Below the centralized level sit the zone managers responsible for making decisions for the specific sections of the building. The zone managers interact directly with the control and sensing units within their section of the building to obtain information about and alter the lighting state of their respective zone. The zone managers, in conjunction with the sensing and control units, compose the distributed control level of the system.

The sensing and control units are spread throughout the building with each control unit affixed to a fluorescent ballast or small set of ballasts and each sensing unit representing a location of importance in the building. As depicted in Figure 2-2, each occupant is entitled to a designated sensor which he or she can use to set a light level preference and which keeps track of the light level and occupancy in the space. The building is subdivided into zones with each zone assigned a zone manager that is responsible for coordinating the sensing and control units within its zone. The centralized building server communicates with all zone managers to allocate the total building energy to the zones. Block diagrams of the elements of this tiered structure and their physical location within a building are depicted in Figure 2-2. As shown in the figure, a zone is an identified section of the building assigned one zone manager and many sensing and control units.

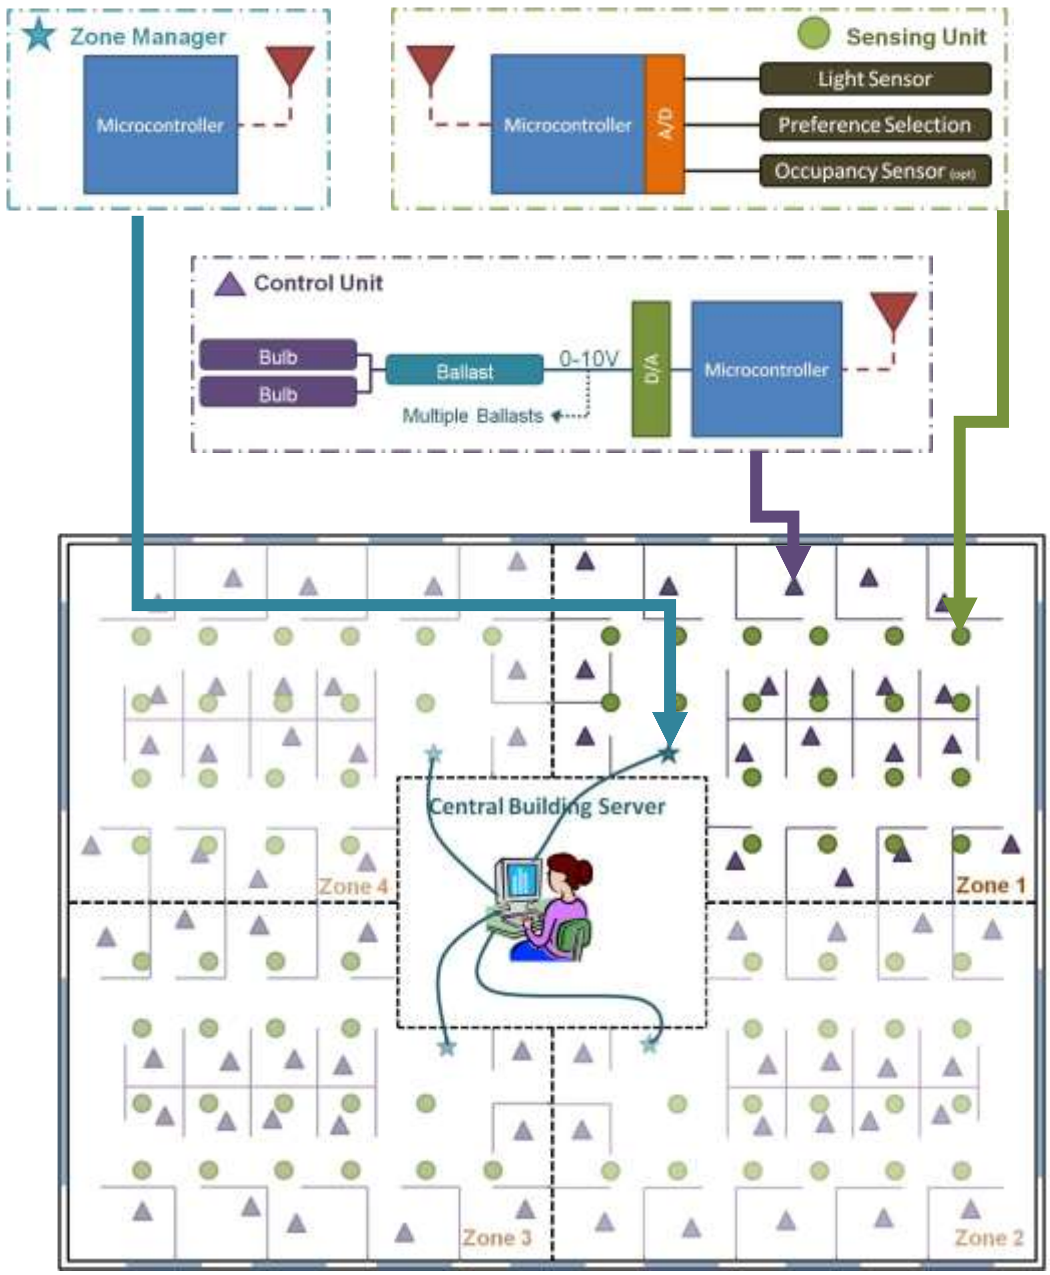

Figure 2-2: Layout of System Elements

2.2.1 BASE LEVEL SENSING AND CONTROL

The sensing and control units are the frontline agents in the system. They are respectively responsible for collecting sensor data and controlling the light output level of the lamp on requests from the zone managers. These units communicate wirelessly with the zone managers. In a real-world implementation, the sensing and control units would be designed separately, specific to their particular functionality. As discussed in Chapter 3, a single unit capable of performing both of these functions was designed for the physical implementation portion of this project due to the desire for flexibility in the laboratory setting.

The purpose of the distributed sensing is to have a dispersed sensor network responsible for collecting information for each point of interest throughout the building. In addition to sensors positioned in common areas, each occupant would have his or her own designated sensing unit that would report back on the local conditions. Information of relevance collected by the sensing unit includes the current illuminance level, occupancy status, and preferred illuminance level. Illuminance level is necessary for calculating daylight contribution and initially for determining the influence of each light to each point of importance. The preferred illuminance level is used as the target light level for the optimization scheme and occupancy status is necessary for determining whether this preference is incorporated into the particular round of optimization calculations. If the space is unoccupied, the individual preference is irrelevant and that constraint can be removed from the system.

The distributed control allows the system to set each lamp dimming level individually as necessary to optimize over the entire system. This flexibility creates the opportunity to determine the precise lighting scene that both meets the occupant preferences and uses the minimal amount of energy. While the prototype sensing and control system is constructed specifically to interface with fluorescent lights, the decision-making and resource allocation structure discussed in this chapter is designed to be flexible for implementation with any type of lighting system.

The physical design of these sensing and control units requires: the ability to communicate with the zone manager, a processor, and the appropriate peripherals for executing the relevant task. For the physical implementation addressed in this thesis, communication is conducted via wireless radio due to the relatively dense nature of the sensing and control network. A wired communication system could alternatively be developed but would require significantly higher installation costs in an existing building. The sensing units require analog inputs for connections to both a light sensor and a potentiometer for selecting the preferred light level setting as well as a digital input for interfacing with an occupancy sensor. The control unit requires the appropriate output type for communication with the selected type of dimming ballast.

2.2.2 Zone Level

The zone level units are the coordinators and leaders of the zones. A zone may be defined as a large shared office space or several smaller offices within a building. Ideally the zones are photo-isolated from one another although strict enforcement of this restriction is considered infeasible and small levels of cross-zone light sharing is anticipated in the real-world case and is not anticipated to significantly impact system performance. All sensor readings are taken and control values are set at the direction of the zone manager, which acts as the information hub of the zone.

The zone level units are the agents responsible for initializing, computing, and executing the distributed control commands. At commissioning, these units are responsible for communicating with the sensing and control units to develop a record of the light level influence between each lamp and each sensor. This information is stored by the zone manager for future optimization calculations.

The zone units are responsible for periodically collecting sensor readings from all sensor units in the zone to calculate the optimal lamp settings for all lamps in the zone. They then communicate with all control units in the zone to update the artificial lighting scheme. These settings are determined using a linear or quadratic programming algorithm discussed in Chapter 4. In determining these optimal settings, the zone constructs an energy use utility curve to be shared with the building server.

The utility curve represents the explicit relationship between energy use and performance for the zone and illustrates potential for performance improvement or decline based on marginal increases or decreases in energy units allocated to the zone.

To perform the necessary computation and communication functions, the zone unit requires only a processor and a radio. If the processing capabilities of the sensing and control units are sufficient, either a sensing or control unit within the zone could be used to perform the zone manager functions as well. For the laboratory prototype, one unit is designed to perform the sensing, control, and zone manager functions.

2.2.3 BUILDING SERVER LEVEL

The building server operates as the central command post of the whole operation. While the zone managers distribute energy throughout their zones through setting lamp dimming levels, the building server allocates energy to the zones in accordance with their self-defined utility curves with a total energy allocation capped at the maximum energy use for the building-wide lighting system. This building level allocation takes place periodically and less frequently than the utility curve calculation. The timing of these building-wide updates can either be set by a timer or be event driven by requests from the zones when they observe significant changes to their needs.

The centralized control level provides relevant building performance information to the building manager and allows performance standards and energy use limits for the building system to be set. At this level, the building manager has the options to set the following building-wide parameters:

- Maximum building energy use under normal conditions
- Maximum building energy use under demand response conditions
- Minimum utility value for each zone under normal conditions
- Minimum utility value for each zone under demand response conditions

These parameters can either be stationary or can be set on a schedule.

The ability to set these parameters explicitly allows the building manager to consider how limiting energy use impacts the occupant experience. For example, the energy use restrictions may conflict with the minimum utility standard. When setting the parameter values, the building manager may choose to set conditional maximums such that more energy use is allowed if the utility of a zone drops below an absolute minimum threshold. The manager may also choose to prioritize some zones over others. For instance, it may be less impactful to the building under an energy constrained scenario to reduce the performance of a hallway or corridor zone instead of an office space. By setting the minimum performance standards by zone, the manager can determine the relative importance of the zones while still enforcing minimum standards.

By viewing the energy use as an explicit tradeoff to a defined performance metric, the incremental cost per performance level change can be evaluated. With emerging real-time pricing information systems, cost-based decision making can be integrated directly into this system. A properly integrated pricing signal would further enable the building manager to set maximum expenditures for energy throughout the day. When energy is cheap, it may be advantageous to set a very high performance standard, but in the middle of the day at peak pricing, a slightly lower standard could significantly reduce operating costs. Monitoring this data over time can give the building owners or lease holders information about the premium paid for various service quality levels.

This system also provides a secondary benefit to the building manager in that it helps in pointing out inefficiently constructed zones or zones where maintenance such as bulb replacement may be necessary. An inefficiently arranged zone will require a higher fraction of total energy use to obtain the same performance level as other zones located in similar areas of the building. Depending on the flexibility of the room layout, an alternative layout may be worthy of consideration if the current layout is performing poorly. Additionally, in looking at the progression of utility curves over time, if a zone requires increasing quantities of energy to maintain the same level of performance, normalized for weather changes, the bulbs may be nearing their end of life and require replacement.

Physically the building server constitutes a centrally located computer. It can be connected to the zone managers either through a wired or wireless connection depending on the ease of wired installation and the communication distance requirements. Wireless communication distance issues can be mitigated through use of a multi-hop network. These network design issues are considered outside the scope of this thesis. The building server for the physical test to be discussed in Chapter 3 is a Windows XP based laptop with a serial connection to a wireless modem.

2.3 System Advantages

The proposed lighting control and resource allocation system integrates occupancy sensing, daylight compensation, and occupant-selected preferences to reduce system-wide energy waste using tailored light settings selected via distributed optimization. A pervasive problem in buildings with new energy control and management systems is poor performance and lack of individual control, often leading to system overrides which reduce system performance or the complete removal of the system. The proposed system incorporates and responds to occupant preferences to give the occupants a degree of control and to prevent the poor performance issues.

While currently available technologies rely on simple feedback systems from a single sensor or a static optimization equation with a built-in weighting of performance versus energy use, the proposed system shifts the focus toward viewing energy use in terms of achievable performance. Instead of viewing energy constraint as though it is in conflict with providing quality service, the problem is reframed as determining the best performance possible with the available resources and identifying what resources are really necessary to achieve the performance goals. Through focusing on the resultant performance and representing the energy use in terms of its utility, the marginal cost, either in terms of energy or converted to dollars, of a performance improvement is identifiable. Knowledge of the performance based on resource use allows a building manager to select parameters to reduce waste and cut operating costs where the marginal cost is high for minimal incremental performance. Knowledge of the direct relationship between resource use and performance over time yields information about areas with maintenance problems or easily resolved design issues which can lead to additional long-term savings.

2.4 SUMMARY

This chapter describes the overall design of a tiered lighting control system that uses mathematical programming methods to define utility curves for energy resource use throughout a building. The system is composed of three levels, a base level of distributed sensors and control points, a middle level of processors responsible for energy optimization across a zone of the building, and a building-wide level which makes high-level decisions and enforces building-wide performance requirements and energy use limits. This system is novel in its combination of centralized and distributed control as well as in its use of optimization techniques to derive the explicit relationship between level of performance and energy use. This system is designed specifically for office building lighting applications but could be adapted for use with other building systems or building types. Given the availability of a similar characterization of energy use and performance for other systems, multiple systems within the building could be integrated into a shared energy resource allocation system where all systems compete for the available resources and are awarded these resources in accordance with their demonstrated utility.

Chapter 3.

HARDWARE AND SOFTWARE DESIGN

A laboratory prototype system is developed to demonstrate the feasibility of implementing the lighting control and optimization system. The physical implementation of the system shows that functional wireless sensing and control units can be built with widely available, off-the-shelf components and can be embedded with the necessary algorithms to implement the full resource allocation scheme. This chapter describes the hardware and software design of this system. Section 3.1 outlines the structure of the hardware system. Section 3.2 details the physical hardware construction of the wireless units, the wireless communication elements, the interface with and design of the external sensors, and the generation of the control signals. Section 3.3 discusses the software embedded in the units that enables the on-board and peripheral communications, defines the roles of the units and establishes their task sequences, and processes the data. Section 3.4 describes the role of the building server, its responsibilities, and how it interacts with the wireless units.

3.1 HARDWARE SYSTEM DESIGN

The backbone of the control and optimization system is an array of wireless sensing, computation, and actuation units which communicate with each other and with a top level server through a tiered command structure. At the base level are sensing and actuation units. These units are coordinated by zone managers which make local decisions and communicate with the building server to provide relevant information for building-wide decision making. Figure 3-1 diagrams the structure of the system.

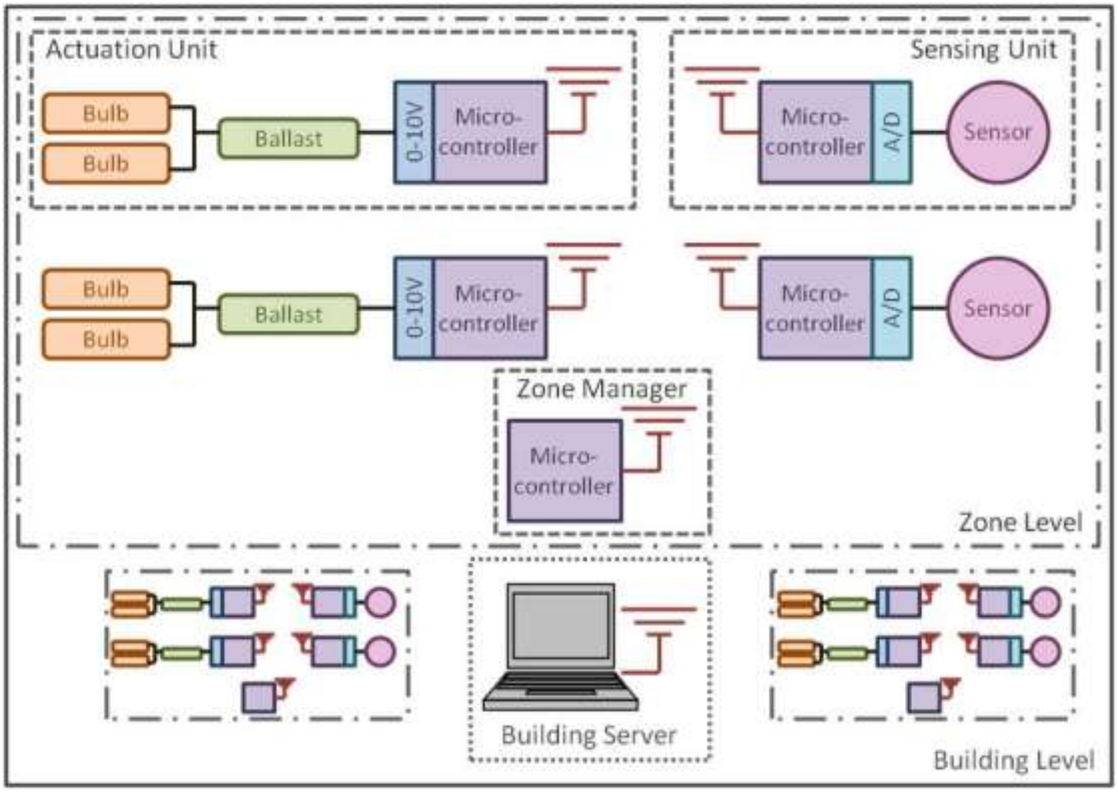

Figure 3-1: System Hardware Structure

The wireless units are endowed with sensing and actuation capabilities for use in gathering data on environmental conditions and in setting the dimming levels of fluorescent ballasts. All data and command communication between wireless units and between the units and the server is conducted via the 2.4 GHz wireless band. For the laboratory prototype system, a single type of hardware unit is designed with the ability to act as a sensing, actuation, or zone manager unit and can fulfill more than one of these roles at a time. While the units are physically identical, for the system implementation they are separately tasked as either base level units or zone manager level units. The base level units in the laboratory setup are assigned both sensing and actuation roles and the zone managers perform both of these roles in addition to the coordination and building server communication roles of the zone manager. While the same unit performs multiple roles, these roles are completed independently and the setup is therefore analogous to the separate unit system diagramed in Figure 3-1. These units, in conjunction with the building server, form the tiered optimization and resource allocation network for the lighting system.

3.2 Hardware Design of Wireless Sensing, Computation, and Actuation Unit

The prototype unit built for use in the laboratory setting is designed to measure the pertinent environmental conditions, compute necessary parameters, communicate required information with other units, and exercise dimming control of a fluorescent lighting ballast. The printed circuit board (PCB) houses the processor, radio socket, digital to analog converter with control signal socket, and four slots for sensor connections. With the sensors mounted externally, flexibility in type and quantity of sensor is preserved. The socket for the radio is designed to fit the Digi Xbee radio modules of which different, drop-in replacement versions are available. Figure 3-2 shows the functional block diagram of the hardware structure.

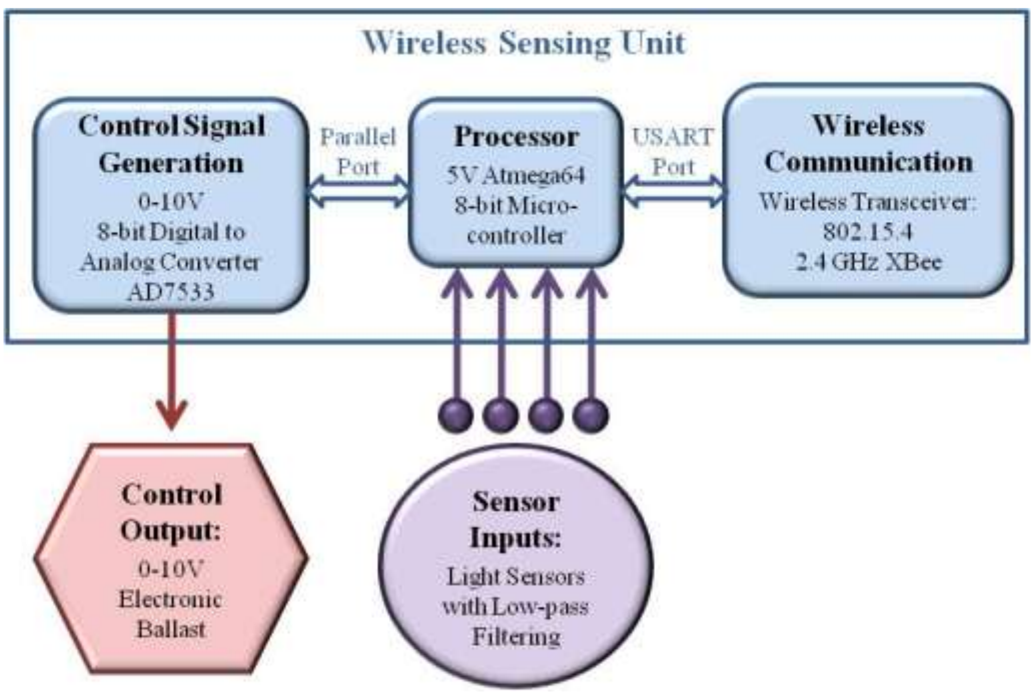

Figure 3-2: Functional Block Diagram of the Wireless Sensing, Computation and Actuation Unit

This prototype unit is specifically designed to meet the following objectives:

- Demonstrate the suitability of low-power consumption wireless protocols for building controls applications
- Accommodate a variety of sensor types to allow for the expansion of the environmental sensing capabilities and to facilitate the use of a variety of light sensors as applicable to a variety of lighting scenarios
- Establish robust peer-to-peer communications to allow for extended laboratory run times
- Store and process moderate quantities of data to support the distributed optimization scheme
- Output a consistent 0-10V analog control signal compliant with the specifications of the commercially available dimmable fluorescent ballasts

The processor, radio, control output connection, and sensor input connections are mounted on a PCB as shown in Figure 3-3 with the components labeled. Table 3-1 lists the published performance specifications of the unit components and the remainder of this section discusses the design of the unit in detail.

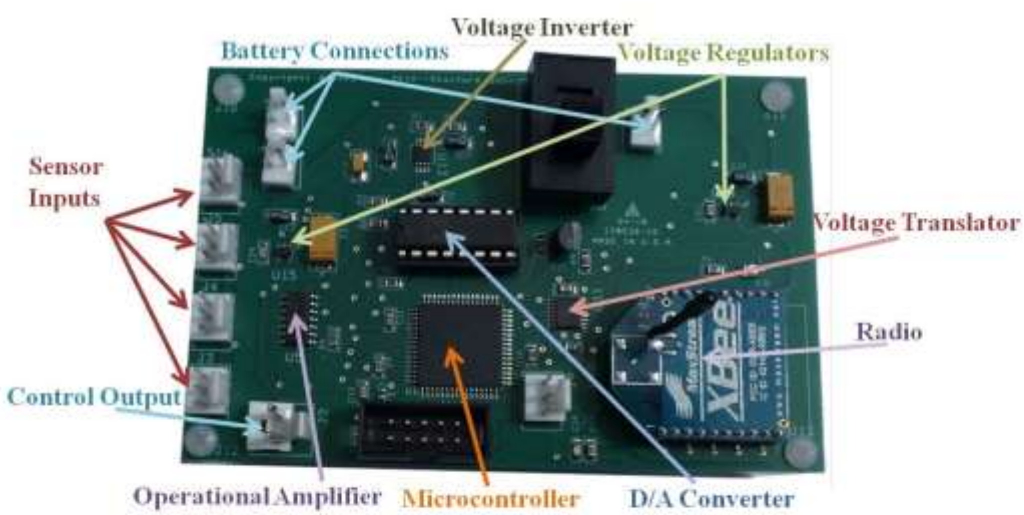

Figure 3-3: Wireless Sensing, Actuation, and Computation Unit

Table 3-1: Wireless Sensing, Computation, and Control Unit Specifications

Parameter	Specification
Processor	
Microcontroller	8-bit Reduced Instruction Set Computer (RISC) architecture
Flash Memory	64K bytes
Internal SRAM	4K bytes
Power Consumption	8.1 mA active, 2.8 mA idle
Wireless Communication	
Operating Frequency	ISM 2.400-2.4835 GHz
Communications Protocol	IEEE 802.15.4
Data Transfer Rate	9600 bps
Power Consumption	Transmit: 45 mA at 3.3V
	Receive: 50 mA at 3.3V
	Idle: < 10μA at 3.3V
Communication Range	Indoor/Urban: up to 30 m.
	Outdoor line-of-sight: up to 90 m.
Control Signal Generation	
Precision and Range	8-bit, 0-10V
Sensor Reading	
Precision and Range	10-bit, 0-5V

The above components are discussed in more detail in the following sections. The discussion includes the part numbers and manufacturers of the major board components.

3.2.1 *Wireless Unit Core: Computation and Communication*

The wireless unit is constructed as a printed circuit board whose major components include a surface mount processor and a socket for connection to a through-hole radio. All communications between units occur wirelessly through IEEE 2.4 GHz 802.15.4 standard transmissions from the radio modules. The processor is tasked with sending information to the radio for transmission with other units and with communicating with all other on-board elements. The processor further performs all digital computations and receives and converts analog sensor signals. The following subsections discuss these components in detail.

3.2.1.1 Computational Core

The core of the hardware is an 8-bit RISC Atmega64 processor. This microcontroller has 64 kilobytes of programmable flash memory and 4 kilobytes of SRAM. The processor is supplied at 5 V to maximize tolerable input signal range while remaining within specification limits. Additional features of this microcontroller package that suit it to this project include:

- 10-bit Internal Analog to Digital Converter (ADC)
- Universal Synchronous and Asynchronous serial Receiver and Transmitter Port (USART)
- In-System Programming Serial Peripheral Interface (SPI)
- 16-bit Timers and Internal Oscillator
- 53 programmable Input/Output Pins

3.2.1.2 Wireless Communication

The wireless communication for this unit is provided by the Digi XBee 802.15.4 radio. Selection of a 2.4 GHz frequency band IEEE 802.15.4 standard radio ensures international compliance with respect to frequency band usage. The 802.15.4 standard is well suited for commercial building energy efficiency applications because it was designed to minimize current consumption while maintaining robust communications at moderate distances, specified up to 30 m. indoor.

The 802.15.4 standard defines the physical and medium access control layers of the radio construction. The physical layer handles the energy and signal control of the physical RF transceiver. The medium access control layer manages the frame transmission, network time slots, and beaconing. These two layers together enable the physical transmission of the wireless signal and avoid message collisions between units.

This standard also forms the basis for the ZigBee protocol which outlines additional layers for interoperability and mesh networking. The mesh capabilities of a ZigBee network are not required for the laboratory scale prototype but might be of service in a large-scale, real-world network. As Digi produces a drop-in, interchangeable ZigBee model to the 802.15.4 model, the prototype board could easily be upgraded to incorporate a ZigBee model if additional networking functionality is desired in a future application. The IEEE 802.15.4 protocol specifies a maximum data signal rate of 250 kilobytes per second and up to 65,000 nodes using local addressing. The necessary data shared within a zone in the wireless network are contained in short packets of less than 10 bytes. The small amount of necessary data transmission per additional node in the system implies that the zones could be scaled to be very large while maintaining fast zone-wide updating. However, the radios are limited in communication range and therefore data may require multiple hops to move between an edge node and the zone manager. With consideration of the need for multiple hops, the expectation of imperfect communication, which requires packets to often be sent multiple times, and the latency both within the wireless infrastructure and within the higher level unit software, a more realistic upper bound on the number of nodes per zone is on the order of 100. Implementing a zone on this scale also requires sufficient memory and processing power at the zone manager level to compute the optimal lamp settings for a problem of this size.

The radio operates at 3.3 V in contrast to the 5 V supplied to the microcontroller. A second voltage regulator is employed to supply a constant 3.3 V to the radio and a voltage translator is used to moderate the communications between the radio and the microcontroller.

3.2.2 Control Signal Interface

The selection of the fluorescent ballast is important to the design of the lighting control interface. Unlike light-emitting diodes (LEDs) or incandescent bulbs, fluorescent lamps require very specific inputs to function properly and therefore require more sophisticated dimming hardware. With an LED, a direct current signal is simply turned on and off at a high rate with the proportion of on versus off time providing the level of dimming. An incandescent bulb can be dimmed by clipping the alternating current input for a selected width after the zero crossover points, twice per sinusoidal cycle, reducing the current through the filament; the quantity of current clipped from the signal provides the dimming in this case. By contrast, the filaments in a fluorescent tube must maintain a level of current flow sufficient to keep them hot enough to emit electrons to ensure consistent flow of gaseous electrons through the tube. Cutting the supplied signal or turning the power supply on and off does not allow for the required temperature to be maintained and will damage the bulbs over time. Dimming ballasts are designed to maintain the requisite current flow through the filaments and preserve the bulbs while simultaneously decreasing the energy use of the lamp and the light output.

Several types of dimmable fluorescent ballasts are available on the market. The dimming options are categorized as either continuous dimming or multi-way dimming. Continuous dimming allows for light output to be set to any fraction of total output. Multi-way dimming is accomplished either by implementing stepped dimming for each bulb in a fixture or by turning on and off separate bulbs in each fixture. Continuous dimming is preferable for this application because it allows for the greatest flexibility to meet the output of the optimization. While the hardware system developed for this project is tailored to continuous dimming, the overarching optimization and control system could be adapted for use with multi-level dimming.

Within the family of continuously dimmable ballasts there are several control signal options. One option is a ballast that receives a 0 to 10 V analog signal on designated input wires and varies the light output level proportionally with the applied voltage level. Ballasts that operate on power line signal controls are also available where the AC power signal is modulated to send dimming information to the ballast. This ballast type was not considered for this project because it is more appropriate for sending the same command to several ballasts at once as opposed to the desired tailored, individual control. The most recent development in ballast control signals is Digital Addressable Lighting Interface (DALI) protocol designed by the International Electrotechnical Commission (IEC), standard 60929. DALI protocol requires a two-wire interface, similar to the 0 to 10 V control signal system; however, as a digital communications protocol, it allows for individual addressing of ballasts and other devices on the network and communication of messages in addition to simple control commands. The 0 to 10 V ballast is selected for this project because the additional features provided by a DALI system are already inherent to the wireless communications system used for the system-wide communications in this project.

The remainder of this section discusses the hardware implementation of the control signal generation and the interface with the selected ballast.

3.2.2.1 Control Signal Generation

The control signal is generated using the Analog Devices AD7533 8-bit parallel input digital-to-analog converter (D/A) that is mounted on a socket on the wireless unit circuit board. The eight input pins on D/A are directly connected to port C of the microcontroller with the least significant bit tied to pin 0 of the port and the most significant bit tied to pin 7. The D/A shares the 5 V power supply with the microcontroller and therefore outputs a voltage in the 0 V to 5 V range. The National Semiconductor LM2687 voltage inverter is employed to create a -5 V signal on the board to increase the available voltage range to a total of 10 V. This -5 V signal feeds an operational amplifier (Op-Amp) quad pack, National Semiconductor LM6484, as the reference baseline. Connecting two Op-Amps in series spreads the 0 to 5 V signal from the D/A across the -5 to 5 V span. The -5 V signal is fed to the ballast as the reference level giving the ballast the requisite 0 to 10 V dimming control signal.

3.2.2.2 0-10V Analog Ballast

The Philips Advance Mark 7 ballast (IZT-2S32-SC) is used for this project because it is a high quality, 0 to 10 V continuous dimming ballast available in T8 32W sizing, a prevalent size and wattage for commercial lamps. This ballast is 120 V 60 Hz compatible which allows it to be powered directly from a standard United States wall outlet. Table 3-2 summarizes the performance specifications of the ballast.

Table 3-2: Mark 7 Ballast Specifications

Parameter	Specification
Lamp Wattage	32 W
Number of Lamps	2
Bulb Size	T8
Input Current	.56 A
Ballast Factor	.05 min to 1.00 max
Power Factor	.99
Ballast Efficacy Factor	1.49
Dimming Control Range	10% to 100%

The ballast is designed such that illuminance, control signal, and energy usage are directly proportional to each other above a minimum start up wattage and control voltage signal. Figures 3-4 and 3-5 show experimentally determined relationships between these values. From these figures it is observed that the increase in wattage and light output stops at a control voltage of approximately 7.7 V. The ballast is designed to set the lights to full output with the control leads on the ballast detached and the power consumption and light output in this condition match those measured at the 7.7 V control signal. This phenomenon is consistent across all six tested ballasts. As the three parameters of interest, namely illuminance, power consumption, and control voltage, remain linear at values below this control voltage level, the control signal computations are rescaled to account for this issue. Three options are available for accounting for the initial start up wattage and illuminance. If the voltage and corresponding illuminance are low, the plateau on the low end can be ignored without much expectation for degradation in performance. Alternatively, the ballast control signals can be restricted to greater than or equal to the minimum control value on proportional section of the curve, in a similar manner to the restriction on the upper end. If the start up energy requirement is large, the most energy efficient option would be to run the optimization problem and successively remove any ballast from the solution that has a setting below the proportional cutoff. With this method, the ballasts are not used in this relatively inefficient window of operation.

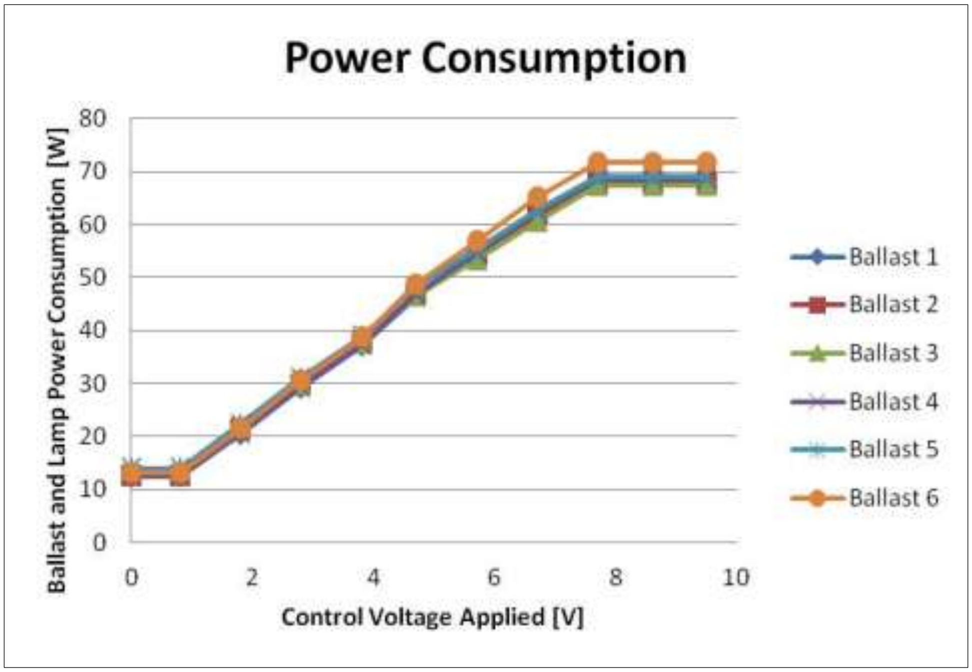

Figure 3-4: Ballast and Lamp Power Consumption

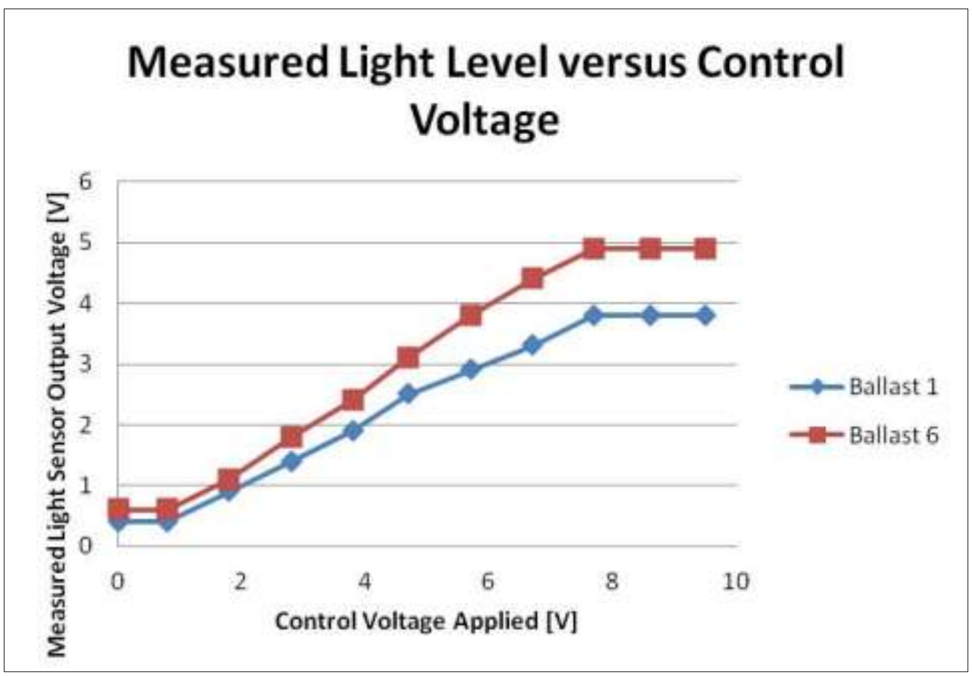

Figure 3-5: Measured Light Level versus Control Voltage

The incident light level is measured using light-to-voltage sensors selected for this project. The plot shown in Figure 3-5 represents a voltage measurement from the sensor output pin which is proportional to illuminance from the lamps. The significance of the light level measurements is in demonstrating the shape of the curve, showing the linear relationship between light level and control voltage for the range of interest, greater than the initial startup voltage and lower than the control signal saturation voltage. The plotted values are a function of sensor positioning and rotation and therefore the specific voltage values are not of interest. Power consumption of the ballast and lamps is measured using a commercially available Kill-a-Watt™ meter with reported accuracy within .2%. The control voltage is measured with a multimeter.

Most critically the above relationships indicate that varying the control voltage settings proportionally scales the incident illuminance at the sensors and the power usage of the system. The linear relationship between illuminance level and energy use by the ballast and lamps depicted in Figure 3-6 justifies the use of linear constraints in the optimization scheme and the design of the influence matrix as discussed in the Multi-Level Optimization and Resource Allocation chapter.

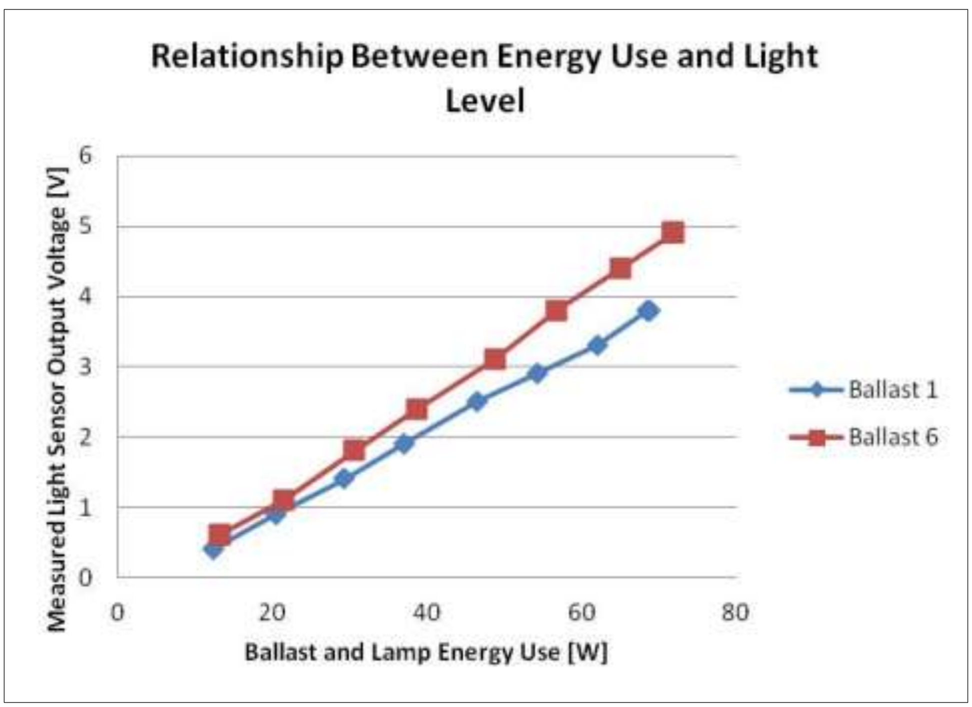

Figure 3-6: Linear relationship between Illuminance and Energy Use.

3.2.3 Sensor Design

In order to physically implement the system, the sensing and control unit requires the ability to detect ambient light levels. A light sensor with a nominal 0-5 Volt output is directly connected to the microcontroller board to be read by the Atmega64 built-in ADC. The sensor connection is a three wire connection which enables the sensor to be powered from and share ground with the microcontroller board allowing for a consistent ground voltage from which to measure the output signal.

3.2.3.1 Sensor Type

The Texas Advanced Optoelectronics Solutions (TAOS) Light-to-Voltage Converter is selected for this application. This component uses a photodiode to sense the light level and an amplifier to linearize the voltage output with respect to incident irradiance. The photodiode is sensitive over the 320 to 1050 nm range which captures the full visible light spectrum as is necessary for this application. The component package requires low supply current (1.1 mA) which aids in reducing the energy use of the control system. Additionally, the voltage input specifications for the sensor are the same as the microcontroller board requirements thus allowing it to be powered directly from the board. TAOS manufactures three products in this family; the TSL14S is selected due to its relatively low sensitivity which yields a greater sensing range prior to reaching saturation.

3.2.3.2 Signal Processing

The rise and fall times of the sensor are on the order of microseconds causing the light sensor to be particularly sensitive to noise and transitory inputs. A series resistor-capacitor (RC) low pass filter was added to the sensor output to minimize these issues. In order to achieve a large time constant relative to the rise and fall times of the sensor, a 3.3kΩ resistor and a 100μF aluminum electrolytic capacitor are used.

3.2.4 Power Consumption

The laboratory prototype unit runs on a battery pack of 5 AA batteries for setup flexibility purposes. A unit adapted for use in a building would require reconfiguration to draw power from the same power lines supplying the ballasts or from a wall outlet power near sensing locations for a long-term installation; therefore battery life is not a significant consideration in this application. However, as the goal of the system is to reduce overall energy use, consideration of energy use overhead for use of this system is necessary. Table 3-3 lists the rated typical power consumption of the board components while the board is active. As shown in the table, the energy use of the wireless unit is on the scale of a few hundred miliwatts. Dimming each light fixture saves energy on the order of watts. While active, the sum of the above component energy use is 239 mW. As an indication of the relative insignificance of the unit energy consumption compared to the potential savings, 239 mW is the amount of energy saved by dimming the lamp approximately 0.3%. Furthermore this unit was designed as a basic prototype; an upgraded unit with the same functionality could be designed with a greater focus on energy use minimization, particularly with inclusion of a more energy efficient radio. Scheduling the units to power off when not in use can also significantly reduce the current draw over time.

Table 3-3: Wireless Unit Component Typical Rated Power Consumption

Component	Part Number	Manufacturer	Power Consumption
Microcontroller	Atmega64	Atmel	27 mW (run) 10.8 µW (idle)
Radio	XBee 802.15.4 OEM	Digi	149 mW (TX) 165 mW (RX/idle) 33 µW (power-down)
Light Sensor	856-TSL14S-LF	TAOS	5.5 mW
D/A Converter	AD7523	Analog Devices	10 mW
Voltage Inverter	LM2687	National Semiconductor	2.5 mW
Operational Amplifier	LMC6484	National Semiconductor	14 mW
Voltage Translator	TXB0106	Texas Instruments	20 µW
3.3 V Regulator	MIC5219-3.3BM5	Micrel	6 mW
5 V Regulator	MIC5219-5.0BM5	Micrel	9 mW

3.3 SOFTWARE DESIGN OF WIRELESS SENSING, COMPUTATION, AND ACTUATION UNIT

The units are designed with identical hardware but are programmed with role-defined software. The majority of the units are limited to sensing, sending actuation signals, and communicating with their zone manager unit. The select zone units have the additional responsibilities of communicating with the main building server, a radio modem connected to a laptop through a USB port, and calculating the optimal settings for their zones. All units have the same capacity for accessing their peripherals and the zone units have additional processing software. The software design is discussed in detail in the following subsections.

3.3.1 SOFTWARE CONTROL OF PERIPHERALS

The peripherals for these units include the radios, the analog to digital sensor reading connections, and the digital to analog control signal output. All units are required to perform these functions.

3.3.1.1 Wireless Communication Module

The wireless communication module is the backbone of the unit-to-unit and unit-to-server communication. It provides the communication between the Xbee radio and the microcontroller.

The Atmega64 has two Universal Synchronous and Asynchronous serial Receiver and Transmitter (USART) ports. USART1 is used to communicate with the Atmel Xbee radio. The modules handling the communications protocol for this communication are adapted from those written by Wang (2007). The baud rate of both the microcontroller and the radio is adjustable and 9600 bits per second is selected as adequate for this low data transmission application.

The radio is set to run in transparent operation mode meaning that any data sent through the Atmega64 USART1 transmit (TX) pin will be added to the queue for RF transmission and any incoming RF signals are sent to the microcontroller through the USART1 receive (RX) connection. In this way the RF signals replace a direct wired serial connection. The radios can be set to either broadcast or send peer-to-peer messages and all radio settings can be changed in situ enabling the radios to change to whom they are communicating as necessary as they execute the individual tasks.

In transparent mode, no additional information is provided to the units about the origin of the messages they receive. To remedy this issue, a packet structure is designed for this project. Packets typically include a byte indicating the length of the message; however, in this case the length information is unnecessary because each possible message is a command or sensor reading of the same length with one exception, the transmission of the utility curves from zone manager to building server. For the curve transmission case, the zone manager sends a "curve will be sent" command followed by a message containing the length of the data stream that is queued up for the next transmission. Upon receiving this specific message, the building server enters a state responsible for handling the subsequent data stream. In place of a message integrity check, the function that reads in the packet information verifies the packet header, the "unit to" byte, the "unit from" byte, and that message itself is an appropriate message for the current state and the units involved. If any of these conditions are not met, the message is discarded. To avoid missing important messages, the units receiving messages send acknowledgement commands to the units from whom they receive messages. The sending unit continues to issue the command until it receives an acknowledgement confirmation. Figure 3-7 depicts the packet structure used for wireless communication throughout the tiered system.

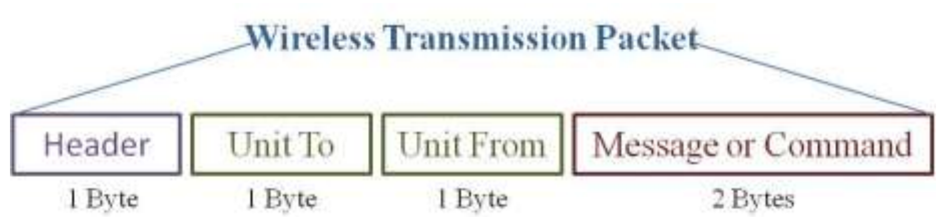

Figure 3-7: RF Packet Structure.

3.3.1.2 A/D Module and Sensor Signal Interpretation

The Atmega64 built in Analog to Digital (A/D) converter is used for taking sensor measurements. This microcontroller has an 8-channel multiplexed analog input. The A/D module first waits to ensure no other lines on the multiplex are in use and then sets the multiplexer to the channel connected to the sensor. The 10-bit A/D is read in as two separate bytes with the high byte holding the two high bit values and the low byte containing the remaining bits. The two bytes are appropriately bit shifted and combined to result in the full 10-bit value. As this is a 10-bit reading, the values range from 0 to 1023. These readings are converted into voltage levels for parsing into sunlight and fluorescent light contributions.

As illustrated in Figure 3-8, sunlight and fluorescent light have distinct spectral power distributions. Accounting for this spectral dissimilarity is crucial to the performance of a photosensor-driven lighting system (Doulos et al., 2008). This is significant as the sensitivity profile of the sensors selected for this project does not identically match the photopic curve which represents the human eye sensitivity response in normal lighting conditions. Normal light conditions here are defined in contrast to low light level conditions in which the scotopic curve is used to model the human eye sensitivity response. As this study is concerned with typical working light levels, the scotopic curve is not used. In order to compensate for the inexact matching of the light sensor to the visible light spectrum, relative factors are employed to compensate for the disparate sensitivity to the sunlight and fluorescent spectra.

In this case, the total sensor response arises from two types of contributing sources, sunlight and fluorescent light. The spectrum for each of these sources is well defined and is shown in Figure 3-8 along with the sensor response curve and photopic curve. The figure shows that the selected sensor is significantly more sensitive in the infrared (IR) range than the human eye and that sunlight irradiance has a large IR component as well. The sensitivity of the sensor to the invisible IR range means the sensor will show a relatively higher reading for solar irradiance than for fluorescent irradiance for the same level of visible light. For this reason the sensor response from each source must be decoupled from the other.

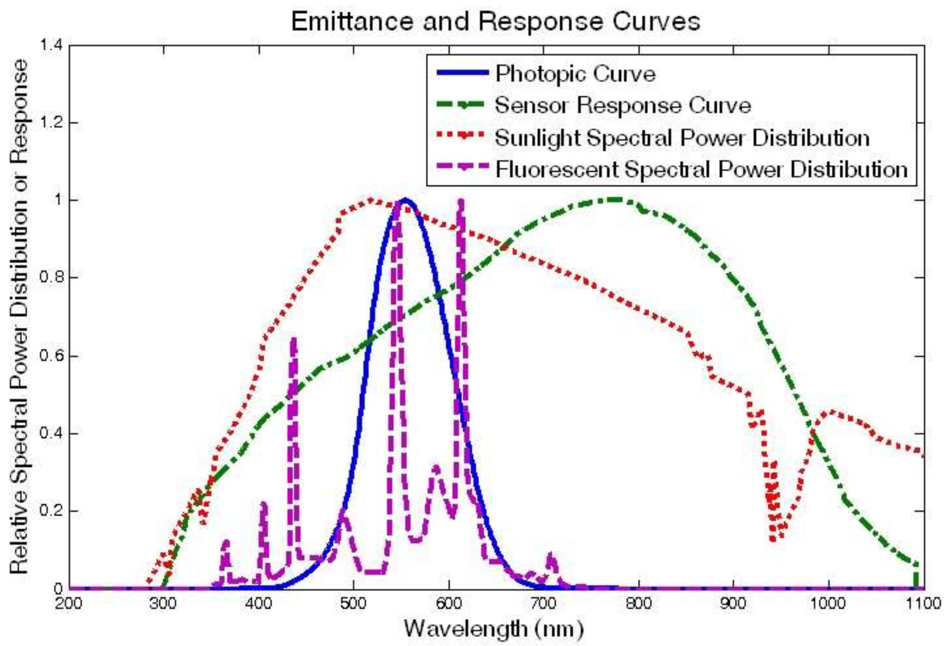

Figure 3-8: Emittance and Response Curves (CIE, 2011; TAOS, 2007; NREL, 2004; Sylvania, 2000)

The sensor response in voltage output is linear with the level of irradiance from each source as both the response curve and the distribution of radiation from both sources are constant. Because the design of the system allows for the determination of the measured incident radiation due to the artificial source, the contribution to the sensor reading from natural light can be determined and thus the total sensor response can be decomposed into artificial and natural contributions. With the sensor response decomposed thusly, the contribution of each source to the human eye response can be evaluated similarly. For a commercial system it would be preferable to select a sensor more closely tailored to the visible light spectrum or incorporate light filtering films to the design of the sensor itself to address the variable spectral distributions of reflected light; however, for use in the laboratory test bed where the largest portions of measured light are direct light components this decoupling procedure is deemed sufficient.

In order to determine the visible light contributions from the two sources, transformation factors are calculated. These factors are evaluated by comparing the numerical integral of the original emittance and spectral power distribution curves with the product of these curves and the relevant response curve. By comparing the results from the photopic curve and sensor response curve computations, the measured sensor response can be converted to the equivalent visible spectrum value. Using the response parameters for the selected sensor, the luminance value is computed using Equation 3-1.

$$\text{Visible Light [lux]} = \left(\frac{\text{visible light response per unit source}}{\text{sensor response per unit source}}\right)(\text{A/D Reading})\left(\frac{5\text{ V}}{1024}\right)\left(\frac{1\text{ W}/\text{m}^2}{1.6\text{ V}}\right)\left(683\frac{\text{lm}}{\text{W}}\right) \qquad [3\text{-}1]$$

3.3.1.3 D/A Module

The hardware design for the digital-to-analog (D/A) converter enables direct and robust software implementation. The eight pins on the Atmega64 port C are connected in parallel with the eight input signal pins on the D/A allowing the port C pins positions, when initialized as digital output pins, to function as the bits of the 8-bit control signal. The desired light level fraction is converted to a fraction of the maximum value represented by the eight bits, 255, and this value is written to port C. Figure 3-9 shows the linear relationship between values written to port C and the measured control voltage output.

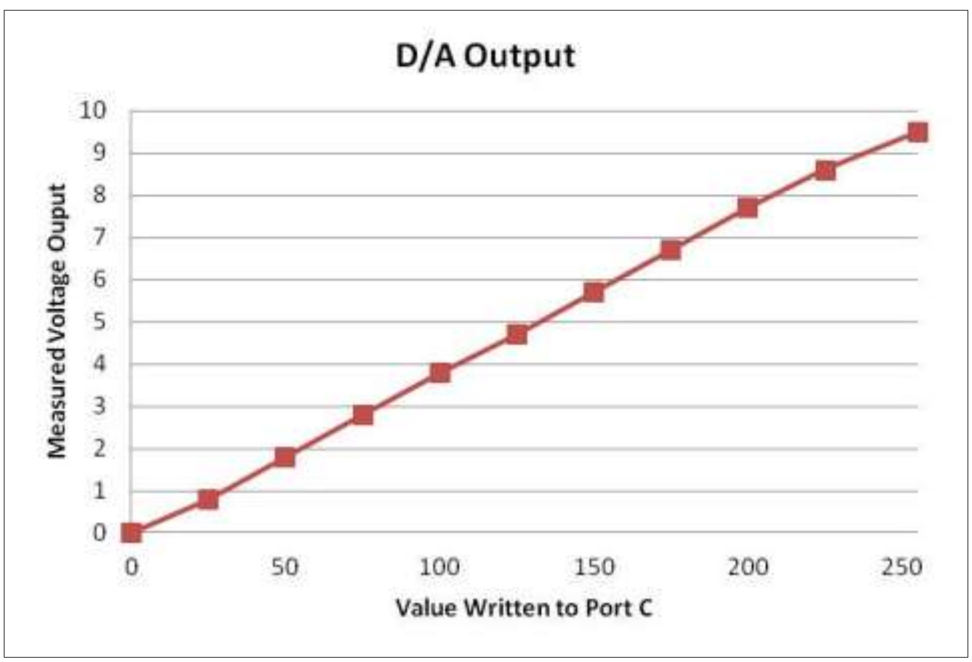

Figure 3-9: Digital to Analog Module Output

3.3.2 PROCESSING MODULES

As noted previously in this chapter, the units are assigned one of two designations. They are either sensing, control, and communication units or they have the additional responsibility of processing data and determining the appropriate light level settings for their respective zones using an embedded linear programming algorithm. In the following sections, the state machines that define these roles and the methods used to compute settings for the zones are explained in more detail.

3.3.2.1 State Machine Design: Sensing, Control, and Communication Unit (Basic Unit)

Figure 3-10 shows all states required for this unit type. In the *Initialize* state all ports on the unit are appropriately initialized for their specific tasks and the unit has been informed of and has acknowledged the start of the exercise. The unit then spends the majority of its time in the Wait state and responds only to command requests from its zone leader. After acknowledging the request and executing the appropriate task, it returns to the *Wait* state to wait for the next command.

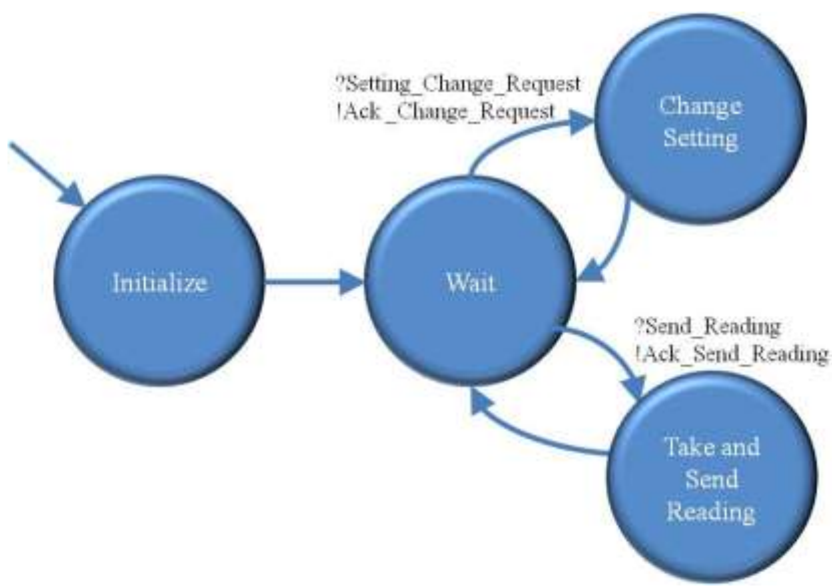

Figure 3-10: State Machine for Sensing, Control and Communication Unit

3.3.2.2 State Machine Design: Zone Manager Unit

The zone unit has the responsibility of managing the zone in addition to serving as a sensing and actuation unit and therefore has more states than the basic unit. The basic unit functions are implemented on a separate state machine path to parallel a scenario where the zone manager is a separate unit. The zone manager states are shown in Figure 3-11 a definition of the states follows. Figure 3-12 shows the sub-states within the *Initialize Influence Matrix* state to illustrate the zone initialization procedure.

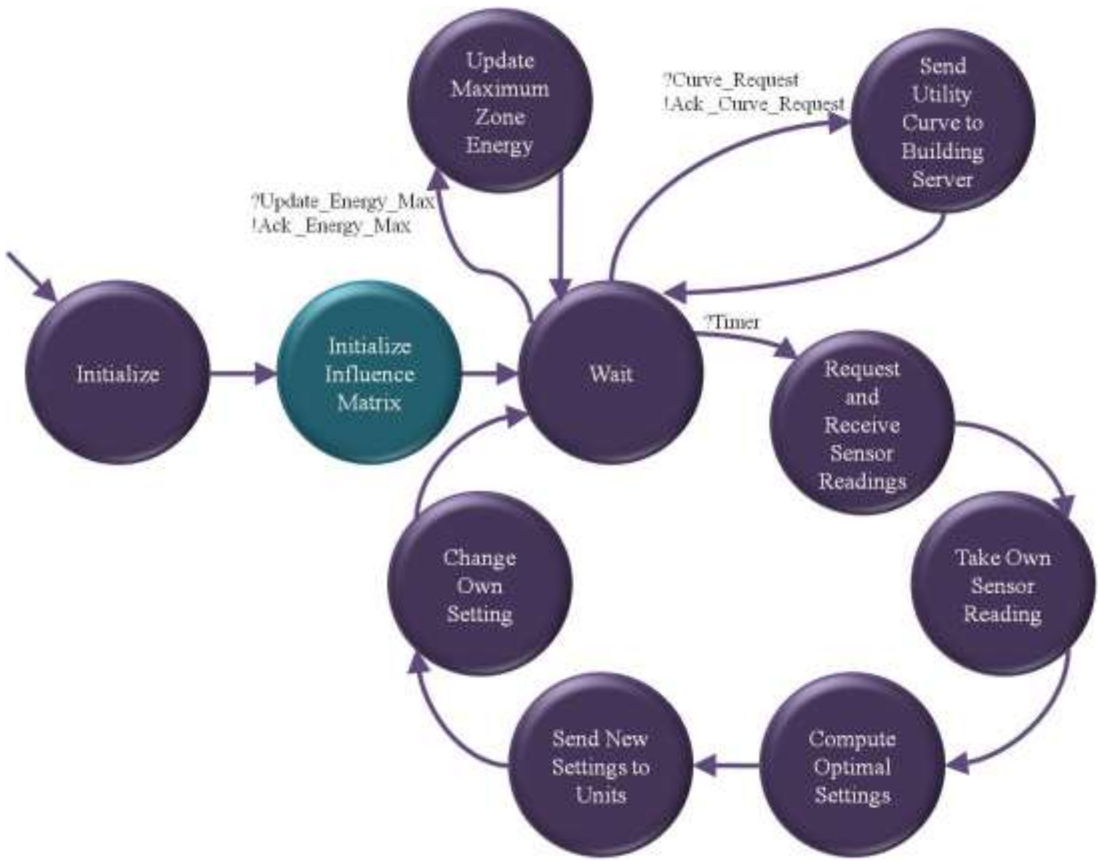

Figure 3-11: Zone Unit State Machine

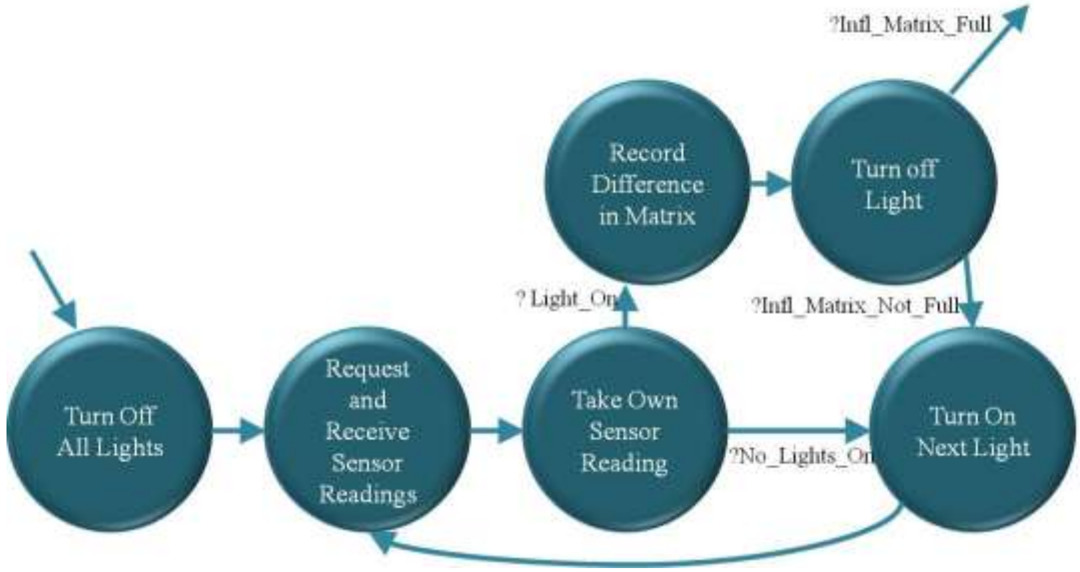

Figure 3-12: State Machine Description of Initialize Influence Matrix State from Zone Manager State Machine

Initialize: The unit sets all internal registers and ports appropriately and is informed of and has acknowledged the start of the exercise.

Initialize Influence Matrix: The influence matrix is set up to represent the contribution of light from each lamp in the zone at full output to each sensor. This process first requires that the lamps are all turned fully off. A measurement is then taken from each sensor and used as a baseline background light level. The lamps are then turned on individually and a sensor reading is taken from each unit. The difference between the off measurement and the on measurement is recorded as the entry in the influence matrix. The state machine for this sub-process is shown in Figure 3-12.

Wait: This is the default state of the microcontroller. This is where the microcontroller passes time between requests from the building server and updating the utility curve with information from the sensors in the zone. This state is exited only when a curve or energy update request is received from the building server or the internal timer signals that it is time for the next curve update.

Send Utility Curve to Building Server: The computed utility curve is transmitted to the Building Server.

Update Maximum Zone Energy: The zone updates the maximum energy it is allowed to use. This information is used to select the final settings for the lights in the zone.

Request and Receive Sensor Readings: The microcontroller requests and receives a sensor reading from each unit in its zone. The communications are acknowledged on both ends to ensure all sensor readings are appropriately updated.

Take Own Sensor Reading: In the laboratory setup, the zone manager additionally functions as a sensing and control unit and takes its own reading for inclusion in the optimization calculation.

Compute Optimal Settings: The microcontroller uses a mathematical programming algorithm to compute the optimal settings for the lights in its zone.

Send New Settings to Units: Each unit in the zone is sent its new optimal light setting. The units acknowledge the new setting commands to ensure the commands are executed. The zone manager continues to issue setting change commands until they are properly acknowledged.

Change Own Setting: Because the zone manager also serves as a sensing and control unit, it must update its own control voltage setting.

3.3.2.3 Linear Programming and Lamp Settings Selection Module

The optimization algorithm implemented on the units is a linear programming algorithm utilizing the Simplex Method. The linear program is set up to minimize the difference between the simulated requested light level and the actual level of light seen at each sensor constrained by a maximum energy level. This maximum energy level is varied from one tenth of the maximum possible energy use in the zone to the maximum allowable zone energy level as set by the building server. The sum of the error in meeting the requested light level is stored as the utility for each energy value.

The problem formulation as embedded on the microcontrollers is designed to minimize the sum of deficit in light level provided to all occupants. The Simplex Method (Dantzig, 1951) is selected for this implementation because the simplicity of the row and column operations is well suited to being embedded and executed on a microcontroller with limited storage space and processing power. The case where excess light is not penalized is exclusively considered in this execution as the goal of the laboratory setup is solely to demonstrate the feasibility of the physical implementation of this type of sensing and control system. Discussion and comparison of alternative formulations is provided in Chapter 4.

3.4 SOFTWARE DESIGN OF TOP LEVEL BUILDING SERVER

The building level server is responsible for interacting with the zone managers to set the maximum zone energy use for all zones. A laptop with a serial connection to a wireless transceiver is used as the building server for the laboratory setup. The laptop runs a compiled executable originally coded in C that coordinates the start of the system, requests and receives utility curve data from the zones, and allocates energy resources in accordance with the utility curves. The building server additionally stores system performance and energy use data.

3.4.1 COMMUNICATIONS MODULE

As previously noted, the building server communicates through a wireless radio with a serial connection. A serial development board is available for the Digi XBee radios and is used to transmit wireless messages to other units from the laptop server. The module enabling the serial communications is adapted from Lynch (2002). The code is written for use on a Linux system requiring the use of a Linux development environment program for use with a Windows operating system. The freely available Cygwin software is used for this purpose. The Cygwin download includes the C-language compiler required to form the executables used in this project.

3.4.2 PROCESSING MODULES

Responsibilities of the building server include communicating with the zone manager nodes and allocating energy. The following subsections illustrate and discuss the contribution of this component to the system and how it interacts with the wireless units. The state machine design for the building server is composed of six states with the majority of time spent in the *Wait* state. Figure 3-13 shows the state machine diagram and the states are summarized as follows.

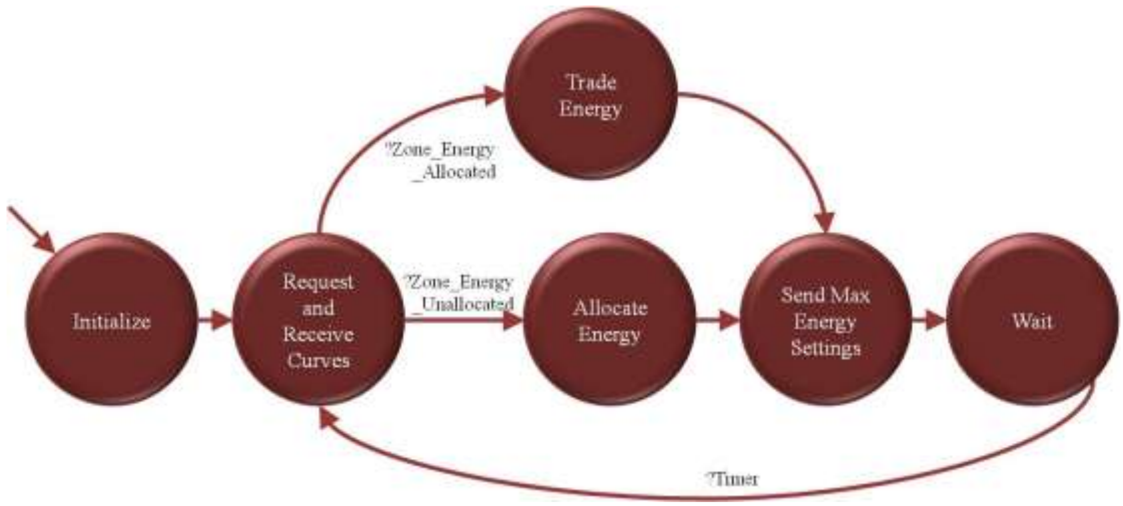

Figure 3-13: Building Server State Machine

Initialize: The building server sets up the serial port for communication and notifies all units that a test is about to begin. After all units have acknowledged the start of the test, the building server proceeds to the next state.

Request and Receive Curves: At the outset the building server has no information about the zones. It waits in the *Request and Receive Curves* state until the zones collect the sensor information from their sub-units, build their respective influence matrices and create their first utility curves. After this information is received, it moves on to the *Allocate Energy* state. As the curves are consistently updated based on varying conditions in the building, updates to the allocation are necessary. Subsequent visits to the *Request and Receive Curves* state are followed by the *Trade Energy* state which assumes energy resources have previously been allocated.

Allocate Energy: The building server uses the preset maximum building energy use and the zone-provided utility cures to allocate energy units to the zones.

Send Max Energy Settings: Maximum energy settings are sent to all zones via the wireless transceiver.

Wait: This state is where the majority of time is spent. The building server exits this state on a schedule controlled by a timer to update the energy allocation.

Trade Energy: Using the updated utility curves, maximum building energy use, and the previous zone energy allocations, the building server allows zones to trade energy units for the next allocation cycle.

For the purposes of the experimental design, an energy unit is defined as equivalent to 10% of fully-on energy for a single lamp, approximately 6.4 Watts, for the building-wide allocation. The zone curves are also constructed at this increment for continuity. Because the zones communicate their utility curves as a vector of utility values, it is important that these two levels are consistent in their definitions of energy unit. The initial resource allocation is computed using a modified greedy algorithm with units of energy assigned to zones with the highest marginal utility. Leftover energy units are distributed evenly to allow zones to update themselves as necessary between building-wide updates. In subsequent rounds, energy units are traded again on the basis of marginal utility increases and decreases. The details of the energy unit allocation are discussed in Chapter 4.

3.5 SUMMARY

A hardware and software system has been designed to physically implement the tiered optimization and resource allocation scheme. The hardware system consists of three levels of components: a building server, a zone manager, and a sensing and actuation unit. All communications between components are conducted wirelessly using 802.15.4 2.4 GHz radios. The building server is composed of a laptop with a serial connection to a wireless radio. The zone manager and sensing and actuation units are physically identical and are comprised of a printed circuit board hosting an Atmega64 microprocessor, a radio socket, connections for external sensors, and a 0-10 V control signal output. While these two types of units are physically identical, they have distinct roles within the system and are embedded with different software to support these roles. The sensing and actuation unit responds to zone manager commands for changes to the control signal and takes sensor readings on request. The zone manager functions as a sensing and actuation unit for its zone in addition to its management role. The management role includes the additional responsibilities of computing the optimal light settings for the zone and responding to requests from the building server for the zone utility curve. The design of the hardware and software system is conducted to demonstrate the feasibility of constructing a working system with readily available hardware components.

Chapter 4.

OPTIMIZATION AND RESOURCE ALLOCATION

Optimization and resource allocation methods form the basis of the system decision making. The zone level uses mathematical programming algorithms to determine the optimal light settings at each lamp for all possible energy use levels in order to construct a zone energy use utility curve. The building level uses a greedy algorithm to make initial resource assignments to zones and subsequently implements a trading strategy to allow the zones to trade for resources as the state of the building changes throughout the day. This chapter discusses the specific algorithms used to develop the zone utility curves and to allocate building-wide energy based on these curves.

4.1 ZONE-LEVEL OPTIMIZATION

Optimization at the zone level has two purposes. The first is to assign the optimal lamp settings to each lamp in the zone considering both energy use and occupant preferences. The second purpose is to create an energy use utility curve for the zone to be shared with the building server. In order to accomplish both of these tasks, the zone manager calculates the optimal light settings given the current sensor measurements for a range of energy use. A metric of the error in meeting the preferences for each of these energy use values is defined as the zone utility. The discrete function of these utility values versus their respective energy consumption values defines the zone utility curve.

Several formulations of the optimization problem were considered. The design of the optimization problem informs the utility metric as each formulation of the problem is designed with a specific goal. A description of these options and their advantages and disadvantages follows. For each of these options, one of the constraints in the program is the maximum energy use for the zone, varied as discussed above. This constraint is selected as less than or equal constraint, as opposed to a strictly equal constraint, to ensure that the resultant curves are monotonically increasing for use in the building-level resource allocation.

4.1.1 One-Sided Linear Programming

The one-sided linear programming problem is the most straight-forward case. This case is developed with the goal of ensuring that all occupants receive at least the amount of light they request. No penalty is assigned for excess light. The linear program can be structured as follows.

Minimize:

$$\sum_{i=1}^{M} p_i \epsilon_i$$

subject to:

$$\boldsymbol{A}\boldsymbol{x} + \boldsymbol{\epsilon} \geq \boldsymbol{b}$$

$$\sum_{j=1}^{N} x_j \leq E$$

$$0 \leq \boldsymbol{x} \leq 1, 0 \leq \boldsymbol{\epsilon}$$

where:

$\boldsymbol{A}$ has dimensions M by N and is the influence matrix capturing the illuminance from each lamp to each sensor for a fully-on setting; each row represents one of M sensors and each column represents one of N lamps

b is a vector of the artificial light levels required to meet the occupant-specified target levels, where $b_i \geq 0$

x is a vector of the fractional settings of all lamps in the zone

ε is a vector of the error in meeting the occupant-specified light level for all sensors

p is a vector of assigned participation weighting factors to be used for two purposes. First, p_i is set to zero if the respective occupant is not present to ignore the relevant constraint. Second, p_i can be used to rank the importance of the sensor locations within the zone if desired.

E is the maximum level of energy (in terms of fractional light settings) allowed for the particular iteration

M is the number of sensors in the zone

N is the number of lamps, or individual control points, in the zone

To form the utility curve, the above program is solved for values of *E* up to the fully-on energy level for the zone or until **ε** reaches zero, whichever comes first. The shape of the curve shows how precipitously the performance of the system declines under energy use restriction.

The influence matrix ***A*** is formed for each zone during building commissioning and is stored by the zone manager. To take into account the deterioration in the system over time, a routine commissioning schedule to redefine ***A*** is recommended for long-term installations. Re-running the commissioning sequence is also required for any major changes to the layout of the room or movement of sensors or lamps. The influence matrix ***A*** is determined through the following steps:

1. Turn off all lamps in the zone
2. Turn on one lamp
3. Record sensor readings for all zone sensors
4. Turn off the lamp

5. Repeat steps 2 through 4 for all remaining lamps in the zone
6. Record sensor readings for all zone sensors
7. Subtract the values recorded in step 6 from all other readings, sensor by sensor.
8. Convert the remaining values to illuminance in units of lux and record in $\boldsymbol{A}$

The light level vector $\boldsymbol{b}$ is computed by taking the current voltage level readings from the sensors, subtracting the expected voltage level due to the current light settings, transforming the remaining voltage values into equivalent sunlight illuminance, and subtracting the resultant sunlight illuminance from the occupant-specified light preference level. If more sunlight is provided than is desired, the computed value for that element of vector $\boldsymbol{b}$ would be negative; however, negative illuminance cannot be provided by the system and as such these values for b are set to zero to better represent the physical system behavior and to allow the zone manager to stop running iterations once the value of the objective function reaches zero. Because the objective function is a sum of the error terms, the solution to the program is unchanged by this alteration.

The vector $\boldsymbol{x}$ contains the dimming fractions for all lamps in the zone. The values in vector $\boldsymbol{x}$ for the iteration with the maximum zone energy allowed, as established by the building server, are used as the next lamp settings unless the value of the objective function reaches zero at a lower energy level in which case the settings for that level of energy use are used.

To make use of the information regarding occupancy, the corresponding row of matrix $\boldsymbol{A}$ and vector $\boldsymbol{b}$ is removed for any non-present occupant. If a minimum light level were preferred in the absence of particular occupants, this minimum level could be substituted for the relevant entry in vector $\boldsymbol{b}$ while the corresponding row in matrix $\boldsymbol{A}$ would remain intact.

One advantage of this formulation is that it requires the fewest number of constraints and variables, which minimizes the computational requirements. This program also ensures that meeting the necessary light levels for the occupants to perform their required tasks is the first, and only, goal as energy use is increased. The main drawback to using this construction is that it may lead to some occupants receiving too much light as excessive provision is not considered. This formulation is selected for implementation for the physical test system described in Chapters 3 and 5 despite its disadvantages because its singular goal of minimally meeting request levels allows for clear visual interpretation of the results. As the physical test is designed to demonstrate the feasibility of the implementation of this type of overall system and to show the general performance capabilities of this system, the one-sided linear programming construction is best suited.

The utility metric used for this formulation is the negative sum of deficit error in meeting the requested light levels. This is the value of the objective function for a specified value of *E*. The Simplex Method (Dantzig, 1951) is employed for embedded processing of the linear programming problem. This method is ideal for embedded systems because of the minimal memory requirements and the simplicity of the calculations.

4.1.2 TWO-SIDED LINEAR PROGRAMMING

In an effort to address the issue of providing too much light to some occupants, a two-sided program has been developed and is stated below. This program is very similar to the one-sided version with the exception that the objective function incorporates both positive and negative errors in meeting the requested demand.

Minimize:

$$\sum_{i=1}^{M} p_i(\epsilon_{1i} + w\epsilon_{2i})$$

subject to:

$$\boldsymbol{Ax} + \boldsymbol{\epsilon_1} - \boldsymbol{\epsilon_2} = \boldsymbol{b}$$

$$\sum_{j=1}^{N} x_j \leq E$$

$$0 \leq \boldsymbol{x} \leq 1, 0 \leq \boldsymbol{\epsilon_1}, 0 \leq \boldsymbol{\epsilon_2}$$

where:

A, b, ***x, p,*** *E*, *M*, and *N* are as defined in Section 4.1.1

$\boldsymbol{\varepsilon_1}$ is a vector of length *M* of the deficit in meeting the occupant-specified light level for all sensors

$\boldsymbol{\varepsilon_2}$ is a vector of length *M* of the surplus in meeting the occupant-specified light level for all sensors

w is a constant weighting factor that allows for specification as to the relative importance of surplus and deficit provision of light; this value is set to one for equal consideration of these concerns and typically would be set to a value between 0 and 1; *w* equal to zero is a special case representing the one-side program discussed in the previous section

This construction is an improvement over the one-sided case in that a penalty for too much light is incorporated directly into the program, thereby truly tailoring the light settings to the preferences of the occupants. However, use of this version may contribute to higher energy use in the zone because the solution is no longer singularly focused on providing at least a minimum light level and is instead required to comply with a potentially competing request of minimally exceeding the light level at all sensors as well. This version is not guaranteed to reach an objective function value of zero as *E* is increased even if the installed lighting capacity is sufficient to minimally meet the requests of the occupants and is therefore reliant on determining the location of the start of the performance plateau in the utility curve to select the appropriate lamp settings for the zone.

In this case, utility is defined as the total sum of positive and negative error, respectively weighted according to *w*. This new utility definition is reflective of the change in design of the objective function and the goal of minimizing error from both sides. As *w* is typically selected to be less than or equal to 1, the possible utility function values range from the negative sum of light levels requested in the zone to zero.

As with the one-sided linear formulation, the two-sided linear formulation is similarly well suited to solution via the Simplex Method. A comparison of the difference in performance of the two linear constructions is presented at the end of this section.

4.1.3 Two-Sided Quadratic Programming

With the linear programming algorithms, all incremental error in meeting the requested light levels is treated equally, meaning those sensors that are close to meeting their demand have equal opportunity to resources as those that are much further away. To refocus the resources toward those with the highest need, a squared sum of error is employed in the objective function. The program for this case is a linearly constrained least-squares problem and is presented below.

Minimize:

$$\sum_{i=1}^{M} p_i({\epsilon_{1i}}^2 + w{\epsilon_{2i}}^2)$$

subject to:

$$\boldsymbol{Ax} + \boldsymbol{\epsilon_1} - \boldsymbol{\epsilon_2} = \boldsymbol{b}$$

$$\sum_{j=1}^{N} x_j \leq E$$

$$0 \leq \boldsymbol{x} \leq 1, 0 \leq \boldsymbol{\epsilon_1}, 0 \leq \boldsymbol{\epsilon_2}$$

where:

$\boldsymbol{A}$, $\boldsymbol{b}$, $\boldsymbol{x}$, $\boldsymbol{p}$, $\boldsymbol{\varepsilon_1}$, $\boldsymbol{\varepsilon_2}$, E, M, and N are as defined in Section 4.1.2

w is a constant weighting factor that allows specification as to the relative importance of surplus and deficit provision of light; this value is set to one for equal consideration of these concerns

The constraint equations in this version are identical to those in the two-sided linear program with the exception that $\boldsymbol{b}$ must not be zeroed out for the cases where excess natural light is already supplied. Changing these values to zero would limit the interpreted significance of the excess supply of light and inadequately incorporate the desire for no additional light in these spaces.

The most significant change is in the objective function. By summing the square of the error, the sensors that are furthest from meeting their desired light levels are more heavily weighted than those that are closer to reaching their goals. This is a preferred method for assigning resources as it encourages a more even distribution of quality performance for all occupants. The drawbacks, however, are that the solution to this program is significantly more computationally intensive for embedding on the microcontrollers and, similarly to the two-sided case, performance plateaus must be identified to select the lamp settings for the zones.

The utility metric defined for this case uses the negative of the objective function values for each specified value of E because this is again a minimization problem. These values are bounded on the lower end by the negative sum of the squares of all requested light levels in the room and by zero on the upper end.

Because this is no longer a linear programming problem, a different solution method is required. This problem has linear constraints and a quadratic objective function of the form:

$$\boldsymbol{v}^T \boldsymbol{Q} \boldsymbol{v}$$

where:

$\boldsymbol{v}$ is a vector of all variables in the problem including $\boldsymbol{x}$, $\boldsymbol{\varepsilon_1}$ and $\boldsymbol{\varepsilon_2}$

$\boldsymbol{Q}$ is a diagonal matrix with diagonal elements equal to 0, 1, and w

As a diagonal matrix with no negative elements, Q is positive semi-definite which indicates the convexity of the objective function. This guarantees that the solution to the minimization problem reaches a minimal objective function value. A quadratic programming problem of this form can be solved by many methods such as interior point, active set, gradient projection, or a modified simplex algorithm. The modified simplex algorithm (Wolfe, 1959) is well suited to implementation on an embedded system because of the relative simplicity of the computations. As with the traditional Simplex method, the calculations are row and column matrix operations which are easily performed by a low-level processor. However, this implementation requires significantly more memory space and computation time than the original Simplex algorithm because the number of modified constraint equations and problem variables are increased.

The results from the quadratic program are presented alongside the two linear programs in Section 4.1.4 for comparison. For this comparison, the built-in MATLAB® quadratic programming solver is used to find solutions to the above quadratic program.

The scaling limitations of the zone with respect to the hardware system were discussed briefly in Chapter 3. On the optimization software side, there is no explicit limit with respect to an ultimate size of the zone, but the larger a zone is, the larger the processor and memory space to conduct the optimization calculations are required to be. Therefore selecting an appropriately sized microcontroller for the desired zone size is a critical hardware design decision for the development of the zone manager hardware.

4.1.4 Illustration and Comparison of Optimization Formulations

In order to compare the three different formulations, the lamp settings calculated, resultant provided light levels, and resultant utility curves are plotted for a sample zone. The number of energy units allocated at the low edge of the plateau in the plot represents the minimum amount of energy required to provide the lighting scene that best meets the preferences of the occupants. Any additional allocated energy would not be used by the zone as the use of more energy would not improve and may actually degrade performance in the two-sided formulations. Parameters used in this illustration include:

- 5 occupants/sensors in the zone
- 5 lamps in the zone
- Energy units equivalent to .1 of the energy use of a fully-on lamp

- Requested illuminance level of 500 lux for each sensor
- Weighting factor *w* between over and under-supply of light equal to 1
- All participation weighting factors p_i equal to 1

For the purposes of this demonstration, the influence matrices (***A***) of the simulated zones are populated with arbitrary values to give an indication of the general performance of the optimization system in a diverse array of spatial setups but are not calibrated to any specific physical setup.

Figure 4-1 shows how the light level at each sensor is affected by the number of energy units used by the zone under lighting allocation assignments provided by each of the three optimization programs. The sum of the artificial light contributions for Sensor 1, as represented by the respective row in the influence matrix (***A***) is less than the target 500 lux. In this case, the sum of this row in the influence matrix is 465 lux which is the value of the line shown in Figure 4-1 representing Sensor 1 at the far right edge of all three plots. As the energy use in the zone is increased, the light level for this sensor moves closer to the target value, but only at the expense of forcing the other sensors higher than their requested level.

A much wider range in light levels exists in the one-sided linear plot than in the two-sided ones. The one-sided linear formulation allows sensors 3 and 4 to well exceed the desired light level early on in the allocation. The two-sided linear construction prevents this from happening. The two-sided quadratic figure keeps the zones most tightly together as the energy level is increased. This behavior occurs because of the squaring of the error terms which heavily weights those zones furthest from the goal illuminance of 500 lux. The corresponding lamp settings are shown in Figure 4-2, which shows how widely the settings vary for different zone energy use levels depending on which optimization formulation is selected.

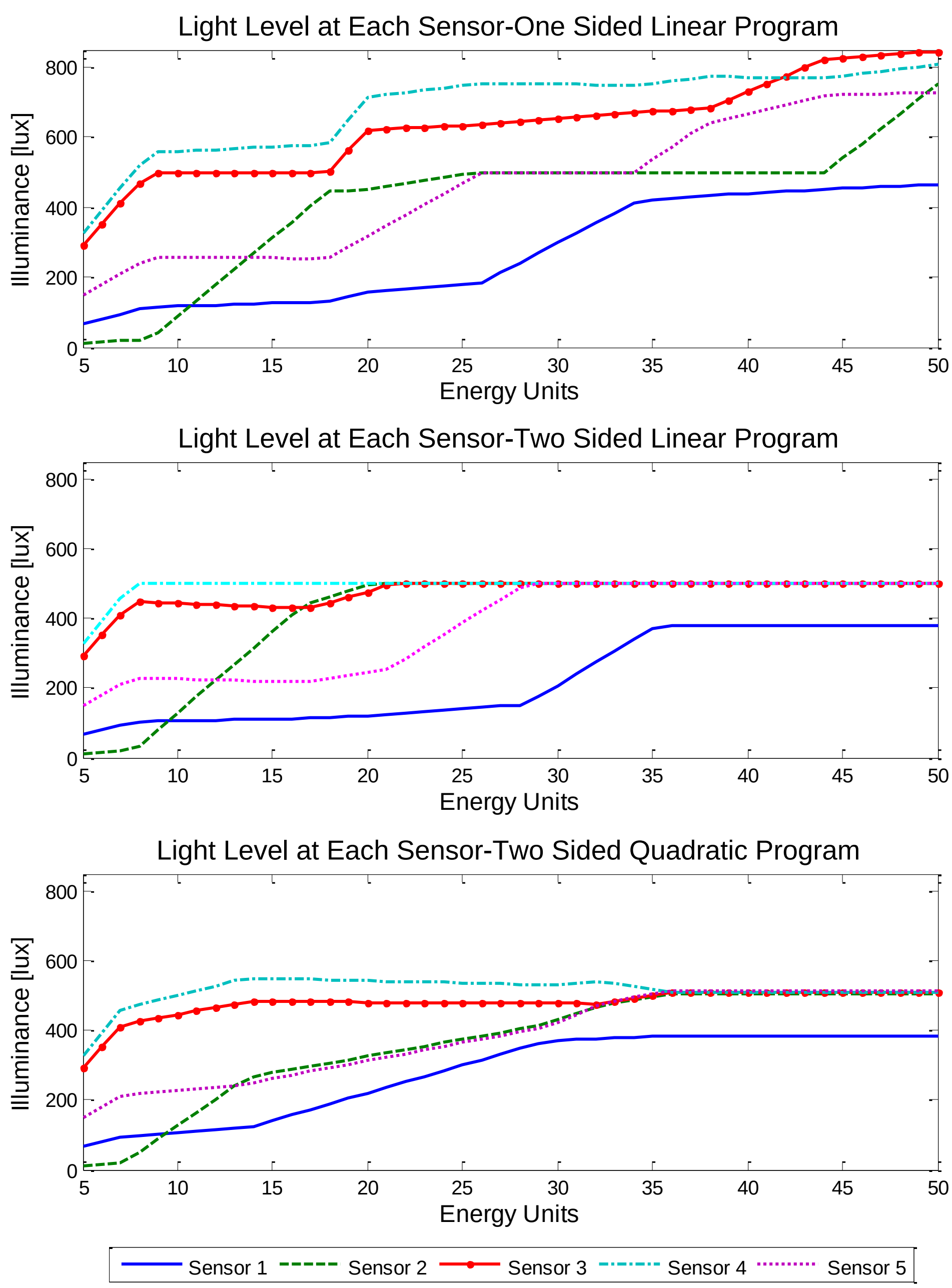

Figure 4-1: Comparison of Optimization Methods: Light Level versus Energy Use

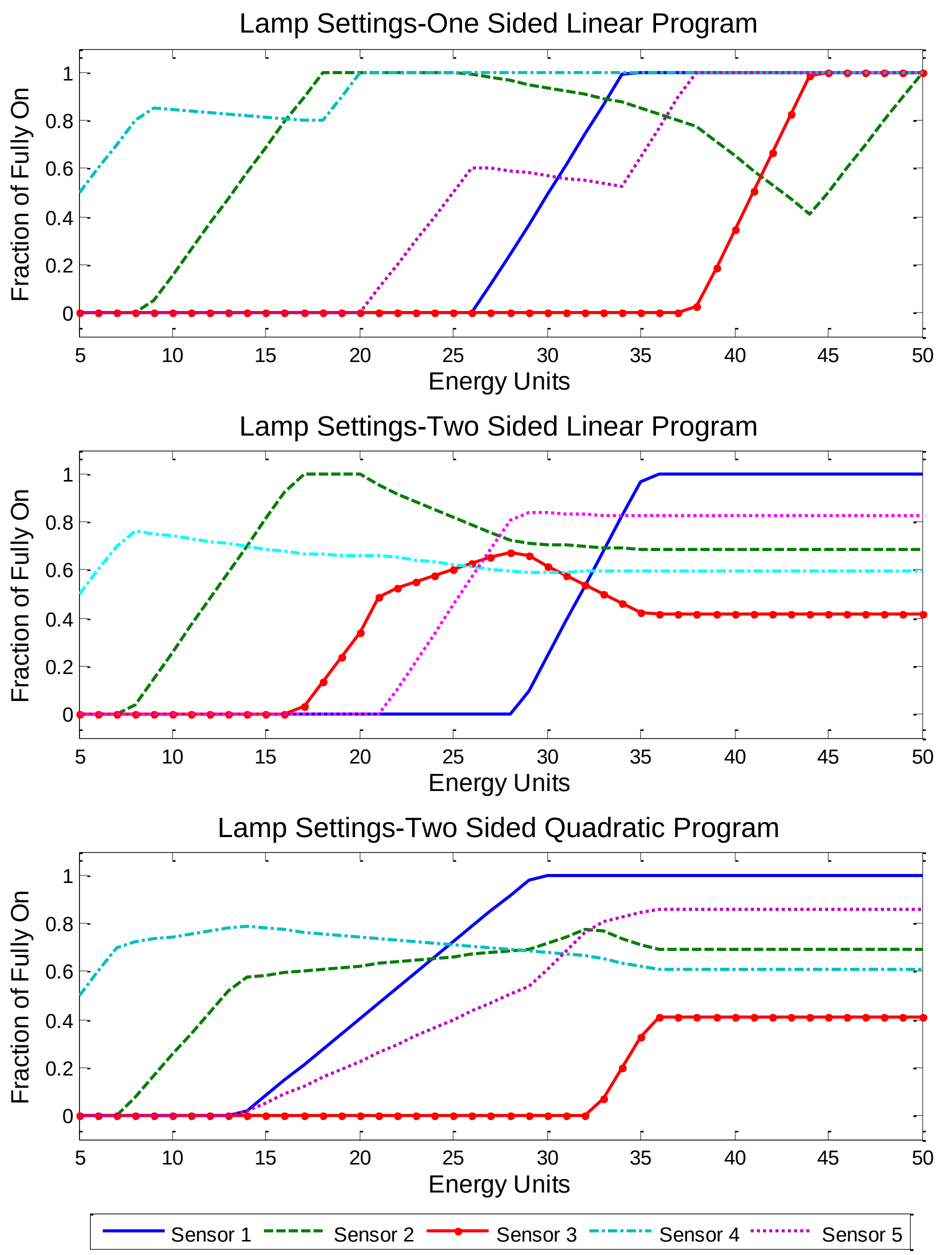

Figure 4-2: Comparison of Optimization Methods: Lamp Settings versus Energy Use

Figures 4-3 through 4-5 show how each optimization scheme performs under the three different utility metrics. As is necessary according to the relationship between the utility metrics and their corresponding objective functions, each formulation performs the best by its own metric and sets the standard by which the other two are compared.

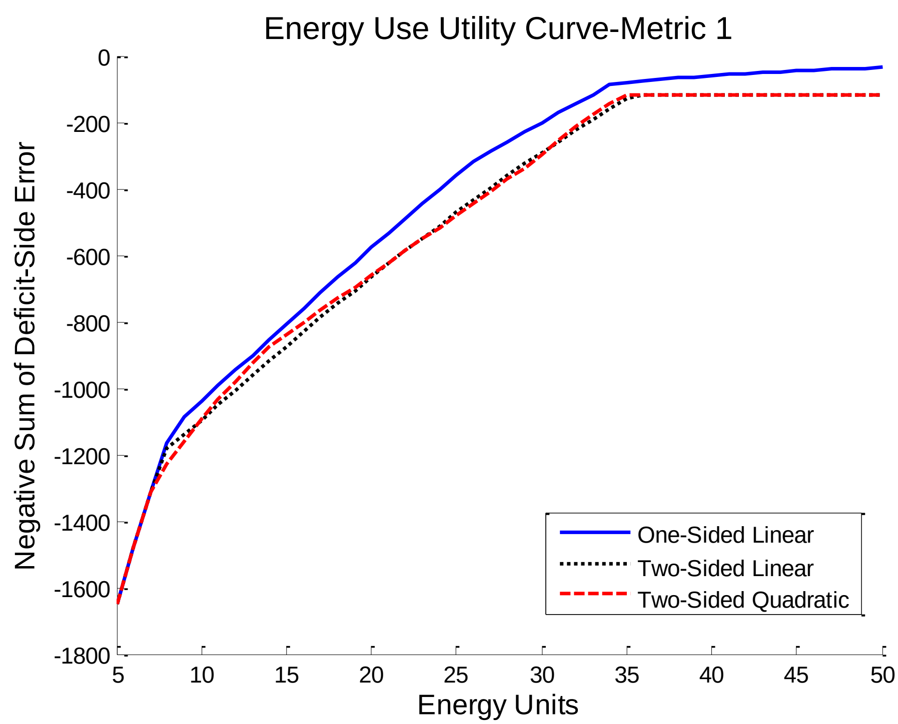

Figure 4-3: Example Energy Use Utility Curve (Negative Sum of Deficit-Side Error)

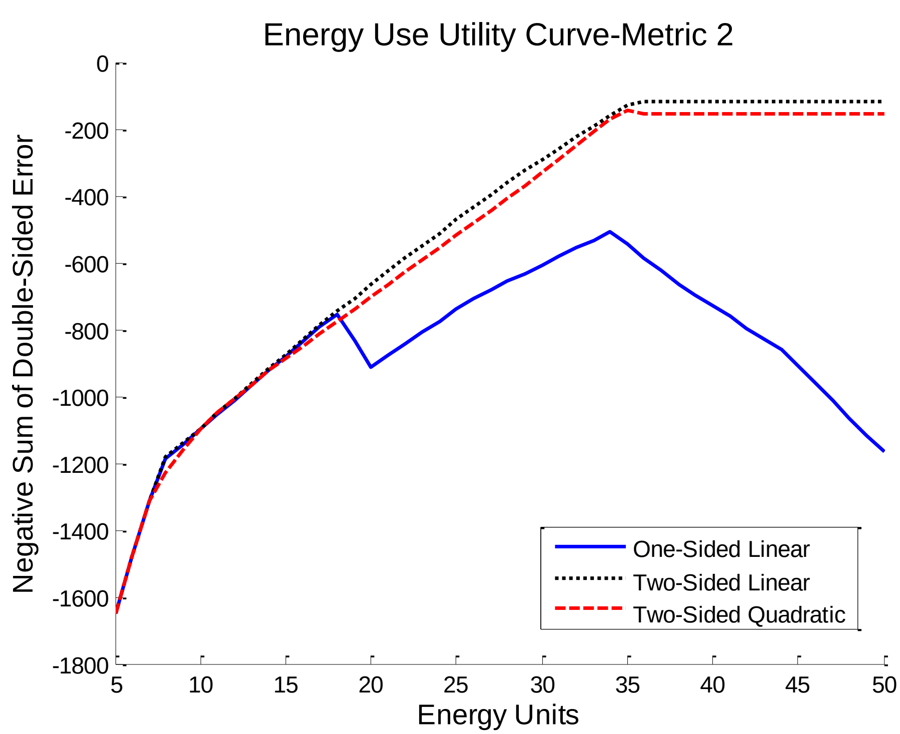

Figure 4-4: Example Energy Use Utility Curve (Negative Sum of Double-Sided Error)

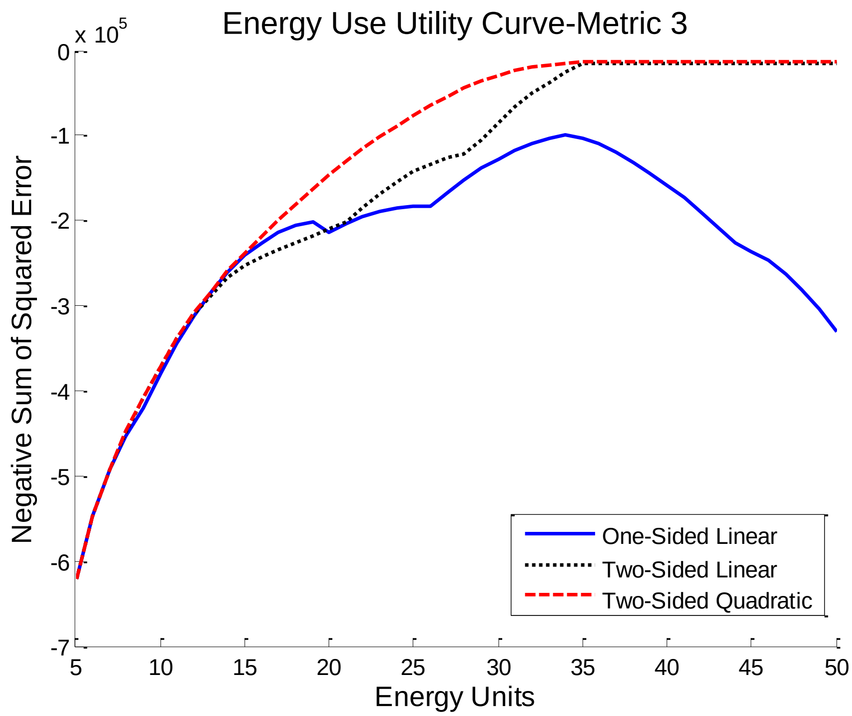

Figure 4-5: Example Energy Use Utility Curve (Negative Sum of Squared Error)

Figure 4-3 shows metric 1 which is the one-sided sum of error. For this case all three optimization constructions perform well with regard to minimizing error on the deficit side. The two-sided method curves are very similar to one another, with both exhibiting a shallow increase toward the upper end of the curve. The one-sided linear method will always continue to increase the energy use in the zone until it reaches the maximum possible energy use, which means it will always use the fully-on energy amount if there is a sensor in the zone incapable of reaching the maximum, unless there are lamps with zero influence on this sensor. In this way, one sensor is allowed to dominate the zone which is not ideal for the other occupants in the zone if they are sensitive to excess light.

Figure 4-4 shows metric 2 which is the two-sided sum of error. Again the two-sided methods show similar behavior as both consistently minimize the linear sum of double sided error and reach minimum double-sided error summation at the same number of applied energy units. The one-sided linear construction does not perform as well in this assessment because it allows some sensors to receive far too much light with fewer units of energy than the other two methods. Because the one-sided construction does not penalize for excessive provision of light, it uses all available energy resources while the preferred light levels of all occupants have not been met, thereby forcing some individuals to receive significantly more light than requested. This behavior decreases the performance of the system as evaluated by two-sided metrics. If over-abundance of light is of concern, a two-sided method is recommended for the reasons shown in Figure 4-4.

Figure 4-5 shows metric 3, the sum of the squared two-sided error. Again the limitation of the one-sided method is shown in that it does not capture the high-side error. The linear two-sided method shares much of the curve path with the two-sided quadratic with the exception of the middle region. Because the quadratic formulation forces the zone to address those sensors with the largest error, it is able to maintain a higher utility value throughout by squeezing the light level met values together for all sensors.

Assuming lamp settings are selected as those which generate maximum utility by the relevant metric, the ultimate settings and resultant light levels that would have been selected with each method are summarized in Tables 4-1 and 4-2.

Table 4-1: Example Resultant Lamp Settings (Fraction of Fully On)

	One-Sided Linear	Two-Sided Linear	Two-Sided Quadratic
Lamp 1	1.00	0.97	1.00
Lamp 2	1.00	0.69	0.68
Lamp 3	1.00	0.44	0.45
Lamp 4	1.00	0.60	0.60
Lamp 5	1.00	0.82	0.86
TOTAL:	5	3.5	3.6

Table 4-2: Example Resultant Sensor Readings (lux)

	One-Sided Linear	Two-Sided Linear	Two-Sided Quadratic
Sensor 1	465	372.7	381.74
Sensor 2	752	500	497.83
Sensor 3	845	500	501.37
Sensor 4	807	500	516.62
Sensor 5	729	500	506.46

Tables 4-1 and 4-2 reinforce the fact that the one-sided construction will always use the fully-on energy quantity if any of the sensors are unfulfilled in their requested light level, even at the detriment of the other occupants in the zone. The tabulated values also show how the quadratic form distributes the error more evenly than the two-sided linear one does. The resultant settings, however, are very similar between the two. The similarity of the results in this energy-unconstrained scenario but difference in performance for lower values of energy use coupled with the computational advantage of the linear methods indicates that the two-sided linear method may be sufficient for scenarios where stringently restricting energy use is not of immediate concern but that the quadratic form may be a better choice for demand response-type conditions or in buildings where the owner is very intent on limiting energy use for cost or other reasons.

4.2 Building-Level Resource Allocation

While the zone-level control is important to ensuring rapid response to local changes in the system, a building-level element is necessary for performance tracking and resource use decision making. The total building-wide energy use over time is of interest for operating cost projection and minimization, demand response, and system-wide performance management purposes. The ability of the building to maintain a high quality occupant experience is also important as the productivity of the occupants is of significant value to the building owner or lease holder. For these reasons, it is necessary to maintain a centralized repository of information and level of control over the system. Without this component, the building manager cannot monitor the status of the building and the building cannot participate in a demand response request due to the lack of a mechanism by which to decrease energy use by a specified amount. In order to exert authority over the building-wide distribution of energy units, the building server disburses units of energy based on the utility curves defined by the zones and supervises trading of these units between the zones as the dynamics of the building change over time.

Information gathering at the building level is a complicated balance between the desire to make effective decisions and to limit the amount of data transmitted and stored.

While the building server has significantly more processing and storage capabilities than the zone-level managers do, both of these capacities are finite. Further, it is expedient to minimize the data transfer requirements between the zone-level and the building-level by taking advantage of the distributed processing capabilities of the zone managers. For this reason, the building-level server is designed to periodically collect the utility curves from the zones and use them to allocate the building-wide energy resources by zone.

The allocation at this level is determined by a modified greedy algorithm on the basis of the zone utility curves. While with a traditional greedy algorithm the next increment of the resource is allocated to whomever demonstrates the largest immediate need, this modified version allows the building server to look up to five steps in the future to assess the potential for increased utility based on increased energy allocation. While only one energy unit is allocated at a time, the unit is allocated to the zone with the highest average utility increase per unit if given one to five more energy units. This means the utility increase is calculated and averaged for one, two, three, four, and five additional units of energy and the highest of these averages is the number the zone uses to compete against the other zones. The incorporation of this forward-looking perspective helps to ensure that zones do not get trapped on low or intermediate level plateaus in their curves and instead are able to take advantage of steeper points in their respective curves. The five step forward cutoff is selected based on the scale of the problems considered and the size of the energy units selected. Five steps is equivalent to half the power of a fully on lamp. The utility curve values are stored as a vector of values for each zone with the vector length determined by the number of occupants, or sensors, per zone, and throughout the energy allocation, a counter variable stores the position along the curve and the number of energy units that have been assigned to each zone. This counter is updated at each unit allocation to reflect the new allocation profile for the next iteration. Energy allocation stops when either the maximum allowable building energy use is reached or the zones are all fulfilled above a threshold utility level.

To facilitate additional demands on the zones that may occur before the next update, any energy units remaining after the zones are assigned enough energy units to reach the top of their curves are divided among the zones. The zone managers automatically limit their zone energy use to the minimum necessary to fulfill the demand so these additional units will only be used if the demand increases during the next interval. The building server, however, stores the actual original energy unit assignment value as the starting point for the next round of energy trading.

In the energy trading rounds, the units start at their previous energy allocation levels and trade energy units based on relative utility level. The zones are divided into prospective buyers and sellers based on where they fall with respect to the average utility level, those who fall below are buyers and those who are above are sellers. There is no actual currency in the exchange of energy units; instead the units are exchanged based on the value each zone places on the particular unit of energy. In order for the buyers to "purchase" energy units from the sellers, their "buying price" is calculated in the same way the value was computed in the original energy allocation. The sales price for the sellers is similarly computed but is taken as the minimum average loss over losing up to three units of energy. In order for the exchange to take place, the buyer has to be willing to pay more than what the seller is offering. The trading could be designed in two ways. In one approach, the highest bidders buy from the lowest sellers until no buyers are willing to pay the purchase price. Alternatively, the maximum number of transactions can be forced to occur by matching as many buyer-seller pairs as possible. This second method is selected to allow as many transactions to occur as possible per round.

This buying and selling process continues iteratively until there are no buyers willing to pay what the sellers require or the trading results in back-and-forth cycling. At this point, zones with more energy units allocated than they need to remain at the peak of the utility curve are stripped of their extra units. After this process is finished, any additionally available energy units are allocated in an identical manner to the original allocation scheme, first according to utility and then disbursed evenly to give room for increased demand.

In the absence of concerns as to the expediency of the processing of the data and data transmission requirements, the building-level server could perform a global optimization of the same form as the zone managers for the entire building, constructing a global influence matrix for the entire building system. As the zones are considered entirely independent of one another, there are no coupling terms between the zones in the global influence matrix. The solution to this problem would be guaranteed to yield a globally optimal solution unlike with the greedy solution. However, this process would be very inefficient and could lead to an intractable problem in large buildings as the number of constraints and variables grows large. A comparison of the results from this global optimization and the modified greedy algorithm is discussed below.

All sample zones presented below have five sensors and five lamps and the two-sided linear method is used to define the utility curves for this assessment. The utility curves for the five zones are presented below. Each zone has a uniquely defined influence matrix resulting in distinct energy use utility curves. Zone 3 has the same composition as the zone used in the previous section.

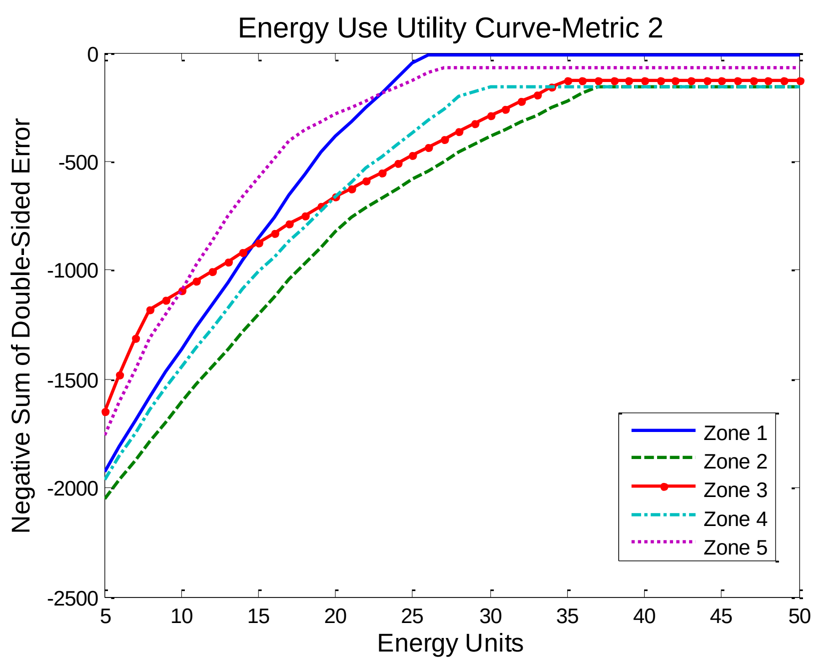

Figure 4-6: Energy Use Utility Curves for all zones used in Resource Allocation Comparison

As shown in Figure 4-6, only Zone 2 in this case is able to achieve a total error sum of zero; all other zones reach their peak utility with still some residual error in the system indicating that not everyone has received the exact quantity of light they desired. Although some of the zones showing residual error may be capable of at least meeting the preferred light level of the occupants, they do not have the flexibility to meet exactly the target light level for all occupants simultaneously because the influence matrices are coupled with each lamp influencing multiple sensor readings.

A comparison of the allocation scheme used in the utility-based resource allocation system with a building-wide two-sided linear program (LP) was conducted to examine the ability of the utility-based system to approximate a building-wide optimization. The constraint matrix for the building-wide LP is a diagonal assembly of the individual zone influence matrices with no coupling terms. A schematic of this matrix assembly is shown in Figure 4-7 where the matrices $\boldsymbol{A_{z*}}$ are the zone influence matrices for the respective zones. The results of using both techniques to assign energy units to the zones progressively are presented in Figure 4-8. The similarity shown in the two plots in the figure arise from the structure of the constraint matrix used for the building-wide LP, in that the zones are considered to be independent and the matrix thus has no coupling terms.

$$\begin{pmatrix} [\boldsymbol{A_{z1}}] & [0] & [0] & [0] \\ [0] & [\boldsymbol{A_{z2}}] & [0] & [0] \\ [0] & [0] & [\boldsymbol{A_{z3}}] & [0] \\ [0] & [0] & [0] & \ddots \end{pmatrix}$$

Figure 4-7: Building-wide Influence Matrix Assembly

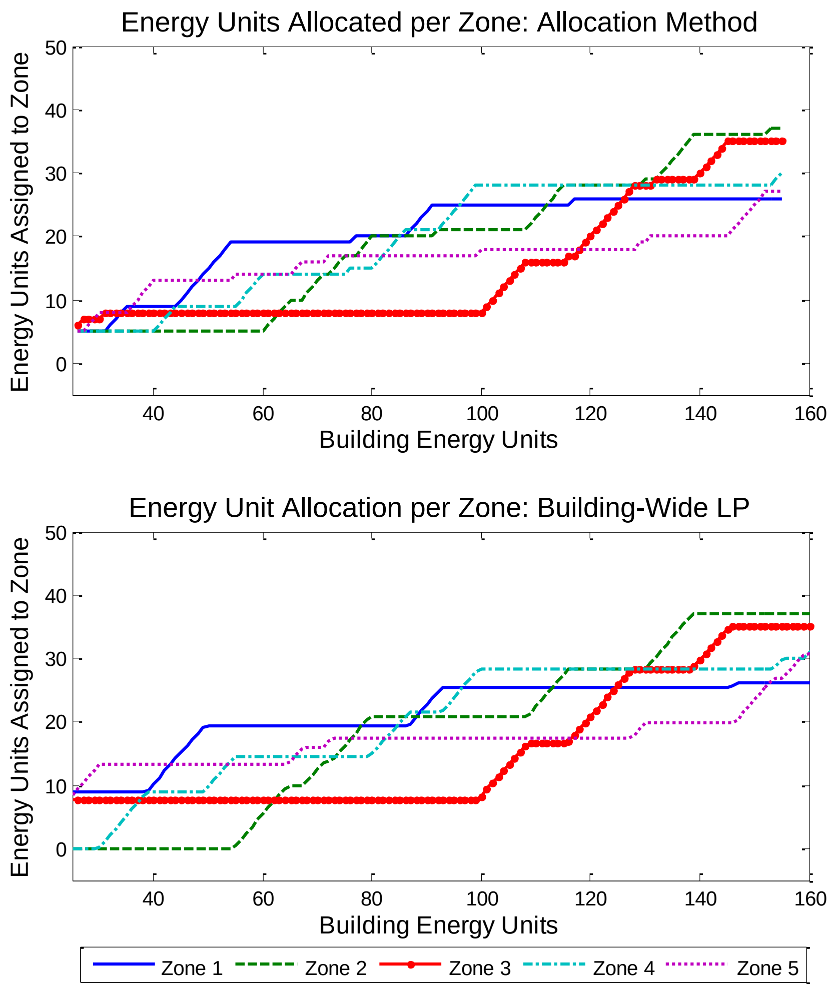

Figure 4-8: Energy Unit Allocation for Allocation Scheme and Building-Wide Linear Program

The resultant utility values during the assignment are presented in Figure 4-9 for each zone and as a summation across the building in Figure 4-10. These figures again show the similarities in the results of the two allocation schemes. The difference in the two sets of curves shown in Figure 4-9 is largely due to the restriction on the allocation scheme that a full energy unit is applied to one zone at a time whereas in the building-wide LP, the energy distribution is calculated in equal sized units of energy but the units can be spread over multiple zones and lamps within the zones. The zones also start at different allocation points. The building-wide LP assignment starts at 0 building energy units and reaches a varied allocation level per zone for a total building wide allocation of 25 units. The zone-level scheme starts at one energy unit allocated per lamp allocated to each zone which starts all zones at the same allocation level at a total building allocation of 25 units.

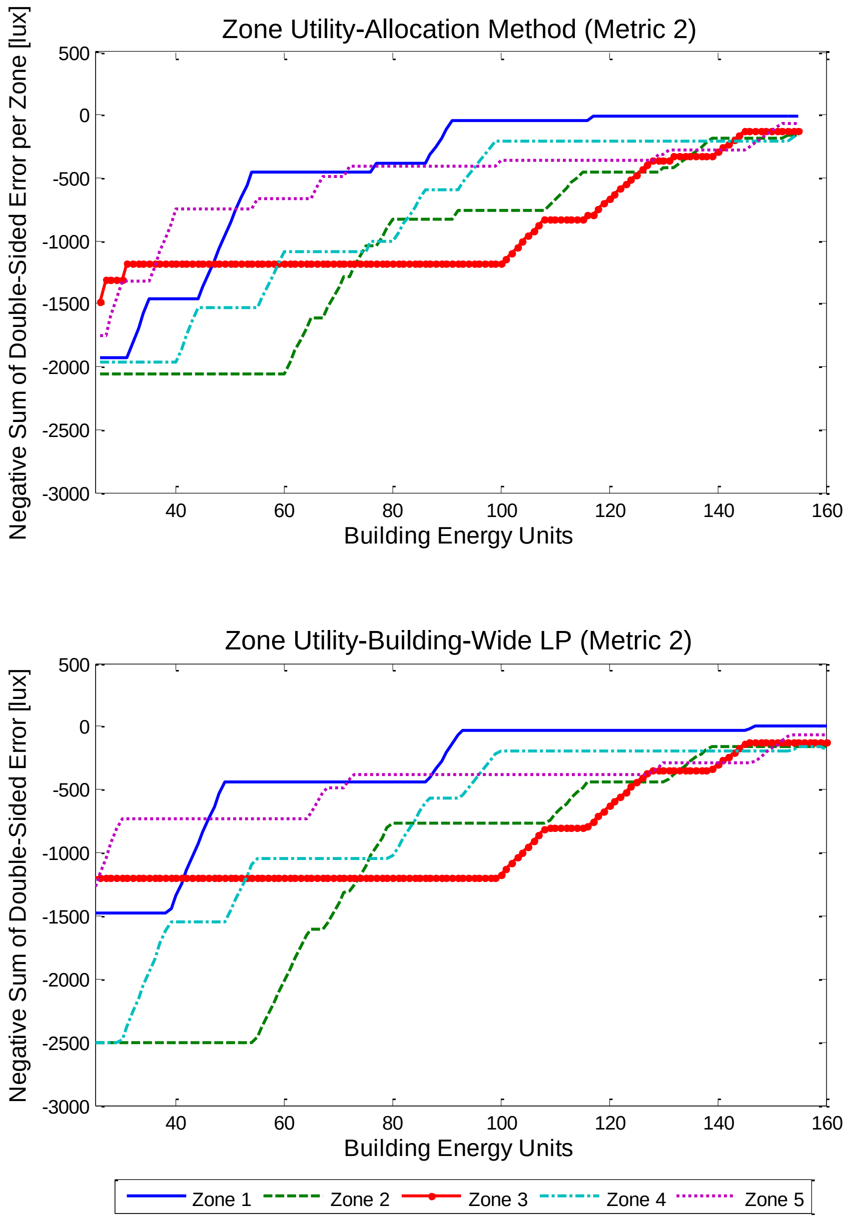

Figure 4-9: Zone Utility per Building Energy Use for Allocation Scheme and Building-Wide Linear Program

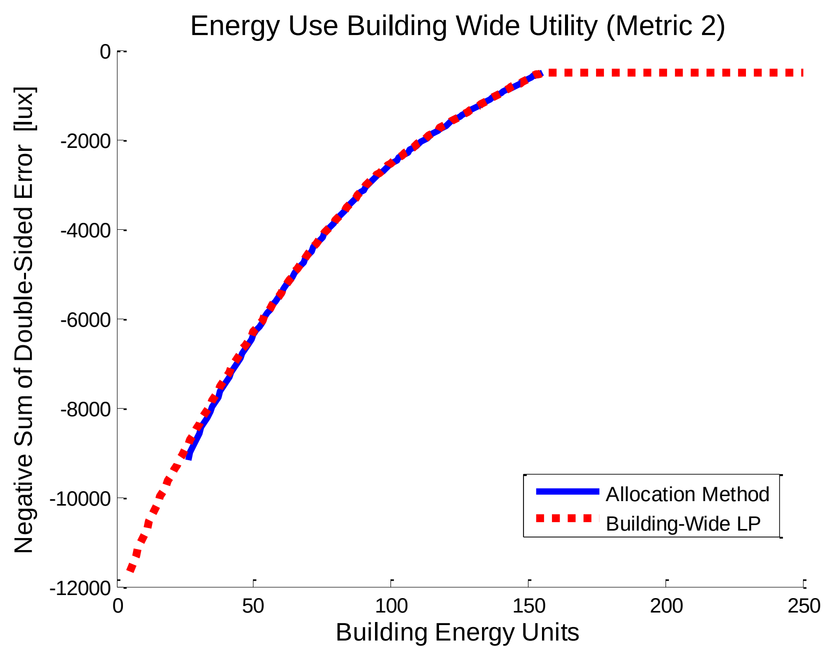

Figure 4-10: Building-wide Utility Curves for Energy Allocation Scheme Comparison

Figure 4-10 reveals the similarity in the resultant building-wide utility values as energy is allocated by the two resource allocation schemes. The Building-wide LP curve spans the full range of possible building-wide energy usage and reaches a plateau at the energy use level indicative of maximum utility. The allocation method which assigns energy units based on the results of the distributed linear programming problems performs very well indicating that tracing the resultant curves of the separated sub-problems is a computationally efficient substitution for running a building-wide program.

Because the building server uses basic arithmetic computations to perform the utility-based resource allocation and only temporarily stores the current utility curve for each zone, it should be capable of distributing resources among thousands of zones. The communication requirements are anticipated to be more restrictive to the scalability with respect to the number of zones. If the zone manager to building server communication is to be conducted wirelessly, it would be subject to the same signal rate, communication distance, and node addition limitations as are discussed in Chapter 3 for the wireless communication within the zones. For a large building where communication distances can be quite long and large numbers of zones would be desired, wired connections between the zone managers and the building server may be necessary to increase communication speed and avoid excessive data hopping over large numbers of nodes. Alternatively, additional tiers of pseudo-building servers, processing units performing the same types of calculations as the centralized building server to create aggregated utility curves, could be inserted to add additional layers of data consolidation with the aggregated utility curves from groups of zones being passed up the hierarchy to the central building server. Aggregating groups of zones together could both further decrease the communication overhead of the system and reduce the necessary computation at the central building server level.

4.3 Summary

A tiered energy use optimization and resource allocation scheme has been developed to assign energy units and dictate lamp settings for all lamp fixtures within a building. This scheme takes into account light level preferences of occupants, the presence of occupants, and the desire for energy savings. The tiered structure allows for rapid updating at a local zone level while maintaining the ability to limit the total building energy use, track performance, and initiate demand response at a centralized location.

The optimization and resource allocation system requires an initialization procedure which maps all lighting fixtures to all sensors within a zone and makes use of this information when setting new light levels for the zones. The initialization process is entirely automated and can be performed at regular intervals to account for lighting system degradation and can be initiated when a zone is renovated or reconfigured to immediately update the system.

The building is divided into zones, each of which has a set of controllers and sensors which provide distributed control and sensing capabilities. The zone server uses the distributed sensing information to create a utility curve for energy use in the zone and optimizes the control settings for all lamps in the zone to use the minimum required energy to achieve the desired performance level.

At the building-level, a centralized server is utilized to allocate energy to the zones using a modified greedy algorithm. This algorithm makes use of the utility curves defined by the zones and awards units of energy according to incremental utility gains. Implementing this level of control enables the building manager to specify maximum energy use for the whole building and ensure that it is allocated efficiently. The building manager can also select a minimum performance standard to be maintained across the building to ensure a quality experience for the occupants. With this scheme, the building manager is able to view the tradeoff between energy use, and therefore operating cost, and occupant-centric utility to better understand the tradeoffs in the decision-making process. A decision to initiate a demand-response energy usage reduction can be viewed in terms of the actual performance degradation, and if carried out, the reduction will be applied in a least impactful manner.

Chapter 5.

PHYSICAL IMPLEMENTATION

A laboratory-scale setup is designed to implement a version of the resource allocation system. The purpose of the setup is to demonstrate the physical implementability and assess the general behavior of the system in a real world environment. The sensing, control, and computation units discussed in Chapter 3 form the backbone of the distributed wireless system and a laptop connected to a wireless modem serves as the building server. The system is arranged in a room with large windows to the outdoors to demonstrate the natural light compensation capabilities. The system evaluation was conducted over several time increments designed to capture sunset or both sunset and sunrise to demonstrate the system response as a result of varying external light contributions. This chapter begins with the motivation behind this physical test followed by an explanation of the design of the test setup. A discussion of the results of the tests follows, including energy and performance implications as well as a demonstration of the types of decision-making information made available through implementation of this system. The chapter concludes with a summary of the physical test and a discussion of the lessons learned regarding physical implementation.

5.1 MOTIVATION

The purpose of this physical experiment is to capitalize on currently available technology to demonstrate the feasibility of the system for real world use. This test includes a prototype hardware system, embedded software and computation, and data logging at the building server level. While Chapter 6 discusses the theoretical performance of the system in a full building installation, this physical implementation is conducted to investigate the specific challenges of designing a functional system that occur outside the scope of an idealized simulation. These challenges include determining (1) how to structure the communication between the wireless units, (2) what information needs to be transmitted and to and from whom it must be sent, (3) how to assign roles to the units and ensure compatibility between levels, (4) how to embed the decision making abilities on the units, and (5) whether sufficient robustness can be developed to establish a consistently functional system. This functional system further demonstrates resilience to non-idealities in sensing and sensor design, in information transfer, and in timing, a competency necessary to a full building-wide deployment. The successful implementation of one instance of this system demonstrates the ability to resolve the aforementioned challenges and translate the initial concept into a usable, real-world system.

5.2 Laboratory Setup and Experiment Design

The room used for the laboratory experiments is a southwest facing first floor room in an office and classroom building in Palo Alto, California. The room is selected due to its external windows which allow varying levels of natural light to enter the space throughout the day. As the goal of the experiments is to show the ability of the system to track and hold a set light level, the ability to successfully compensate for sunlight availability is essential.

5.2.1 Components and Parameter Selection

The physical components of the experimental system are:

- (6) Wireless sensing, control and computation boards with XBee radios
- (6) Light-to-Voltage Converter photodiodes each with an RC filter
- (6) Mark VII (0-10V) Phillips Advance Ballasts
- (6) Constructed lamp fixtures each consisting of a wooden base, (4) Rapid-Start fluorescent sockets, (2) T8 32W Lamps
- Windows OS Laptop with serial connection to an XBee radio

The laptop is designated as unit 1 for the purposes of the test. The other components are divided into 6 physically identical units numbered 2 through 7. Each of these units is managed by a wireless board and is complete with sensing and actuation capabilities.

The units are grouped into two zones; the zone managers are units 2 and 3. Units 4 and 5 report to zone manager 2 and units 6 and 7 report to zone manager 3. The zone managers additionally communicate with the building server, unit 1. For clarity, the zones are numbered in accordance with the unit number of the zone manager.

The one-sided linear formulation is used for the physical implementation to show the capabilities of the system in minimally meeting a set light level. The participation weighting factors for all sensors are set equal to 1.

5.2.2 Physical Layout

The laboratory room is approximately 30 feet by 15 feet with the longer dimension along the external wall. The general layout of this room is shown in Figure 5-1. The room has six southwest facing windows arranged in two horizontal rows spaced very tightly together. Toward the east end of the room are two sets of windows situated adjacent to one another and toward the other end of the room is the third set of windows. Two tables of standard height are placed in the middle of the room as repositories for the constructed lamp fixtures and associated components. The lamps are placed perpendicular to the windows and are spaced roughly evenly parallel to the external wall. Units associated with zone leader 3 are placed on one table and units associated with zone leader 2 are placed on the other table thereby creating two physically separated zones. A series of poster boards is placed between the two tables to minimize the light transferred between zones. The poster boards are excluded from Figure 5-1 for clarity. As the zones are not entirely isolated from one another, a small amount of reflected light from each zone is allowed to enter the other zone. While this does not match the idealized case where zones would be photo isolated from one another, it allows for the examination of the more realistic case where a small amount of spillover from one zone to another is anticipated.

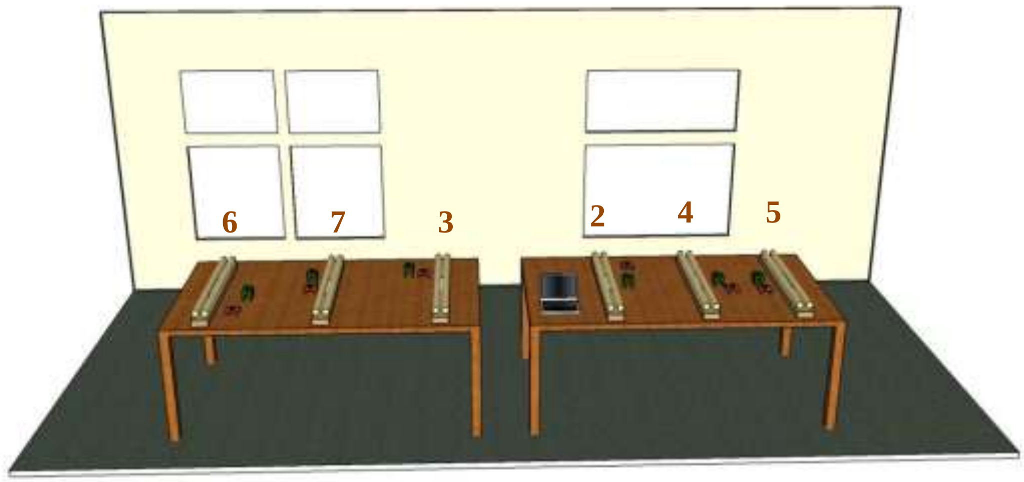

Figure 5-1**: Laboratory Room Layout**

The constructed lamps are spaced roughly two feet apart on center and the tables are approximately four feet from the windows. The taller units in the image above represent the wireless units while the shorter units represent the sensors. Throughout the tests, the sensors are moved around and reoriented to evaluate system behavior. The orientation and location of the sensor is important because the sensors used in the test are sensitive to incident light within approximately 40 degrees to normal from the sensor face. Sensors facing directly into the lamp or window will therefore be more sensitive to the incident fluorescent or sunlight respectively. This narrow view range is ideal for this application because it allows for the creation of multiple scenarios within a limited testing space. For a real-world application, a sensor with a wider view capable of capturing more of the diffuse light contribution would be preferable, or perhaps multiple small view sensors coupled and oriented to view a wider array of angles.

5.2.3 Natural Lighting Conditions and Timing

The experiments were conducted in Palo Alto, California (37.43 N Latitude, 122.17 W Longitude) between May and August of 2011. Sunset at this time of year occurs around 8:00 PM PDT. For this reason, the sunset tests are conducted from the late afternoon and to the early evening. All times reported on figures are local pacific daylight time and therefore true sunset occurs near 20:00 on these plots.

5.2.4 PARAMETERS FOR PHYSICAL TEST

For all of the test runs, the zone managers are asked to maintain an illuminance level at all sensors of at least 500 lux. This value is selected because it represents the standard office desktop illuminance requirement in the United States. While this value is important in a real office scenario, due to the constraints in the laboratory setup, the value is more arbitrary as realistic distances and orientations between the sensors and the lamps are not used. As the system has been designed to accommodate any level of individualized user preference and the prototype setup has no actual occupants to define their preferences, this value is simply used as a placeholder value for a preference that would be selected by an individual occupant. The occupant is assumed to be constantly present as an occupancy sensor was not designed for the test and the laboratory test space is an unoccupied room. Keeping the occupancy and desired light level parameters constant throughout the tests contributes to a clear investigation of the system behavior without the distraction of arbitrarily changing parameters.

The timing of the updates at the building level and zone level are built in to the design of the system. The zones update their own curves and the settings of the lamps in their zone every 20 seconds. The building server is set to request curve updates and allocate zone energy every 100 seconds. While these nominal times are explicitly defined by the code, the actual times vary due to the state machines implemented in the respective codes. Upon entering the state to request a new reading and compute a new utility curve, the zone manager cannot receive a curve request from the building server until that new curve has been computed. Similarly, the timer for the next update in the building server will not start until the curves have been properly received and recorded from both zones. These delays cause a slowdown in the system that is acceptable because the system update frequencies are set primarily for data gathering purposes and maintaining acceptable performance does not require rapid updating.

When the zone manager calculates the utility curve for the zone, it computes the curve until the utility level reaches at least 98 percent of the total utility range. In this case there are three sensors each requesting 500 lux so the zone manager stops allocating energy when the utility value reaches -30 lux. This value is selected to ensure a high level of conformance to the simulated occupant preferred light level setpoint while providing a stopping point for the curve calculations. The building server is set to allocate available energy resources to the zones until their utility level reaches 95 percent of the total utility range, -75 lux. Selecting this parameter for this prototype experiment is analogous to a building manager setting a minimum building level performance standard with respect to occupant satisfaction. Maximum energy use in these tests is unrestricted to allow observation of the performance of the system with regard to tracking and holding the setpoint.

5.3 RESULTS FROM SYSTEM EXPERIMENTS

The following sections detail results gleaned from multiple, independent runs of the system. The tests are performed on different days with varying sensor orientations and positions to display results from a variety of setups. All tests were conducted using the same equipment and room configuration. For all tests, unit 2 serves as the zone manager for units 4 and 5 and unit 3 serves as the zone manager for units 6 and 7.

5.3.1 TEST A

Test A was conducted on May 11, 2011. This was a sunny day with a clear sky. The test was conducted from 4:04:53 PM to 7:44:17 PM PDT. True sunset on this day was 8:08 PM.

5.3.1.1 Influence Matrices

The first step in running a test is to allow the system to calculate the influence matrices for use in determining the light settings for the remainder of the test. This phase is termed the commissioning sequence. To form these matrices, the zone managers request all of the lamps in their zones to be turned off initially. Subsequently each lamp in the zone is turned fully on individually and the sensor levels are measured and recorded. Once all recordings are collected and with all of the lamps in the fully off state, the residual light level is recorded and subtracted from the earlier readings to give the net increase in light level due to each lamp. This subtraction is crucial in the test setup as the initialization procedure is conducted daylight hours and although the window blinds were closed manually for the duration of the initialization procedure, some sunlight penetration into the space was unavoidable.

The influence matrices for the two zones are listed in Tables 5-1 and 5-2. As shown in these tables, the sensors are oriented largely toward the light controlled by their respective units with a smaller lighting component received from the neighboring lamps. In zone 2, unit 4 is located between units 2 and 5 resulting in larger contributions from lamps 2 and 5 to sensor 4 than to sensors 5 and 2 respectively. In zone 3, unit 7 resides between units 3 and 6 and shows a similar pattern of light contribution from its neighboring lamps with the exception of the low influence of lamp 6 on sensor 7 which occurs due to the orientation of sensor 7. In all cases, the sensors are capable of receiving well over their assigned preference level of 500 lux as the sum of each row greatly exceeds this level. The task of the system, therefore, will be to determine the lamp settings which provide this level of service while using minimal energy resources.

Table 5-1: Test A Zone 2 Influence Matrix [lux]

	Lamp 2 at 100%	Lamp 4 at 100%	Lamp 5 at 100%
Sensor 2	1311.7	154.7	41.3
Sensor 4	26.5	1243.9	47.2
Sensor 5	30.9	89.9	1361.8

Table 5-2: Test A Zone 3 Influence Matrix [lux]

	Lamp 3 at 100%	Lamp 6 at 100%	Lamp 7 at 100%
Sensor 3	1295.5	19.2	153.3
Sensor 6	19.2	896.1	60.4
Sensor 7	89.9	7.4	1195.2

5.3.1.2 Light Sensor Readings and System Control Settings

During the test run, the sensor readings and lamp setting commands were recorded. Figure 5-2 plots the voltage measured by the light sensors for all of the units over the entire testing time post commissioning. The readings do not track a stationary value even though the targeted illuminance level is stationary. The reason for this lies in the spectral composition of the two types of light sources. As discussed in Chapter 3, the sensors are more sensitive to sunlight energy outside of the visible spectrum than they are to fluorescent light outside the visible spectrum. Because the system is designed to separate these two components and translate them to their visible light contributions, it is expected that the sensor measurements will be lower for a light level provided mostly by fluorescent light than for one provided mostly by sunlight for a given visible spectrum target. As anticipated, the sensor voltage readings decrease as the sun sets late in the day as fluorescent light supplies a larger percentage of the total available light. Plotting the measured voltage over time provides insight into the location and orientation of the sensors relative to the windows. As sensor 6 was placed closest and most directly facing the external windows, it shows the largest change in sensor reading over the time span due to the aforementioned effect.

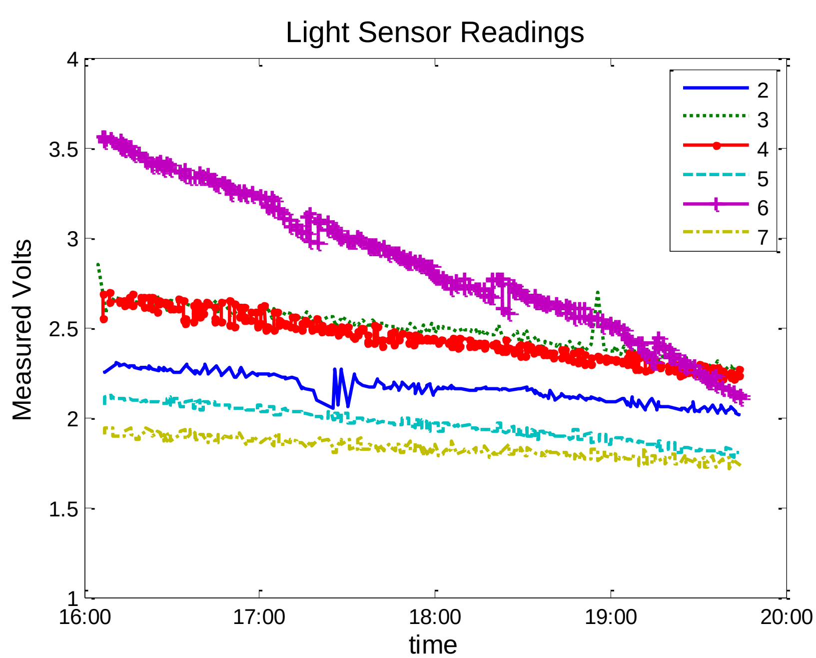

Figure 5-2: Test A Light Sensor Readings

Figure 5-3 shows a calculated estimate of the visible light maintained by the system throughout the duration of the test. The values plotted in the figure are computed in a manner similar to the processing of voltage measurements into light settings as conducted by the zone manager microcontroller. The building server eavesdrops on all communications throughout the system to store the light settings and sensor readings throughout the test duration. The values recorded during the initialization procedure are recorded for use in reconstructing approximate visible light level values. The visible light levels displayed in Figure 5-3 are calculated at regular 2-minute intervals using the nearest sensor readings and lamp settings because the readings and settings do not occur at simultaneous time steps. Because the calculations are performed using assumed influence matrices based on the recordings of the building server eavesdropping which includes dropped and incomplete transmissions, they do not necessarily represent identical calculations to those performed by the zone manager. Additionally, as these are calculated values and not measured data, the values plotted are indicative of what the system believes the visible light level to be and is not an independent assessment of the accuracy of this computation.

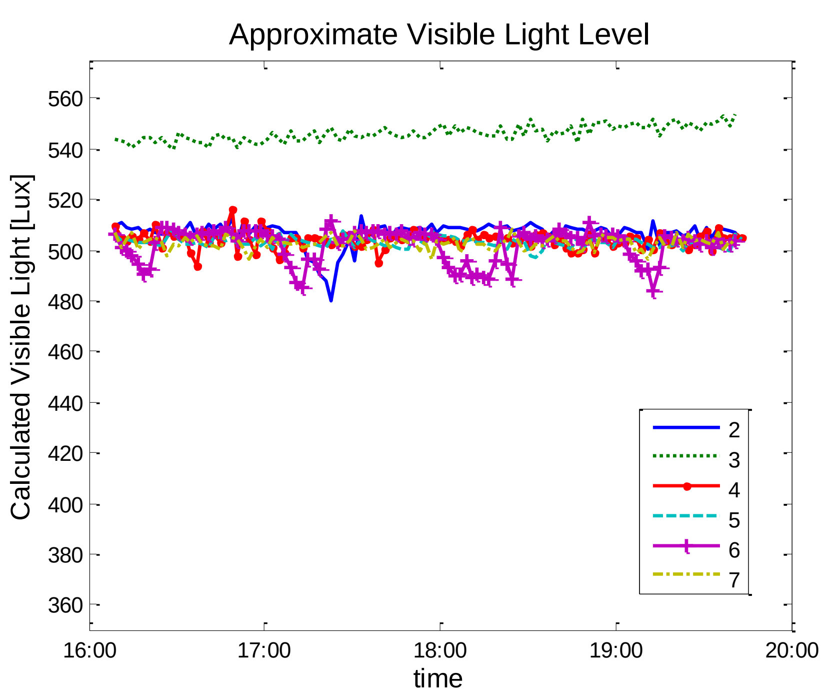

Figure 5-3: Test A Approximate Visible Light Levels

The calculated values plotted in Figure 5-3 show that throughout the duration of the test, the target value of 500 lux is well maintained for most of the sensors. Sensor 3 shows a consistently higher visible light level than is required by the 500 lux target, but as the one-sided implementation is used for this test, this surplus of provided light does not contribute to the error computation. The light surplus shown for this particular sensor may be contributed to by an error in the assumed influence matrix used in the computation of the figure values.

The time series plot of lamp settings shown in Figure 5-4 further indicates the light level tracking abilities of the system and demonstrates the natural light compensation capability of the lighting system. The settings of all lamps are increased as the sun sets to compensate for the loss of sunlight. As sensor 6 is the sensor most affected by sunlight due to its location and orientation, lamp 6 which has the largest direct influence on the light incident on sensor 6 has the largest change in light setting to compensate for this loss of natural light. The other lamps also increase their settings throughout the duration of the test, but as the sensors they most affect are not oriented to receive large natural light contributions, their changes are more moderate.

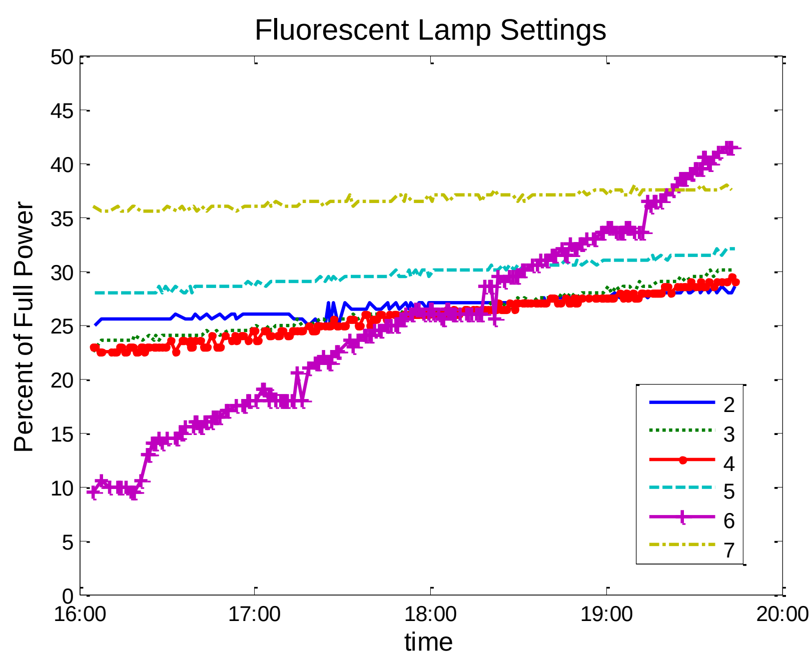

Figure 5-4: Test A Lamp Level Settings

5.3.1.3 Utility Curve Progression

The utility curves from the test are also recorded. The progression of these curves is plotted in Figures 5-5 and 5-6 with earlier curves in dark black fading to light grey for the later ones. Zone 3 shows a greater variability in curve shape due to the presence of sensor 6 which causes the zone to be more sensitive to the presence of natural light. As shown in Figure 5-6, zone 3 requires less energy during the daylight-prevalent hours than zone 2 to fulfill the requested light levels. The impact of the presence and orientation of the sensors to the energy demand for meeting the requested light level for the zone underscores the importance of intelligent design in daylit spaces in buildings to maximize the integration of the daylight for artificial lighting use mitigation. The point where the curves stop increasing is dictated by the minimum utility level set for the test. In this case, the curves reach this minimum utility at 98 percent of the total available utility range, -30 lux. The energy use at the point where the zone reaches this standard is the quantity of energy the building allocates to the zone in an unconstrained situation. Some curves surpass this 98 percent standard because the discretized energy increments in the computation do not yield a 98 percent result exactly.

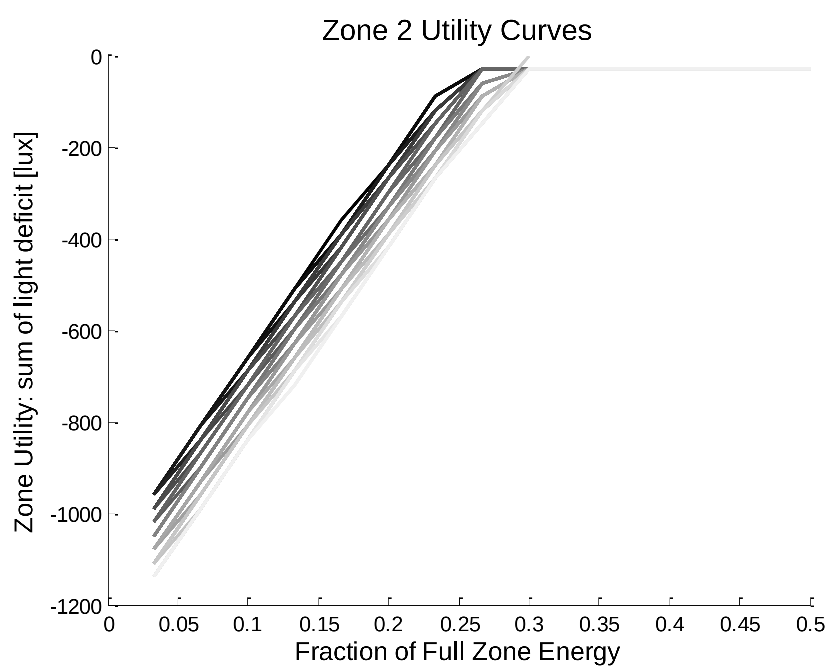

Figure 5-5: Test A Zone 2 Utility Curves

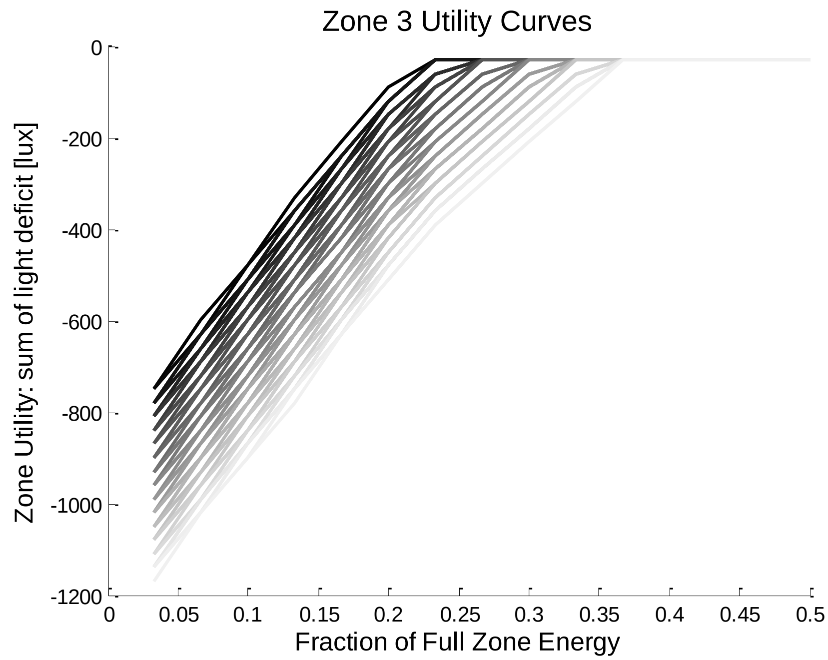

Figure 5-6: Test A Zone 3 Utility Curves

5.3.2 TEST B

A second test was conducted to evaluate the scenario of multiple sensors oriented toward the windows with the intent to capture a larger sunlight component. Relocating the sensors requires the recalculation of the influence matrices in a similar way to how a building would be recommissioned after a reconfiguration of the interior space. The ability to perform this type of system update automatically is one of the major benefits of this type of system in a building that may have tenant turnover during the lifespan of the structure. This test was conducted on May 13th, 2011 with a sunset time of 8:09 PM PDT.

5.3.2.1 Influence Matrices

The new influence matrices for this test are shown in Tables 5-3 and 5-4. The matrices show strong relationships between the sensors and lamps of the same unit number due to the design of the experiment.

Table 5-3: Test B Zone 2 Influence Matrix [lux]

	Lamp 2 at 100%	Lamp 4 at 100%	Lamp 5 at 100%
Sensor 2	1006.6	206.3	20.6
Sensor 4	19.2	655.8	35.4
Sensor 5	32.4	109.1	1380.9

Table 5-4: Test B Zone 3 Influence Matrix [lux]

	Lamp 3 at 100%	Lamp 6 at 100%	Lamp 7 at 100%
Sensor 3	1139.2	19.2	69.3
Sensor 6	45.7	778.2	98.7
Sensor 7	78.1	145.9	1064.1

The matrix for zone 2 also shows strong cross relationships between lamp 4 and sensors 2 and 5 and between lamp 5 and sensor 2. It also shows a relatively low peak contribution from lamp 4 to sensor 4. These two pieces of information lead to the expectation that lamp 4 will be used as a primary source in this zone as increasing energy applied to this lamp most efficiently increases the light level for all sensors and is required for enabling sensor 4 to maintain the requested light level.

Zone 3 has a smaller degree of coupling between lamps and sensors connected to different units than zone 2 as demonstrated by the large values along the diagonal and relatively small values in the off diagonal. The low degree of coupling is similar to the zone layout in Test A.

5.3.2.2 Light Sensor Readings and System Control Settings

As expected from the influence matrices, Figure 5-7 shows that lamp 4 maintains a relatively high setting throughout the experiment. The two other lamps, 2 and 5, in zone 2 are able to maintain lower settings due to the contribution of light from lamp 4 to all sensors in the zone. In zone 3, lamp 6 exhibits interesting behavior. Because lamp 6 yields the greatest influence on sensor 6, it has the lowest diagonal matrix value, and sensor 6 is highly sensitive to sunlight levels, lamp 6 starts out with the lowest setting and ends with the highest setting in its zone after the sun has set. Figure 5-8 shows the corresponding voltage measurements from the sensors.

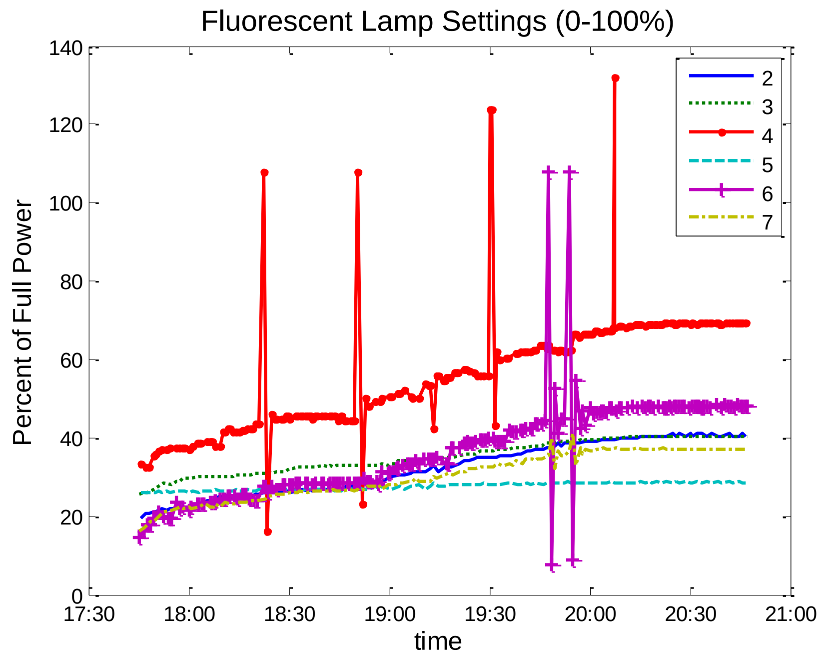

Figure 5-7: Test B Lamp Settings

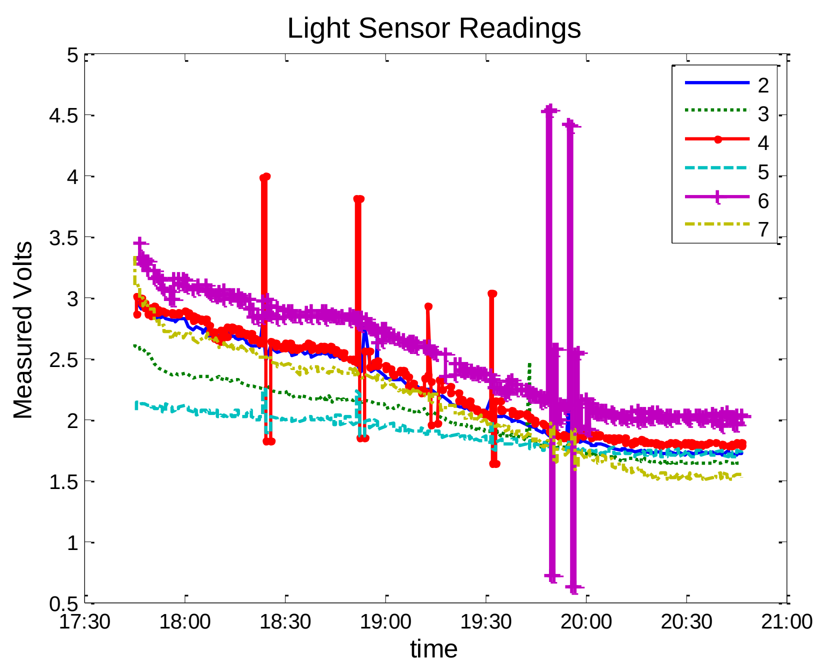

Figure 5-8: Test B Sensor Readings

This test is an interesting case because it shows both the natural light compensation abilities of the system and the ability of the system to recover from transmission errors. The prototype system is designed with very basic transmission error checking which allows for a small number of erroneous transmissions to propagate through the system. Figure 5-7 shows the fluorescent lamp settings transmitted to the units as recorded by the building server. It is important to note that the settings as plotted are the values transmitted by the zone manager during the test, not directly the values calculated by the zone manager. As is evident from this plot, there are false settings requests transmitted that are outside the scope of feasible control commands (0 to 100). As the errors are on the high end of the control spectrum, the lamps respond by switching to 100 percent of their output which vastly exceeds the required light level. In each case the control system overshoots on its response in the next query but then quickly returns to tracking the requisite light level. A corresponding jump in sensor reading, as shown in Figure 5-8, matches the lamp setting increase most prominently in the closest sensor but smaller responses are visible in the neighboring sensors as well. This overshoot response is due to the system using the internally recorded previous setting of the lamps to calculate the current natural light contribution to the space. As the zone manager bases this calculation on the correct previous setting for the lamps and not the erroneous transmitted one, it perceives some of the artificial light contribution as natural light and therefore sets the lamps too low on the next iteration. Because the required past information for the system only goes back one time step, the system is able to quickly recover to the proper settings. While this set of errors is specific to this implementation and test run, the ability of the system to be resilient to error is important as even a commercial grade implementation cannot be entirely free of transmission and execution errors.

In this test, the overall trend toward decrease in energy use and decrease in sensor readings is more pronounced than in Test A because the sensors are oriented to capture more of the daylight contribution. As the sunlight contribution represents a larger portion of the total light available at the start of this test as compared to Test A, the sensor readings show a larger decrease over the course of this experiment as sunlight is less visible to the sensor and fluorescent light contributes more heavily to the total sensor reading. The maintained visible light level for this test, calculated using the same steps as described for Test A, is shown in Figure 5-9.

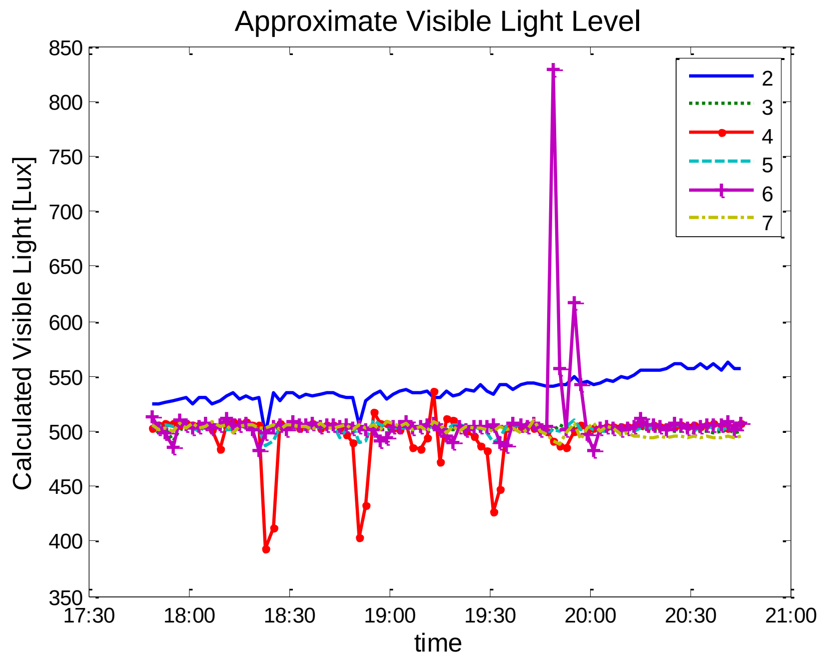

Figure 5-9: Test B Approximate Visible Light Levels

Similar to the results of Test A, one of the sensors, sensor 2, receives more than the preferred light level. This calculated light surplus is a result of both the lack of penalty for excess light provision and the possibility of inaccuracies in the generation of the assumed influence matrix. In a real-world scenario excess light as well as insufficient light can be bothersome to occupants so an alternate formulation of the optimization algorithm is considered in Chapter 6 that incorporates consideration of the provision of excess light.

5.3.2.3 Utility Curve Progression

In this test, because the sensors are oriented more directly facing the windows, the utility curves show a larger change over time than in the previous setup. The utility curves for both zones are shown in Figures 5-10 and 5-11. Again the darker curves are computed earlier in the test and the lighter curves later.

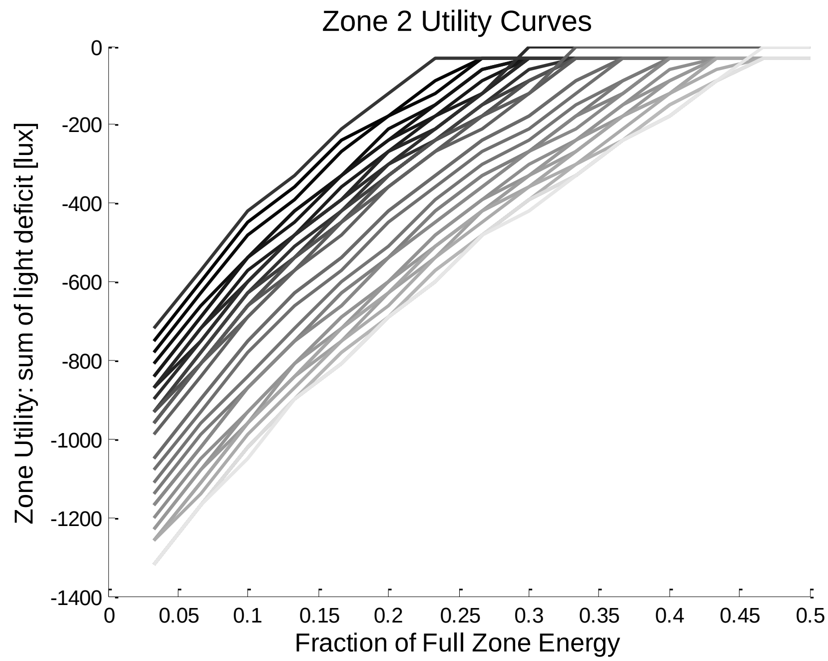

Figure 5-10: Test B Zone 2 Utility Curves

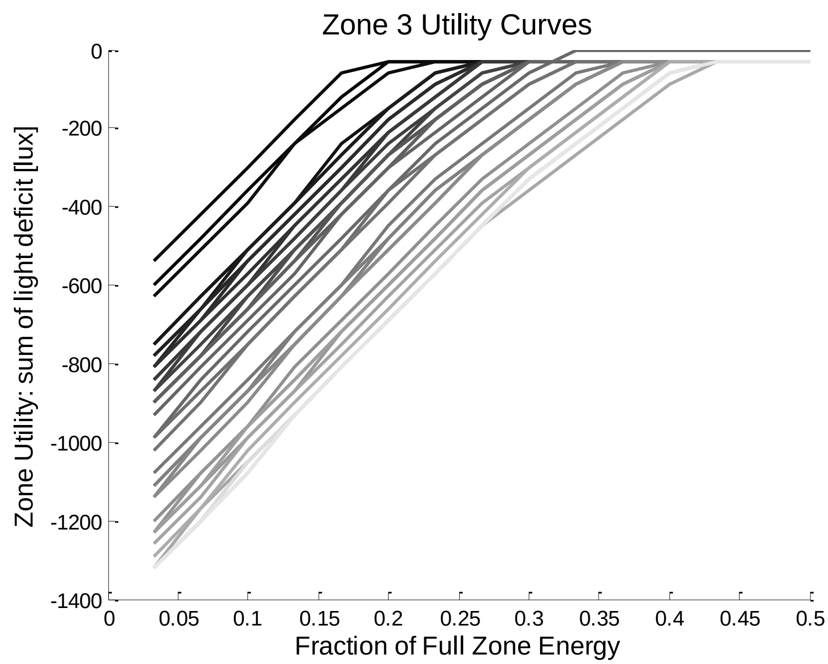

Figure 5-11: Test B Zone 3 Utility Curves

In both zone 2 and zone 3 there is a progression of the curves from left to right during the test. This is due to the fading sunlight during the sunset span of the test. As with

Test A, the diminishing sunlight level causes a corresponding drop in utility for a constant level of energy used and more energy is required to maintain the requested light level.

The larger off-diagonal terms in the influence matrix for zone 2 mean that an increase in one light setting significantly affects more than one lamp. This effect causes a steeper initial slope to the curve indicating a faster increase in utility per energy applied from zero and a slower rate of increase at the top. This is advantageous in an energy constrained scenario as the utility level traces the shallower at the high end of the curve. This information is important for design consideration in that it may be important to consider not only the lamp with the largest light contribution to a work surface but also to consider the secondary contributions as well in the layout to create an efficient system.

5.3.3 TEST C

Test C is conducted to show a longer-term test run of the system. The same basic setup is used for this test, but the sensors are again moved to show a new set of relationships between sensors and lamps. This test runs through a full 24-hour period to show the full cycle of performance over a typical day. This test runs from 12:00 AM PDT on August 29, 2011 to 12:00 AM PDT on August 30, 2011. August 29th was a partly cloudy day with a sunrise time of 6:36 AM and a sunset time of 7:42 PM PDT.

5.3.3.1 Influence Matrices

For this test the diagonal components are again significantly larger than the off-diagonal elements. As shown in Tables 5-5 and 5-6, the influence of each lamp to the sensor of the same unit ranges widely due to purposeful variation in relative sensor location and orientation. The variability is designed to show difference in resultant settings based on the varying influence levels.

Table 5-5: Test C Zone 2 Influence Matrix [lux]

	Lamp 2 at 100%	Lamp 4 at 100%	Lamp 5 at 100%
Sensor 2	1110.0	62.9	29.8
Sensor 4	11.6	1333.3	18.2
Sensor 5	21.5	44.7	1556.6

Table 5-6: Test C Zone 3 Influence Matrix [lux]

	Lamp 3 at 100%	Lamp 6 at 100%	Lamp 7 at 100%
Sensor 3	1290.3	23.2	110.8
Sensor 6	21.5	1028.9	61.2
Sensor 7	94.3	31.4	1205.9

5.3.3.2 Light Sensor Readings and Lamp Control Settings

With the system running at nighttime, the performance of the system in the absence of external influence is considered. Both the light sensor readings and the lamp settings remain constant during this time, as would be expected without perturbations to the system. The variation in nighttime lamp settings is due to the orientation of the sensors with respect to the lamps.

Figures 5-12 and 5-14 show the system response to a large gradation in natural light influence in the system and show the lamps dimming to compensate for this additional light. The degree to which the sensors and lamps are sensitive to the change in daylight availability is relative to the orientation and the location of the sensors with respect to the external windows in the room. Figure 5-13 shows the calculated visible light level in the system for the duration of the test which tracks the preferred 500 lux level as the total light level measured by the voltage sensors varies throughout the day. The pronounced peak in visible light level for sensor 6 in the middle of the data set corresponds to the provision of natural light in excess of the preferred light level.

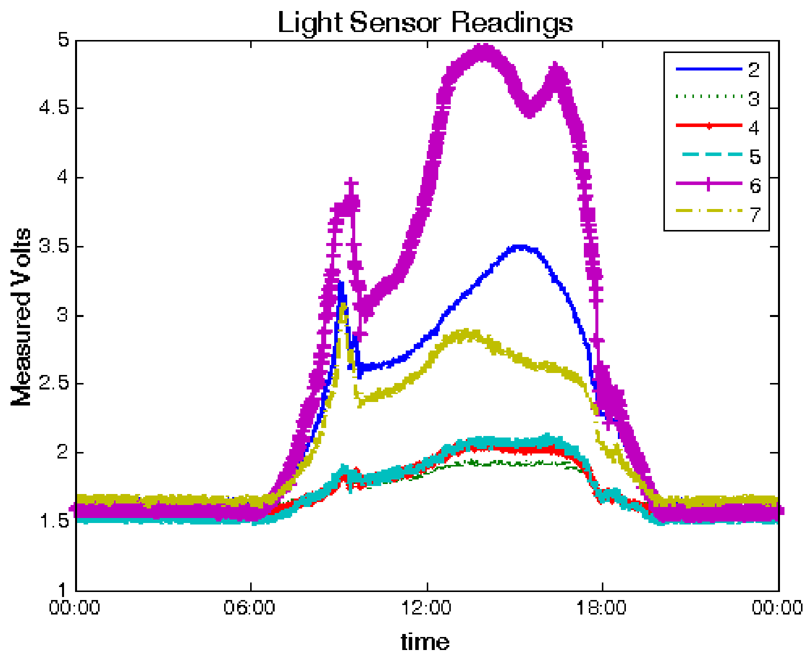

Figure 5-12: Test C Sensor Readings

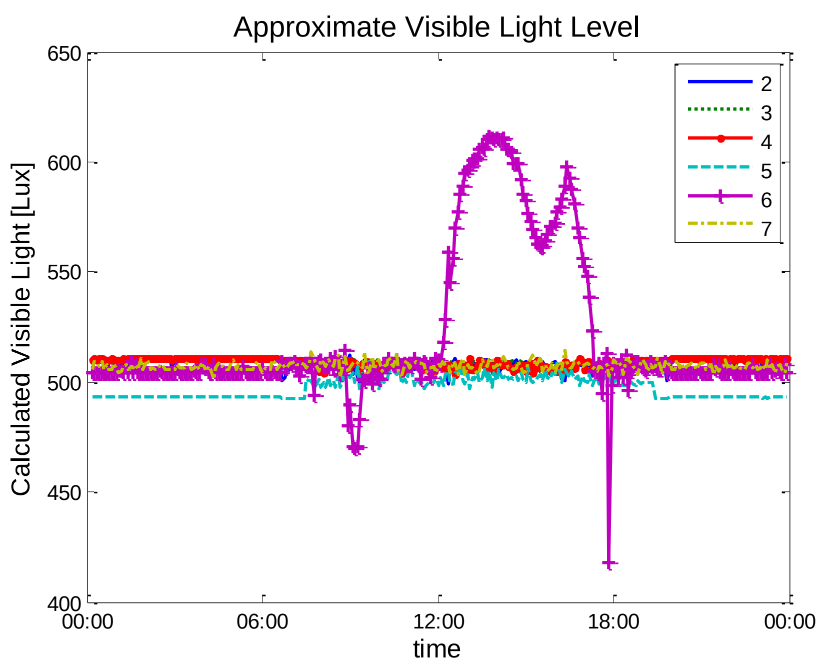

Figure 5-13: Test C Approximate Visible Light Levels

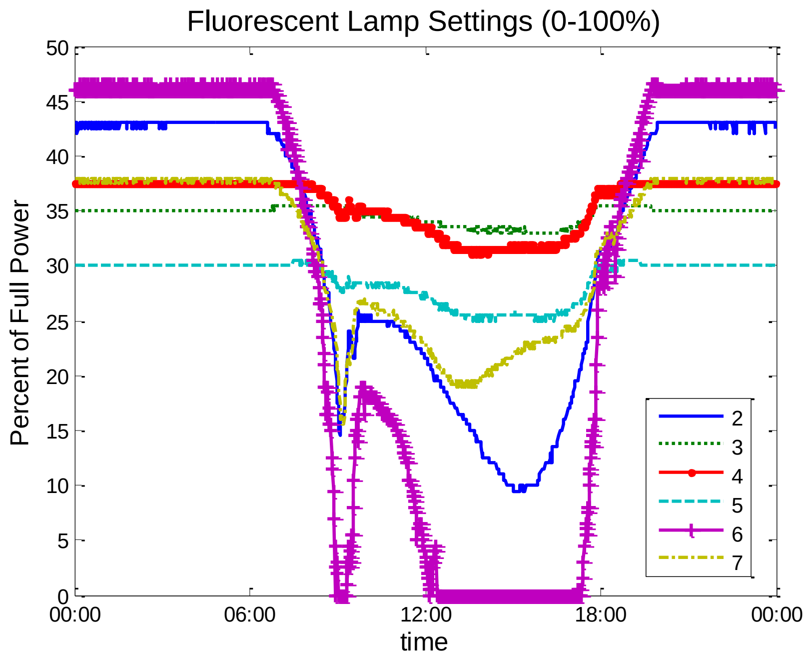

Figure 5-14: Test C Lamp Settings

5.3.3.3 Utility Curve Progression

For clarity, the utility curve progression for the two zones is separated out into two distinct time increments so that the direction of the progression is clear. The first interval includes 12:00 AM to 2:00 PM and is shown in Figures 5-15 and 5-17 and the second includes 2:00 PM to 12:00 AM and is presented in Figures 5-16 and 5-18. Because the natural light level first increases during sunrise and then decreases as sunset approaches, if the plots were not separated the trending would be difficult to see. For each plot the darker lines occur toward the beginning of the interval and the lighter lines occur toward the end. As shown in these figures, over time interval 1, the curves shift upward and to the left as more natural light is available. Time interval 2 shows the reverse of this trend as it takes place over the afternoon and evening hours as the sun lowers in the sky and declining daylight is available.

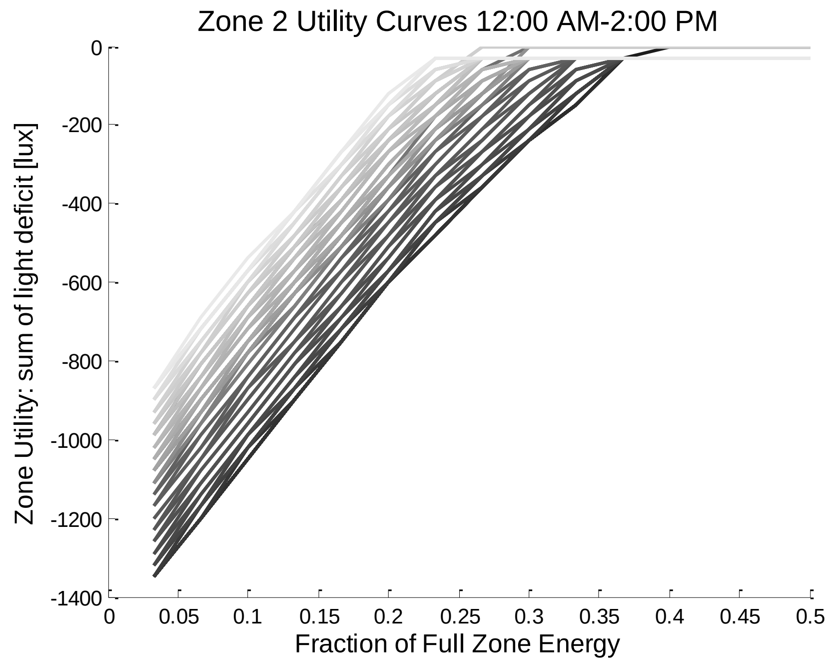

Figure 5-15: Test C Zone 2 Utility Curves for First Time Interval

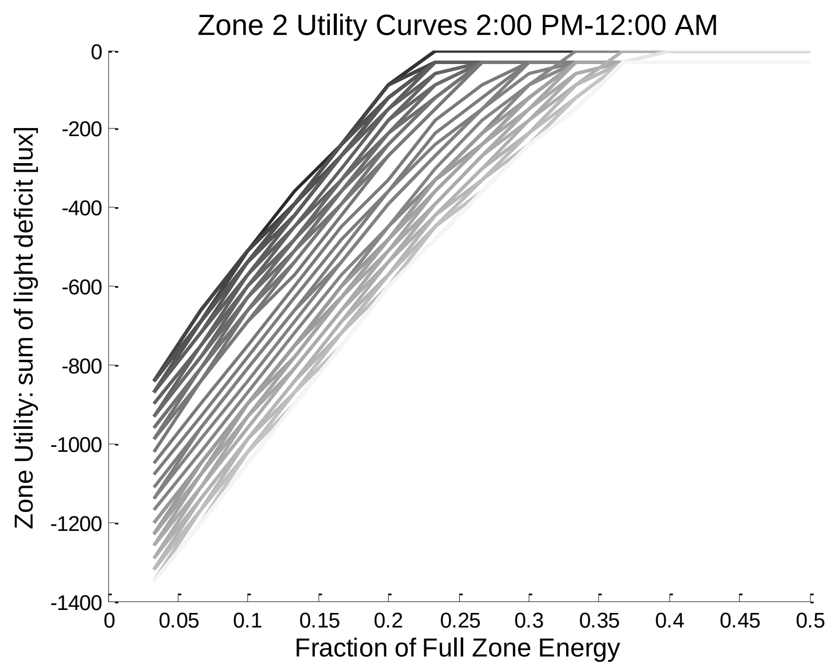

Figure 5-16: Test C Zone 2 Utility Curves for Second Time Interval

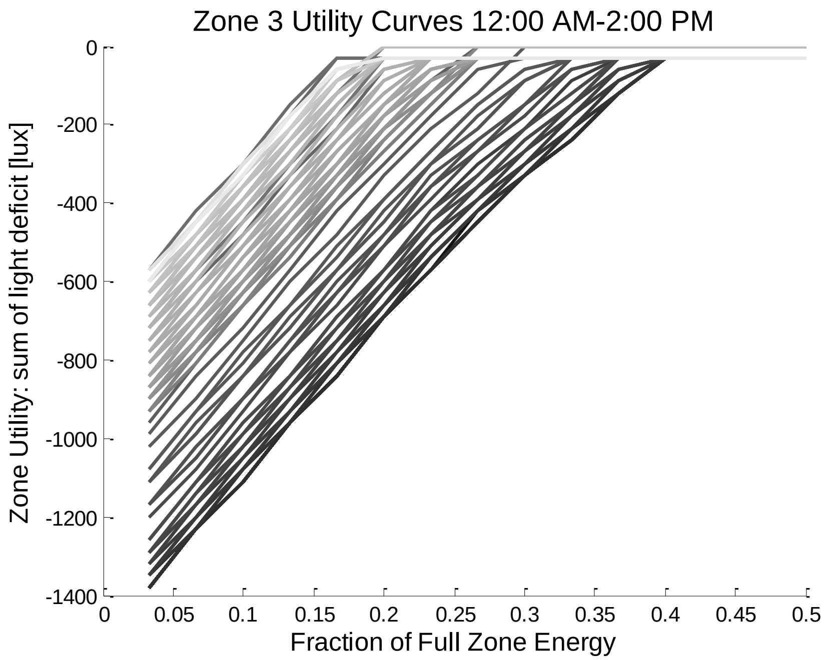

Figure 5-17: Test C Zone 3 Utility Curves First Time Interval

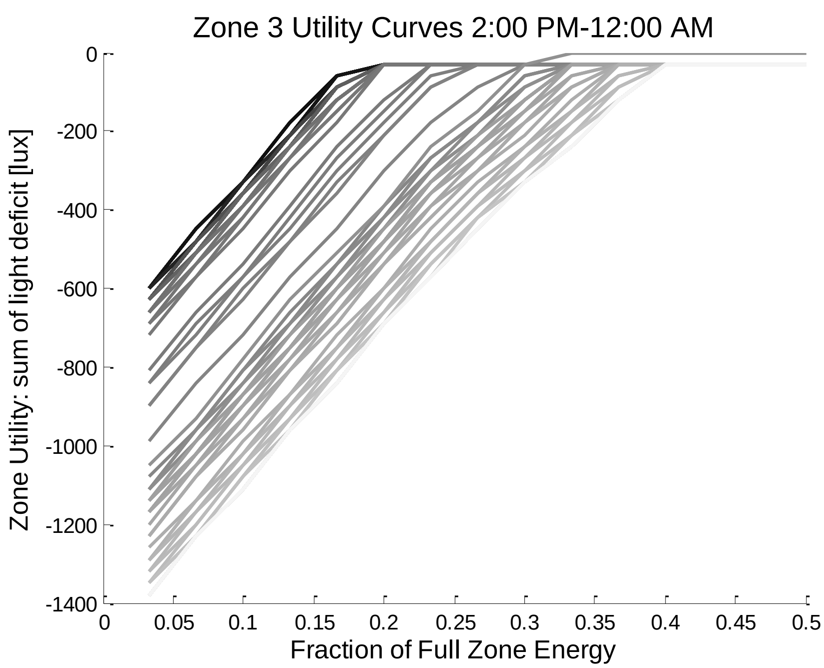

Figure 5-18: Test C Zone 3 Utility Curves Second Time Interval

5.4 IMPLICATIONS OF THE LABORATORY TEST

The purpose for the design and implementation of the physical test is to demonstrate the feasibility of creating a system that runs the tiered resource allocation scheme detailed in Chapter 4. Physically implementing this system evokes many issues that would not be apparent from a computer simulation alone. While the computer simulation does not have to manage non-idealities, the physical system is forced to endure many such issues. Through finding solutions to handling these problems, the potential for this system in a real world simulation is demonstrated.

As implemented, the system demonstrates resiliency to communication errors and to the impact of errant light entering from other zones. A more robust communication protocol would be desirable for use in a real building environment as light levels are important to occupant productivity. However, no system is without errors and as such the ability of this system to rebound quickly to compensate for errors demonstrates an important inherent robustness. Because the two zones are not entirely separated for the experiments, some light is shared between zones. This is not ideal because the effects of zones on one another are not directly accounted for within the system. The extra ambient light is treated as additional sunlight which is inaccurate as it does not share the same spectral distribution as a sunlight contribution would. However, this issue did not appear to have a significant impact on the performance of the system during the experiments which is positive as realistic building zones would likely have some light bleeding between zones. Use of light sensors well matched to the photopic spectrum would eliminate the miscounting of the additional artificial light contribution.

Running these experiments reinforces the importance of separating the sunlight and fluorescent light contributions to the sensor readings. For any sensor with a spectral sensitivity distribution not identically matching that of the human eye, treating these components separately is essential. The goal is uniformity of visible light and the change in voltage readings in the sensors over the course of the tests shows how the contributions of the two types of light are interpreted differently by the sensor. The amount of variation over the course of the test shows the error that would be inherent if the components were treated equally.

The importance of a proper commissioning cycle was also underscored in these tests. Due to the arrangement of the test setup, it is possible to situate the sensors such that they are saturated by the lamps when they are fully on. This causes a problem during setup in that the sensor reads its maximum value for the "fully on" state of the lamp but this is not an accurate basis for the influence matrix as the true sensor reading should be higher. This invalidates the calculations for the remainder of the test. Careful selection of sensor sensitivity and sensor location can ameliorate these concerns. In a scenario where saturation of the sensors at full light output is of concern, the commissioning sequence should be reconstructed to test the readings of the sensors at multiple settings of individual lamps to ensure a saturation value is not recorded in the influence matrix.

In addition to concerns about sensor sensitivity for commissioning, it is important that the selected sensor does not saturate at a voltage output level less than the voltage output for the setpoint provided entirely by fluorescent light plus the voltage output for the setpoint provided entirely by sunlight. The reason behind this design requirement is that if the lights are fully on, the sensor still needs to be able to measure the available sunlight. Because the system only cares about the level of sunlight up to the requested level, the sensor needs to be able to measure sunlight up to that value. Failing to meet this requirement should not disable the system but could result in prolonged response times to changes in sunlight level.

The system as implemented in the laboratory was timer driven in that all updates at the zone and building level were initiated by timers. This design selection was made in accordance with the desire for data collection and to evaluate the consistency of the light settings over time. As is shown in many of the figures from the experiment, often little has changed between updates. To design a more efficient communication system, the zone level updates could be rescheduled to occur when a sensor notices a prescribed degree of change from the previous state or when occupancy status changes. Building wide reallocation could be rescheduled to occur when a demand response request is received or when the energy needs of a zone change above a threshold value and the zone requests an update. Restructuring the communications thusly would decrease the amount of active time of the units and thereby cut the energy overhead of the control system.

5.5 SUMMARY

The implementability of this system is thoroughly demonstrated by the test results described in this chapter. A circuit board was designed to demonstrate that with few off-the-shelf components, a functional electronic hardware system can be developed. The software to control the hardware system was developed to demonstrate the sunlight compensation abilities of the optimization and resource allocation scheme. The hardware and software system were evaluated on a laboratory-scale test bed in a sunlit room and the system was required to track a set light level over the course of several hours on multiple days. In all three experiments the system was successfully able to compensate for the sunlight and maintain communications throughout the experiment duration. Test A shows that the system can continuously update to reflect external lighting conditions. Test B demonstrates that the system can recover from transmission errors and that more variation in lamp settings is employed when the sensors are positioned to take advantage of sunlight. Test C shows that the system can provide reliable tracking over a full day span. The success of this prototype system indicates the potential for design of a real-world, commercial system.

Chapter 6.

BUILDING-SCALE SIMULATION

A typical commercial building is simulated to evaluate the potential energy performance of the light control and optimization system in a real building. While the prototype system discussed in the previous chapters was developed to demonstrate the implementability of the system and general performance, the simulation is designed to assess the realistic implications of the implementation of a building-wide system. Assessing the energy and occupant preference matching performance of the resource allocation system requires the ability to model the system response to changing natural light, occupancy, and occupant preference parameters. Detailed lighting simulation software, most notably RADIANCE© designed by Lawrence Berkeley National Laboratory, and building energy use simulation software for building design, such as EnergyPlus and DOE-2, are available, however these programs were developed for use by designers in creating building spaces and estimating annual energy use but are not designed for the evaluation of new control systems. The lighting simulation software performs detailed analysis to determine the lighting scenes under specified conditions in a particular space. The model output is highly accurate, but simulation is typically performed on a single room requiring detailed information about all surfaces in the space and is very computationally intensive due to the precise ray-tracing modeling and rendering. This type of software is insufficient for application to assessing the performance of a new lighting control system due to its level of detail, the focus on individual spaces within a building, and the inability to integrate the optimization algorithms directly into the simulation. Building energy simulation software is less detailed in its consideration of lighting and typically provides annual lighting energy use with or without a built in lighting control system. Similar to the detailed lighting software, the energy modeling software does not allow for the integration of a new type of control system. Both types of software allow for the consideration of lighting controls; however, the control decisions are selected based on internally defined algorithms. Because the available software packages are designed for different purposes, a simulation environment specific to the needs of evaluating the performance of a new control system is developed and discussed in this chapter.

For the purposes of the simulation, a computer model of a three story office building is created. The building is divided into 12 zones including open-plan shared office spaces, hallways, and enclosed private office spaces. The building can be relocated and reoriented to demonstrate the performance under varying external conditions. The simulation uses a year-long set of recorded irradiance data to provide natural lighting information and stochastic models of occupant preference and behavior to incorporate occupant interaction with the building.

6.1 MOTIVATION

The building-scale simulation is created to demonstrate the effectiveness of the tiered system in meeting lighting demand and minimizing energy use. A typical office building with a rectangular footprint is designed to examine the expected behavior of the system in a real building. The implementation of this simulation shows how the zones interact with one another, how inputs from the building manager affect system performance, and how building performance can be monitored over time. This building model and simulation is constructed to give a realistic view of the performance of the tiered resource allocation system in a full building implementation. The model is not intended to represent any specific building but instead to represent general, plausible conditions to test the system responsiveness to demand and potential as an energy savings tool. The full-year simulation using recorded irradiance data shows the interaction of the building with the surrounding environment under a range of external conditions as the available natural light varies over the year with weather conditions and relative sun position. The lighting models used for converting this data to internal natural light contributions at sensor locations are approximate methods typically used for preliminary design estimates. These methods provide realistic values for internal natural light levels with reasonable computational requirements for the quantity of time steps used in the simulation. Use of both a stochastic occupancy model and light level preferences selected from a distribution matching reported results from other studies incorporate typical occupant behavior dynamics. The results from this simulation are an indication of the general effectiveness and behavior of control and optimization system. The utility of the system in any specific building is dependent upon the individual building characteristics and this model is put forth as one standard example. A more detailed analysis for a specific building would require knowledge of the anticipated occupant behavior as well as a detailed analysis of daylighting and artificial lighting conditions.

6.2 Simulation Design

The simulation model is composed of several separate modules. The modules are independently designed such that each module can be individually changed without affecting the other components of the model. The building module houses all physical design aspects of the building. Behavior of the building occupants and the light levels they prefer are specified by the occupancy module and occupant preference module respectively. The artificial lighting module defines the relationship between the lighting fixtures and the sensors in the building. The natural light contribution to the system is defined in the natural lighting module. All model components are assessed at 5 minute time steps for a full year. The occupancy, preference, and natural light modules are computed for all time steps in advance of running the simulation to decrease simulation run time and to allow the same parameters to be used for each simulation run for direct comparison. The two-sided linear formulation is used to determine the zone utility curves and to determine the optimal light settings within the zones.

6.2.1 BUILDING MODULE

The building module defines the physical parameters of the building model including the geometry of the space and the locations of all sensors and lamps. The module also defines the relevant design parameters of the building: building orientation and size, zone assignments, room depth and internal building layout.

The building is designed as a three story, rectangular structure measuring 50x100 feet. The long edge of the building is oriented along the east-west axis. The floors have identical layouts with the exception that the conference room space on the second and third floors is designated as a reception area on the first floor. Figure 6-1 shows the floor plan layout of the building and the zone assignments of the spaces for each floor. For clarity, the floor plan is reproduced in Figure 6-2 along with the 2-dimensional location of the light sensors and the centers of the lamp fixtures. The lamps are affixed to the ceiling and the sensors represent light levels at the desk surface, or floor surface in the case of the hallway sensors.

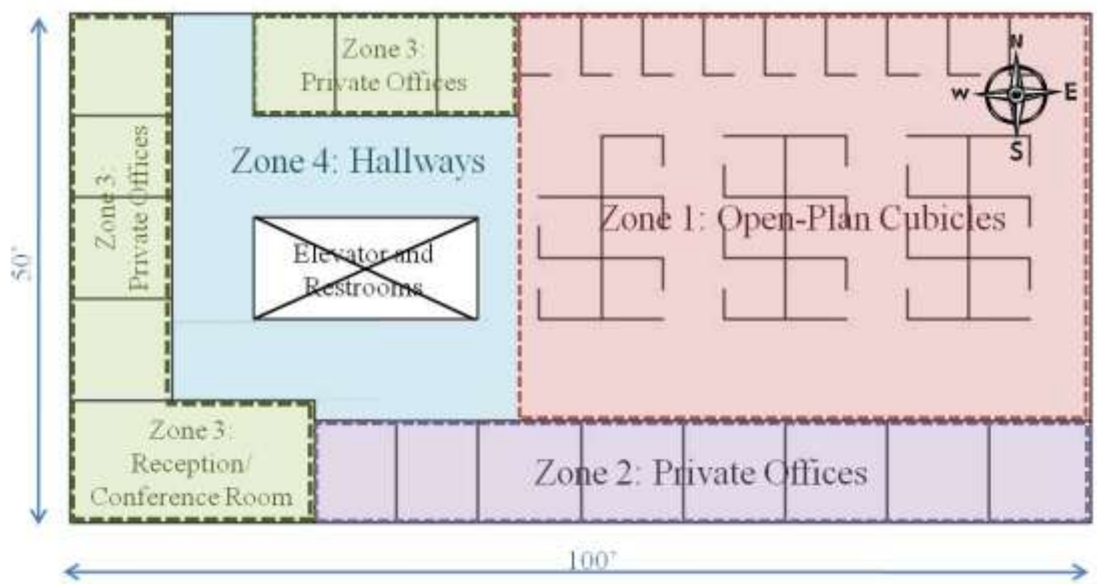

Figure 6-1: Simulated Building Floor Plan with Zoning

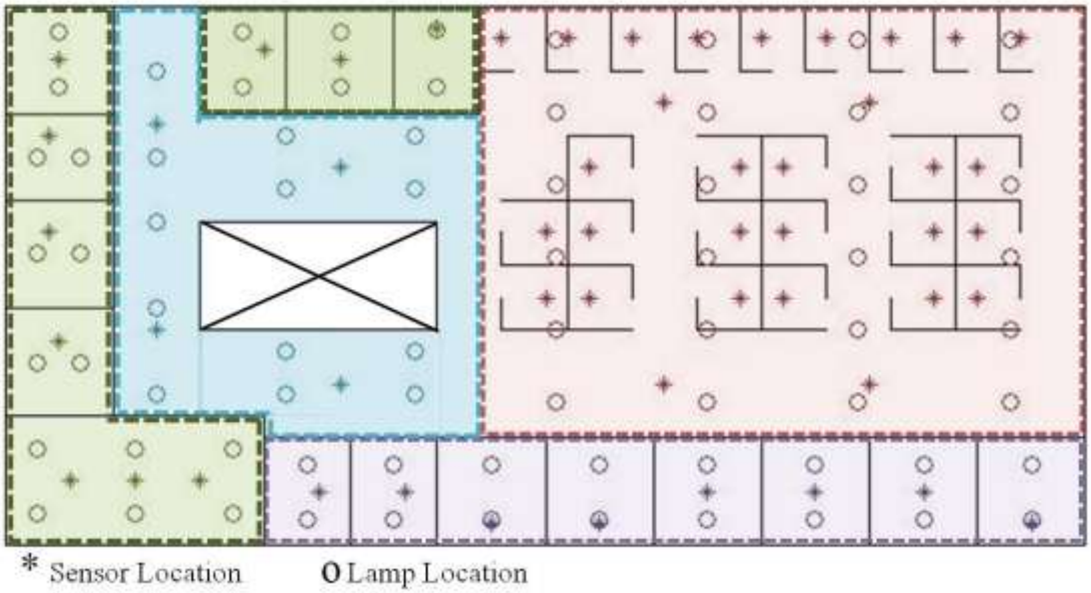

Figure 6-2: Sensor and Lamp Locations for Simulated Building

The bathroom and elevator space is not considered as part of the simulation scenario because the lighting needs in these locations do not require the incorporation of occupant-specific light level preferences. These spaces are internal to the building and therefore do not have significant natural light contributions to compensate for and thus do not require a more complex optimization scheme to meet demand.

Table 6-1 shows the number of potential occupants (equivalent to the number of desk top sensors), sensors, and lamps in each zone. Each floor contains the same four types of zones but the zone types themselves widely range in size and composition to illustrate the applicability of the tiered system to a variety of zone definitions. The building has a total of 156 sensors and 219 control points.

Table 6-1: Zone Design Parameters

Zone Number	Number of Occupants	Number of Sensors	Number of Lamps
Zone 1	26	30	24
Zone 2	8	8	16
Zone 3	7	10	20
Zone 4	0	4	13

6.2.2 OCCUPANCY MODULE

Occupancy information is used by the system to determine which constraints in the optimization equations are considered. When a desk or space is unoccupied, the corresponding constraint equation is effectively removed from the system by setting the participation weighting factor on the respective error terms to zero. Occupancy is defined as a binary 0 or 1 value with 0 indicating unoccupied and 1 indicating occupied, and these values are used directly as the participation weighting factors p_i as discussed in Chapter 4.

The model defines desk occupancy in accordance with the stochastic LIGHTSWITCH model (Newsham et al., 1995) which serves as a standard model for office occupancy in building simulation (Reinhart, 2004; Bourgeois et al., 2004; Mahdavi & Pröglhöf, 2008). The occupancy of the other building spaces is determined based on the known occupancy of the desks. The hallway sensors are assigned occupied status if any desks on the floor are occupied. The first floor entrance is considered occupied from 7:00 AM to 8:00 PM each day and the conference room spaces are given a 50 percent chance of being occupied for all time increments during the same hours of the day.

The LIGHTSWITCH model uses defined probability distributions for probability of arrival within a specific time increment and probability of departure within a specific time increment given not previously departed to determine the working hours of a specific employee. The arrival and departure rates are 100 percent over their specified, respective windows. The arrival and departure windows do not overlap, avoiding the need for additional conditional statements. This model does not take into account vacation or sick days, holidays, or occupied time for cleaning staff at night. The parametric values for the distributions used in this simulation are adapted from the LIGHTSWITCH empirical model. The probability of arrival over time is modeled as a truncated normal distribution. The probability of departure in a time interval given no previous departure is set at .025 from 5 PM to 8 PM and .5 for each time increment thereafter until 1 AM where it is instantaneously equal to 1, forcing a reset of the occupancy model for the following day. The probability distributions used for determining occupancy are shown in Figure 6-3.

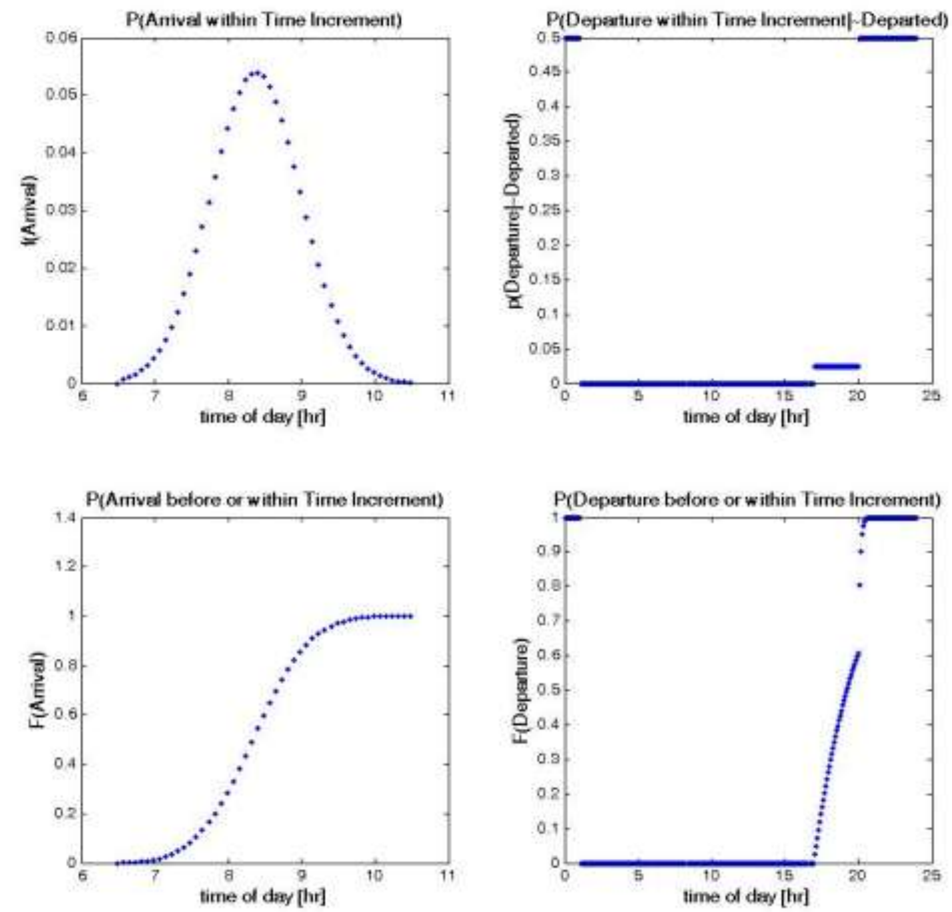

Figure 6-3: Occupancy Probability Distributions (adapted from [Newsham et. al, 1995])

Time-of-day dependent probabilities are defined for whether the occupant is or is not at his or her desk based on whether he or she was at his or her desk in the previous time increment. These time-dependent probabilities are plotted in Figure 6-4. The probability of an unoccupied desk for any time increment following an unoccupied time increment is higher than that for one following an occupied time increment. The probability is time-of-day dependent as lunchtime represents the most likely time an occupant might spend away from his or her desk.

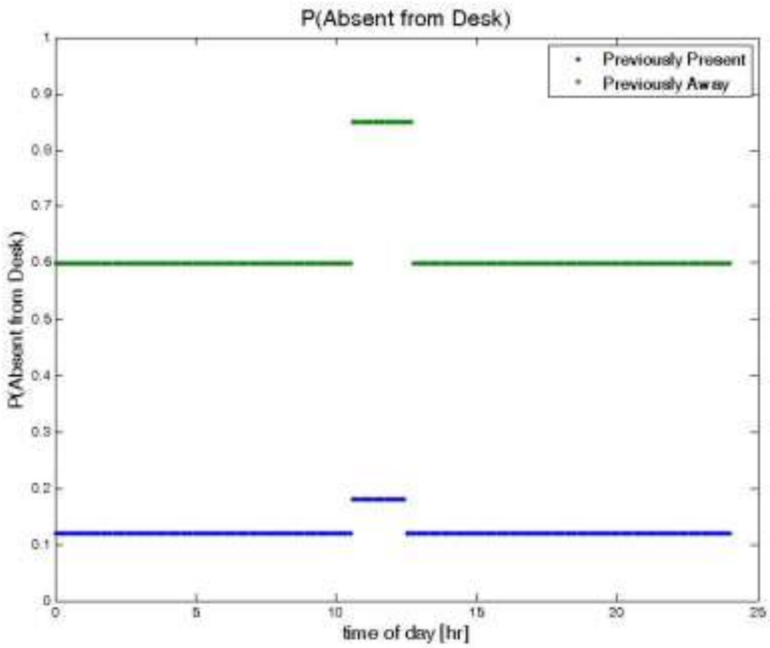

Figure 6-4: Probability Distribution for Temporary Absence (adapted from [Newsham et. al, 1995])

6.2.3 Occupant Preference Module

Simulated occupant preferences are used to determine the necessary light level at the sensor locations. The preference distribution is defined by a beta distribution with mean and standard deviation matched to the study results by Veitch & Newsham (2000).

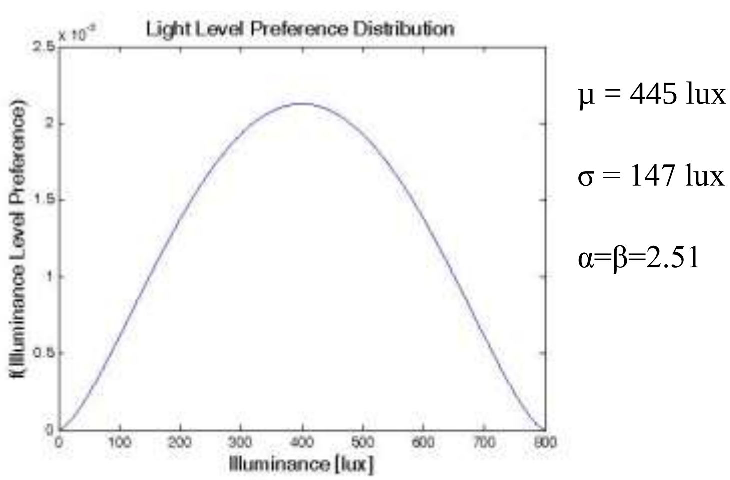

Figure 6-5: Occupant Illuminance Preference Model

The simulated preferences for the occupants at their desks are assigned in accordance with the distribution above. For each associated sensor, a time increment is randomly selected each day for a preference change. If the time increment falls within the working hours of the respective occupant, a new preference is selected at that time of day. If not, a new preference is established at the start of the next day. The maximum of one preference level change per day is based on the results of a several month experiment conducted in an office building that showed over time the average number of changes made to an installed personal control system was less than one per day per occupant (Newsham et al., 2009).

The illuminance level for the hallway illuminance preference is specified to 100 lux. The first floor entry way preference levels are set to 200 lux. The distribution of illuminance preference in the conference room space is uniform from 100 to 500 lux to allow consideration of the multiple uses of these types of spaces. Daily changes to the specified illuminance request for the conference rooms are processed the same way as for the office spaces previously discussed. These values are selected in accordance with the Chartered Institution of Building Services Engineers Code for Interior Lighting (CIBSE, 1994) and the Illuminating Engineering Society of North America Lighting Handbook (IESNA, 2000).

6.2.4 Artificial Lighting Module

The artificial lighting module determines the influence matrices for use in the optimization algorithm. A real-world implementation of this system would construct the influence matrices directly from sensor measurements; however, in the absence of real data for this model, the isolux method is selected for creating the simulation influence matrices. The isolux plots for the typical two 35 Watt T16 lamp with a 3000K color temperature fixture (Sylvania Rana Surface 50275) are used. The isolux method provides precise values for illuminance at all locations on an idealized work plane ignoring both reflected light contributions and shadows. Information regarding the relative locations and orientations of sensors and lamps is used to select the appropriate values from the isolux plot for insertion in the influence matrix. This idealized estimate is considered acceptable for this model as the goal of developing the model is to create a realistic, general building environment to evaluate the potential of the control system, not to create an exact representation of any specific existing building.

6.2.5 Natural Lighting Module

The natural lighting module simulates the sunlight contribution to the model. The module relies on annual irradiance datasets collected from locations across the United States to simulate natural light contributions in a diverse array of environments.

6.2.5.1 Datasets

Cooperative Networks for Renewable Resource Measurements (CONFRRM) solar energy resource data from the National Renewable Energy Laboratory are used (NREL, 2010). The datasets are comprised of solar irradiance measurements taken at 5-minute intervals over several years. Full year datasets are necessary because sunlight contributions are subject to seasonal variation. The recorded measurements of interest are: global horizontal irradiance (GHI), direct normal irradiance (DNI), and diffuse horizontal irradiance (DHI). The data set for Elizabeth City, 2008 is selected for use in the simulations due to its completeness and data quality (NREL, 2010). February 29th is removed from the 2008 dataset to give a typical 365 day year. The data set for El Paso, 1999 is selected for use as a comparison case (NREL, 2010). The El Paso data set is of poorer quality as measured by the associated data quality flags as provided by CONFRRM. In order to have a complete data set for the year, erroneous data points are replaced with the data point from the previous day at the same time. In this way, the seasonality of the data is preserved.

6.2.5.2 Computation of Internal Light Levels

From the CONFRRM data and the relative sun position, the sky clearness is assessed (Perez et al., 1990) and vertical illuminance is determined based on the sky conditions. The horizontal illuminance is calculated directly from the horizontal irradiance data. For overcast sky conditions, Krochman's formula (Stein et al., 2006) with a horizontal to vertical illuminance ratio of 2.5:1 is used. For a clear sky, the vertical illuminance is determined from empirical curves relating vertical surface illumination to solar altitude and bearing angle (Stein et al., 2006). The clear sky values are scaled up by 15 percent for partly cloudy conditions to account for additional reflected light from cloud edges (Stein et al., 2006).

Natural illumination is introduced to the building via a row of windows at each floor level that wraps around the perimeter of the building. The horizontal and vertical surface illuminance values are translated into natural light contributions at the interior sensor locations using the IESNA Lumen Method (IESNA, 2000). The Lumen Method is an approximation method initially developed to aid in the early stages of building design. The method provides values for daylight level at points of varying depth through the space at workspace height based on window height and width, window transmittance, light loss factor, and the net glass area on the building façade. This approximation method is considered adequate to provide the general daylighting behavior for the model. Typical values for the input variables are selected for the model building and are listed in Table 6-2. Results from (Reinhart, 2004) and (Newsham, 1994) are used to incorporate window blind activation behaviors. Reinhart shows blind activation when direct solar irradiance exceeds 50 Watts per square meter and Newsham claims window blind activation reduces visible light by 80 percent. These values are incorporated into the model preventing the underestimation of artificial light use that would occur if the blind action were ignored.

Table 6-2: Daylight Parameters

Parameter	Value
Window Transmittance	.8
Light Loss Factor	.9
Net Glass Area	.92
Window Height above working surface	4 ft.
Window Width	6 ft.

6.3 Simulation Results

Energy performance of the system is evaluated solely with respect to energy used for lighting and all subsequent references to energy use are in reference to lighting system energy use only.

6.3.1 Annual Energy Performance: Fully and Partially Implemented System

To evaluate the effectiveness of the light control and optimization system in energy conservation, the fully implemented control system model is compared to a base case with standard preference values, no distributed occupancy sensing, and no daylighting compensation. The base case scenario assumes centralized occupancy sensors for the shared spaces and individual office and conference room occupancy sensors. The occupancy of the large, shared office spaces is therefore limited to either fully occupied or fully unoccupied; fully occupied status is determined by the presence of at least one occupant in the space in the occupancy model used for the fully implemented case. The no distributed occupancy sensing scenario is also representative of the case of no occupancy sensors in the building but with perfect use of room light switches. For the base case, the standard light preference level is set to 500 lux, the design standard for the desk surface and the natural light contribution is set to zero for all time steps to represent a condition lacking daylight compensation. The linear programming algorithm is used to set the lamps to the best-fit lamp settings given this base case scenario to provide a baseline for comparison. The daily building energy use for the fully implemented case and the base case are shown below. In Figure 6-6, weekend days have been removed as building use is assumed to be zero.

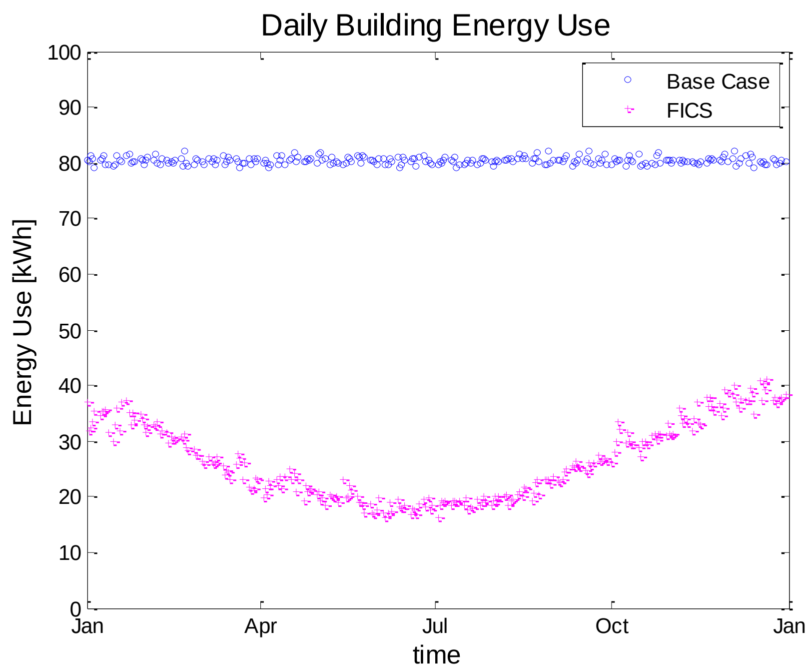

Figure 6-6: Comparison of Daily Energy Use for Fully Implemented Control System (FICS) and Base Case

Figure 6-6 shows the base case requires a consistent level of energy throughout the year while the FICS case varies seasonally based on the quantity of available daylight. The difference in magnitude of energy use required for these two cases is evidence of the energy saving capabilities of the control system. The cumulative impact of these energy savings is tabulated below along with other considered cases.

Partially implemented cases are also considered with each individual system capability removed, one at a time, to evaluate the contribution of each piece of information to the overall total system performance. The removed capabilities are modeled using the analogous scenario from the base case. The annual energy use for each of these partially implemented cases is shown in Table 6-3.

Table 6-3: Annual Energy Use Comparison of Fully and Partially Implemented Control System, Elizabeth City, 2008

Implementation Case	Annual Energy Use (kWh)	% Energy Use Increase over FICS
Fully Implemented Control System (FICS)	6903	--
Centralized Occupancy Sensing	7000	1.4%
Standard Light Level Preference	7688	11.4%
No Daylight Compensation	19,187	177.9%
No Control Implementation (Base Case)	20,925	203.8%
Reference (1W/ft^2 for 8 hour weekday)	29,200	323.0%

The 1 W/ft^2 reference is provided as a comparison of the performance of an office building designed according to the Building Area Method standard from ASHRAE/IESNA 90.1-2004. This reference point refers to a usage of one watt per square foot of space throughout the building, excluding the spaces not accounted for in the simulations, for 8 hours per weekday. This reference does not account for extended working hours, late night custodial services, or weekend building use and is intended as a relative reference point for the simulation values. The difference in energy use between the reference case and the base case is due to the specifics of the base case design. For the base case, the lamp settings are selected using the same linear programming algorithm as is used in the other cases, which means the lamps are set to dimming levels which best meet the required light levels. This tailored lighting scenario reduces building energy use. Occupancy controls are also implemented for all private offices and conference room spaces which yield additional savings from those spaces when they are unoccupied. The central occupancy sensors in the shared spaces also reduce energy use by preventing lights from being left on when no one is present. The reference case, by contrast, assumes a consistent 1 W/ ft^2 usage throughout the 8 hour day.

The simulation results show a relatively small savings from the distributed occupancy sensing. This outcome is partly attributable to the specifics of the centralized occupancy sensing implementation and the occupancy model. The change in occupancy as viewed by the system exclusively pertains to the large, shared office spaces. The small offices, conference rooms, hallways, and reception area retain their individual room occupancy sensors for both scenarios. As the occupancy model yields finite arrival and departure windows, much of the work time remains well staffed throughout the day, with a median occupancy in the shared office spaces of 67 percent for all time steps with at least one building occupant, limiting the impact of the individual sensors. The large, shared office spaces are also well served by daylight throughout the day which lowers the overall use of light limiting the impact of the occupancy on the system performance. In a shared space with little or no daylight, the influence of individual occupancy sensing would be greater. For low occupancy levels with low daylight contributions, the system itself is limited with regard to energy use minimization for shared office cases where the occupants are physically spread out in that the light use for each occupant will be provided largely from distinct lamps requiring a higher average energy use per person than for occupants that are closer who share influence from common lamps. While the energy savings for this case were relatively small, occupant preference performance is also impacted by occupancy sensor distribution and should improve with the larger degree of freedom the system has available to meet the demands of the occupants that are present.

The incorporation of individual preferences results in an energy savings of 10.2 percent over the standard preference case. This savings occurs because the standard recommended light level is greater than the preference of the average occupant as defined in the preference model. For any specific building, the potential for savings is variable with the preferences of the occupants and the installed lighting. For older buildings with lighting designed for mostly paper-related tasks but which now have workers performing more computer-related tasks, significant savings may be achieved as computer workers generally prefer lower light levels to enhance screen contrast. However, a building housing an older segment of the working population would have higher lighting needs than that of the general population resulting in reduced savings. The impact of preference incorporation to meeting occupant demand is discussed in the next section.

As is evident from Table 6-3, daylight compensation represents the largest contribution of lighting energy savings, with a 64 percent lighting energy savings as compared with the no daylighting case. With the large windows on the façade of the building, the narrow floor plan, and open interior space allowing for deep daylight penetration, daylight provides a significant portion of the necessary lighting throughout the building. As an idealized estimate of interior daylight provision with the exclusion of obstructions, the model may overestimate the quantity of daylight provided thus contributing to an overstatement of the energy savings for the building. However, the daylighting model also excludes the contribution of internally reflected light due to the obstructions as well which serves as a mitigating factor in the estimation. The specific energy savings potential for a building due to daylight compensation is dependent on the architectural design of the particular building. Buildings with more internal divisions restricting daylight penetration and with wider floor plans taking occupants further from windows will not achieve the same fractional savings as the type of building designed for this simulation. With new and emerging technologies designed to bring natural light deeper into buildings, the potential for energy savings by using natural light to offset the need for artificial light is continually increasing.

The comparison between the base case and the FICS demonstrates the full potential of the system with all features implemented. From the values displayed in Table 6-3, the combination of all of these features in this building represents a 67 percent energy savings over the base case. As evidenced by the intermediate cases, most of this savings comes from the daylight compensation feature. Harnessing the available natural light and reducing the artificial light contribution accordingly saves resources when they are not needed and prevents flooding the occupants with additional, unwanted light. The large size of the energy savings is due to the extensive daylight availability throughout the building.

In addition to saving total annual energy use, reducing peak demands on the electricity distribution system is important to the longevity and cost-effectiveness of the aging electricity transmission and production system. Figure 6-7 shows a comparison of peak energy demands per day of the base case and FICS systems. The decrease in peak demands is of particular importance during the hot summer months when the grid is taxed by space cooling demands. As is shown in the figure, significant afternoon peak energy use reductions are achieved even during normal system operation as compared to a building lacking similar controls. Further reductions achieved via demand response capabilities are discussed in the demand response section. Seasonal variation in peak energy use is also observed as available natural light availability varies throughout the year, with more sunlight available during the long summer days. Afternoon peak energy savings for the FICS case are most prevalent in the summer months when savings are typically most critical as cooling loads are higher and total demand on the electricity distribution system increases. The potential for energy savings during the most critical time of day and year indicates the benefit of this system in supporting critical energy distribution system infrastructure in locations with heavy summer energy demand.

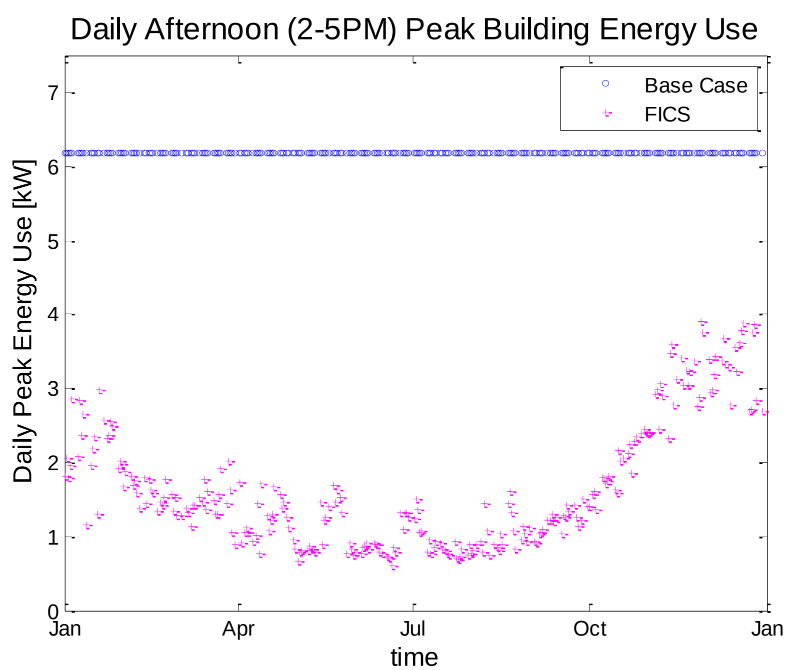

Figure 6-7: Daily Afternoon Peak Lighting Energy Use Comparison, EC

In order to examine the influence of climate on the building energy use and performance of the system, the daily building energy use results from the same simulation but with irradiance data from El Paso, Texas for 1999 are presented alongside the values from the Elizabeth City, 2008 simulation in Figure 6-8. The daily energy use values plotted in the figure show that the majority of the additional, climate-based energy savings occurs during the winter months when El Paso continues to have abundant sunlight availability due both to its weather and its latitude. Summary values for the total annual energy use are presented in Table 6-4. Results for Elizabeth City, North Carolina 2008 are repeated in the table to emphasize the comparison. The no daylight case and base case are the same for both as the only difference in the two simulations is the natural light contribution. The results presented in the table show less annual lighting energy use for the El Paso case than for the Elizabeth City case. This reduction in energy use is a result of the greater sunlight availability in El Paso.

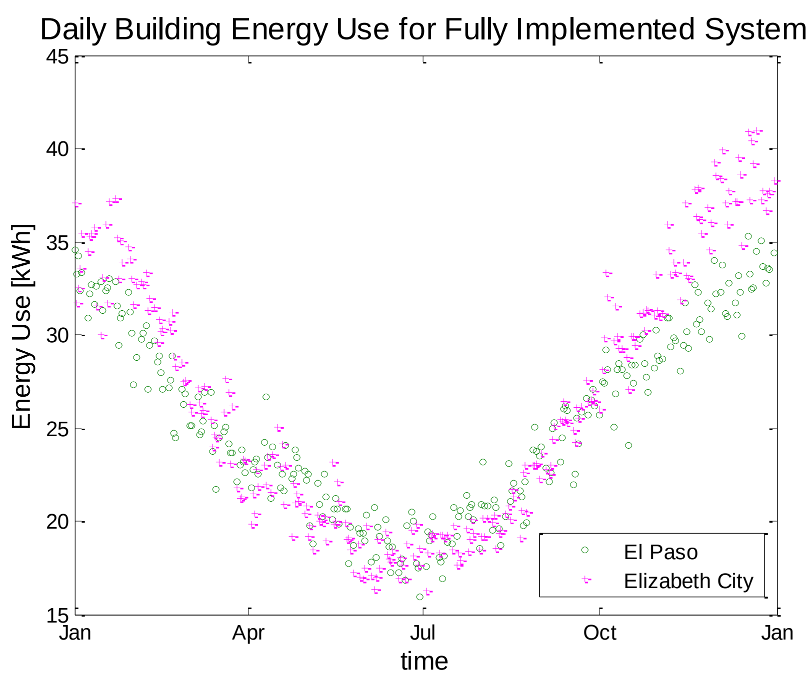

Figure 6-8: Daily Energy Use Comparison for Elizabeth City and El Paso

Table 6-4: Annual Lighting Energy Use Comparison of Elizabeth City, North Carolina 2008 and El Paso, Texas 1999

Implementation Case	Annual Energy Use (kWh)	
	Elizabeth City	El Paso
Fully Implemented Control System (FICS)	6903	6599
Centralized Occupancy Sensing	7000	6689
Standard Light Level Preference	7688	7381
No Daylight Compensation	19,187	19,187
No Control Implementation (Base Case)	20,925	20,925

The smaller difference in energy use seen in the summer months occurs partly due to the greater similarity in sky clarity during these months, but also because of the simplified window blinds model used. The window blind condition is activated in the model whenever incident irradiance exceeds a maximum value. Reducing the high peak daylight values during the sunny summer months by multiplying them by a reduction factor has the effect of bringing the simulated natural light levels closer together. From Figure 6-8 it appears the El Paso case requires more daily energy than the Elizabeth City case during parts of the summer which is due to the increase in necessary artificial lighting energy when the blind are activated, which occurs during a higher percentage of the day for this location. Implementing a window blind model that allows for blind settings other than fully open or fully closed would show increased energy savings for El Paso during the summer months and would thereby increase the relative annual energy savings shown in Table 6-4.

6.3.2 Annual Preference Tracking Performance

Preference levels are met by compensating for natural light contributions with reduced artificial light contributions. For illustration, a sample occupant is selected and the location of the occupant within the building is shown in Figure 6-9. An example day is selected to show the performance of the system from the perspective of this occupant. The light level experienced by the occupant is the sum of the natural and artificial light curves which approach the preference level set by the simulated occupant. As shown in Figure 6-10, the occupant has a constant preferred light level of 40.2 foot-candles (500 lux). For clarity, the light and preference values are removed for the time slots where the desk is unoccupied. For the sample day illustrated in the figure, the preference of this occupant is tracked well but not perfectly. The error in meeting the preference of this occupant is caused by balancing the demands of this one individual with those of the others in the zone. Because both excess light and insufficient light are considered as contributions to error in this simulation, the zone manager has to balance the preference of each occupant in minimizing total overall error. The shape of the artificial light curve shows the system successfully compensating for the varying sunlight throughout the day to provide a close match to the specified preference, indicating a high degree of satisfaction for this occupant and a reduced dependence on artificial lighting.

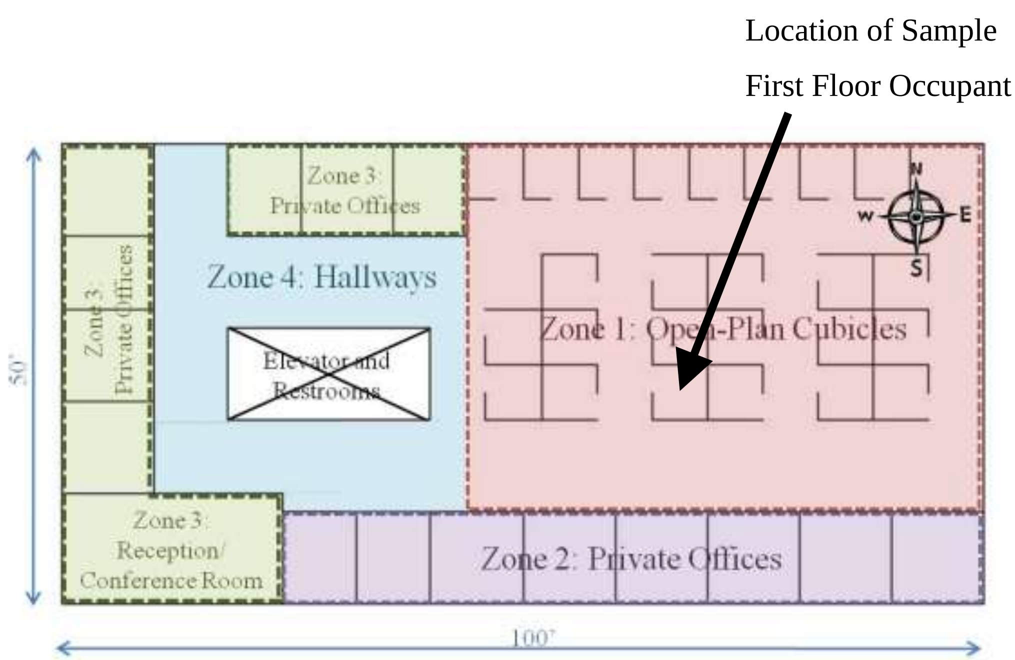

Figure 6-9: Location of Sample Occupant

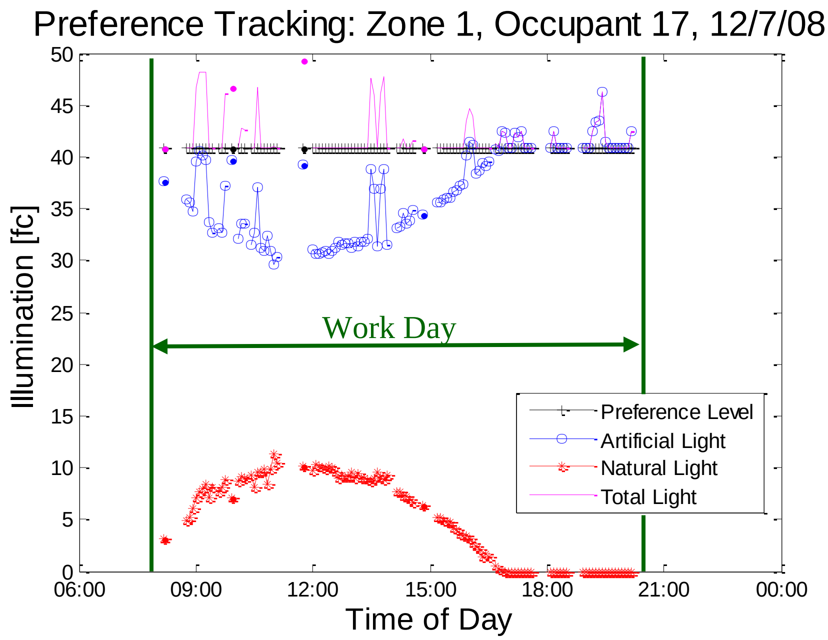

Figure 6-10: Occupant Preference Tracking

As an indication of the general preference tracking performance over the course of the year, the total error in meeting preference demand daily for occupied timeslots divided by the total number of occupied timeslots is plotted below for both the base case and the FICS case. The plotted metric is an average experienced error in meeting preferred light level across all occupants. The artificial lighting results from the base case are used with the occupancy, natural light, and preference levels from the fully implemented case, as these vectors are considered representative of the actual building conditions. Weekend days are excluded from the comparison plot below as there are no preferences or occupancy for those days.

As shown in Figure 6-11, the FICS is a significant improvement over the base case in meeting the preferences of the occupants. Although the base case has a specified target light level to provide to the sensor locations, because it does not have information about the individual preferences of the occupants or the current natural light level, it cannot incorporate this information into its decision-making. The base case system therefore performs very poorly in providing the occupants with their preferred light levels. With this additional preference information, the FICS is able to tailor the light levels throughout the building to better fit the requests.

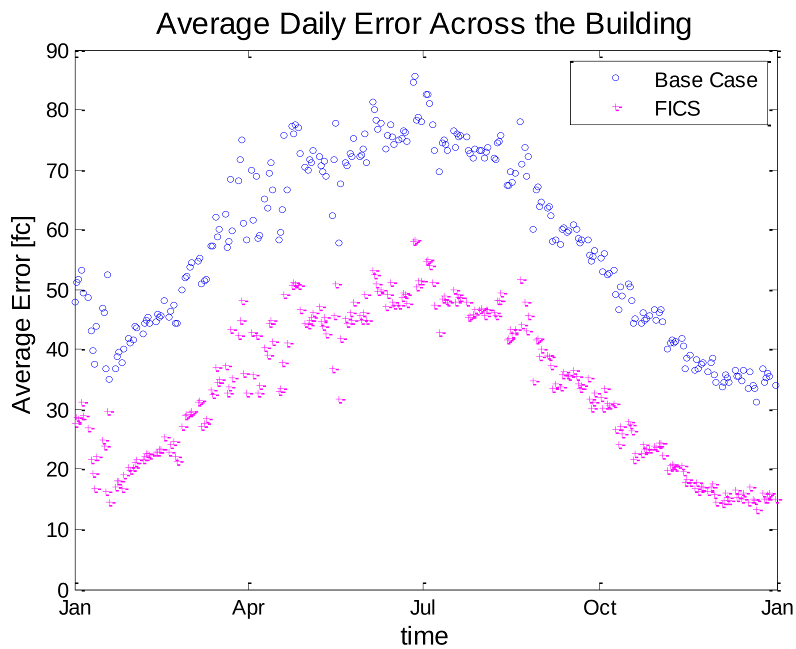

Figure 6-11: Average Daily Preference Error per Occupant Timeslot Comparison

Both the base case and FICS systems perform worse in the summer with regard to preference matching because the occupants closest to the windows receive natural light in excess of their preference level. The use of blinds to avoid excessive direct sunlight is incorporated into the natural lighting module, but simulating additional occupant behavior with respect to limiting incoming natural light is beyond the scope of this simulation. The lighting control system, however, responds by minimizing the additional light for these occupants, both providing them with a better experience and reducing building energy use. The excess sunlight provided to the outer occupants is the main driver of error for both cases shown in the figure. In order to focus on the capabilities and performance of the system, an adjusted measure of error is considered which compensates for timeslots with an overabundance of light. For time steps where more natural light is available than is preferred, the total error is counted as the additional artificial light for that occupant for that time step. For time steps where natural light alone does not meet the needs of the occupants, the error is the difference in supplied light and preferred light for that individual at that time step. The plot of average adjusted daily error is provided in Figure 6-12.

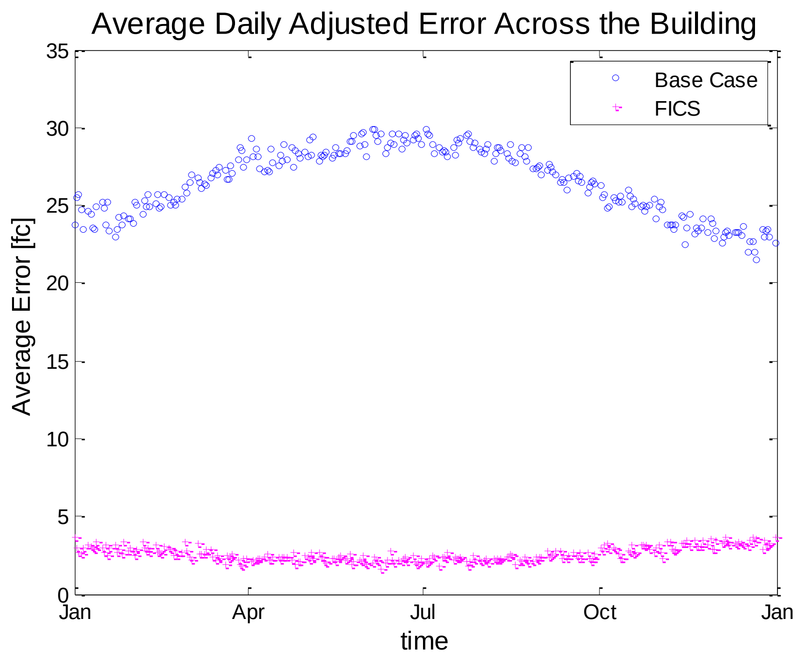

Figure 6-12: Average Adjusted Daily Preference Error per Occupant Timeslot Comparison

The adjusted error shows the error throughout the building that the system is not able to remove and therefore provides a better metric for system performance where shading controls are available to the occupants. As shown in Figure 6-12, the average error throughout the building that is within the control of the system is kept very low throughout the year indicating an improved experience for the occupants.

To assess the relative importance of each component in the system, the summation of adjusted error for all time steps for the year is listed in Table 6-5 for all cases. Table 6-5 indicates preference tracking performance is best achieved for the FICS case. The use of distributed occupancy sensing allows the system to focus on meeting the needs of the present occupants allowing for greater flexibility in the tailored lighting system design. Incorporation of individual preferences targets the system at meeting the true preferences throughout the building to provide overall better performance. Compensating for daylight minimizes the oversupply of light by reducing the artificial lighting contribution when natural light is available. As expected, implementing none of these enhancements leaves the base case with poor comparative performance.

Table 6-5: Comparison of total annual adjusted preference-meeting error

Implementation Case	Total Adj. Error for all Time Steps [fc]	% Increase Over FICS Adjusted Total Error
Fully Implemented Control System (FICS)	1.094×10^{7}	--
Centralized Occupancy Sensing	1.355×10^{7}	23.9%
Standard Light Level Preference	2.334×10^{7}	113.3%
No Daylight Compensation	8.661×10^{7}	691.7%
No Control Implementation (Base Case)	1.118×10^{8}	921.9%

6.3.3 *Demand Response Performance*

A key feature of this system is its ability to respond to a usage curtailment request as part of a demand response system. This capability allows the system to reduce usage at peak cost times and when the grid is at risk of being overloaded. In currently available systems with demand response capabilities, usage curtailment requests are achieved through uniform fractional dimming throughout the building. However, with the tiered system, building-wide energy use reduction can be achieved in a minimally invasive manner. The zones with the smallest loss in utility value for loss of energy resources are successively required to give up units of energy until the building reaches the amount of energy allowable at that time.

In order to demonstrate the comparative advantage of the utility based resource usage prioritization, the FICS case from the previous section is used as the new performance base case. This case represents the energy use of the building assuming no energy usage restriction. The level of preference-meeting error for the FICS case is compared to the error when the lights in the building are uniformly dimmed to 90 percent, 70 percent, and 50 percent as well as to the preference meeting error where the total building-wide energy use is reduced to 90 percent, 70 percent, and 50 percent with the energy allocated according to the utility-based system. In the restricted energy use cases, the light settings, occupancy, and preference settings are updated at 5 minute intervals as with the unrestricted case, but the building-wide energy allocation is performed at 30 minute intervals. The allowable total energy use at these allocation points is the average energy use for five consecutive time steps in the unrestricted case centered on the current time step multiplied by the respective fraction. For consistency, the universal dimming computation is performed at the same 30 minute intervals and is calculated by multiplying the relevant fraction of energy use by the average light settings over the same five step interval. The comparison is conducted from June 1 to August 31 to provide insight into the effects of the usage reduction during the hottest time of the year where total building electricity usage increases dramatically due to the use of cooling systems.

Table 6-6 shows the comparative performance of the two types of energy use restriction, universal dimming and utility-based resource allocation. The adjusted error values presented in the table are calculated similarly to those for the year-long results presented earlier. The tabulated values assume a continuous energy use restriction fraction over the full 3 month time period. The purpose of using a constant level of restriction is to show the typical performance trend for over a wide range of possible scenarios, though any given energy use restriction window would typically be on the order of a few hours. Keeping the reduction fraction constant over an extended period of time also provides information as to the level of degradation in service that could be expected if a building manager chose to implement a long term energy use reduction policy in the building for cost savings purposes. Because energy use is confined to the working day hours for the building simulation, no change in performance is modeled for the nighttime hours which also are not typical load shedding hours. Similarly the weekend days are excluded as weekend building use is not incorporated into the simulation.

Table 6-6: Comparison of energy use curtailment performance

Energy-Use Reduction	Total Adjusted Error for All Time Steps [fc]	
	Universal Dimming	Utility-Based Allocation
0%	2.360×10^6	2.360×10^6
10%	3.796×10^6	3.216×10^6
30%	4.751×10^6	3.912×10^6
50%	5.860×10^6	4.907×10^6

The results show the utility-based allocation is better equipped to meet occupant demand than a universal dimming scenario. Approximately a 20 percent increase in total adjusted error is seen for all dimming fractions for use of universal dimming over utility-based allocation. As expected, a decrease in performance is demonstrated for all increases in dimming fraction. The relatively large decrease in performance between the 0 percent and 10 percent energy use reduction is partially attributable to the switch from a system able to update itself freely every 5 minutes to one with restrictions imposed at 30 minute intervals. This result may suggest that implementation in a real building would benefit from event-driven updating and reallocation instead of a time-based system. For example, the system could be redesigned to update whenever a particular level of demand was unmet and a request is generated by a zone manager. Also, instead of dividing leftover energy units evenly between zones at the end of the building-wide energy allocation procedure, the building server could retain ownership of the leftover units and assign them as they are requested by individual zones.

6.4 Discussion and Evaluation

The results of this simulation study indicate the ability of the new light control and optimization system to both improve occupant performance and save significant amounts of energy in a full building-wide installation. For the simulated building located in Elizabeth City, a 67 percent reduction in energy use and nearly an order of magnitude reduction in adjusted daily average preference-meeting error for the year are indicated as compared to a building with no advanced features. These results indicate the value of installing this type of system with distributed occupancy, light level and light level preference sensing, and tailored control. While the precise performance in any specific building is dependent on the architecture, location, orientation, specific occupants, and use of the building, simulation results indicate that for some buildings this system is highly advantageous.

The analysis described here is subject to several limitations. The simulation study evaluates one type of building. The floor plan of the simulated building is conducive to the incorporation of natural light and therefore the building can capitalize on available daylight to a greater degree than some other building designs could. The narrow and open floor plan allows daylight to penetrate the majority of the office and hallway spaces; wider buildings with more interior walls would reduce the daylight-related energy savings. As new buildings are constructed with advanced daylighting systems and with floor plans built around daylight maximization, the tiered lighting control and optimization system will become more useful for these new buildings and older buildings retrofit with the new technologies. Existing buildings designed without effort made toward daylight incorporation could still benefit from the light control and optimization system, but the energy savings daylight compensation are limited to the daylight the designer allows into the space.

The simulation study is also conducted using data from only two locations. Climate zones with different weather patterns and at different geographical locations have distinct irradiance levels which yield different natural light conditions for a building. As demonstrated by the El Paso case, a building in a typically sunny location has more daylight intensity over more hours than a building located in a typically rainy location. Likewise a building at a latitude with sunrise and sunset times best coinciding with building arrival and departure times year-round will require less artificial lighting energy than one with widely variable day length as the extra hours of sunlight on the long days are essentially wasted as they occur outside the occupied building hours. In addition to changing the geography of the building, changing the orientation of the building on the same site will affect daylight availability and the related potential for energy savings.

The accuracy of the models used to define the lighting and occupancy of the building are limited. The artificial and natural lighting models are basic models used in preliminary building design that provide approximate values for expected light levels throughout the space. The occupancy model is based on a study of the behavior in one specific building. Depending on the particular individuals occupying a building, the occupancy model used here may not be representative. The occupancy model also does not account for late night building use by either desk workers or cleaning staff. Artificial lighting use to facilitate this work during hours with no daylight availability would increase the total building energy use. As the simulation study is conducted to evaluate the realistic performance of a building and not to represent the exact behavior of a specific building, the models are considered sufficiently accurate for this application. However, neither model captures the wide variability throughout the building stock thereby limiting the direct translation of the particular results to individual existing buildings.

With respect to the precise value of natural light compensation, physical measurements or a detailed lighting model created in an advanced lighting simulation program such as RADIANCE© would be necessary in order to evaluate the potential for energy savings potential and occupant preference performance for a particular building. For a more detailed assessment of the composite artificial and natural light performance in the system, building-specific information for forming accurate influence matrices between each artificial light source and each desk surface would be required in addition to a detailed analysis of natural light contribution throughout the year.

While the suitability of the lighting control and optimization system to any specific building is dependent on the resolution of the aforementioned simulation limitations, it is clear from the results of the simulation study that buildings with significant natural light availability and with a diverse population can benefit from the installation of the tiered lighting control and optimization system.

The demand response results further show the suitability of the tiered system to managing the energy distribution throughout a building during energy restriction times. Energy allocation based on utility has the ability to incorporate the needs of the building occupants and assign resources to areas of the building that are most in need. The demand response assessment of the energy allocation system uses the unrestricted energy results as a baseline and bases energy restriction on fractional decreases from this baseline to show the relative degradation in performance for specific degrees of restriction. In a true demand response scenario, the building would need to cut energy use by a specific amount or to below a maximum limit and hold that energy use level for a specified amount of time. The results presented from the simulation study provide an indication of the type of reduction in performance that could be expected when responding to a demand response request, but the simulation scenario and a real demand response request scenario are similar but not identical. However, inherent in the design of the system is the explicit description of the tradeoff between energy use curtailment and building-wide performance which can provide real-time information about the impact of the energy use restriction for a real demand response request.

The overall system performance for both the restricted and unrestricted energy use cases is limited by the systems it is able to control. The optimization algorithm can only provide more light to the sensors, but cannot mitigate when excess natural light is provided. Integration with an automated blinds system would help to alleviate the overabundance of light provided to some occupants. Due to the motion of the sun and the variation in externally reflected light, the incorporation of an automated blind controller to the simulation is not trivial. However, in a real-building installation where blind control is either left to the occupants or is potentially automated, the system itself is designed such that it readily responds to the resultant change in available daylight and updates the lighting scene accordingly. Accounting for daylight distribution pattern modifications due to variable use of blinds is the same as the already incorporated ability to respond to other types of daylighting changes and is therefore a capability already inherent in the system.

6.5 SUMMARY

A simulated 3-story building is designed to evaluate the energy use and occupant performance behavior of the tiered sensing and resource allocation system. A year-long suite of irradiance measurements from Elizabeth City, North Carolina are used to model a year of natural light availability. A stochastic model of building occupancy is incorporated and a distribution of light level preferences is used to select the occupant preferences throughout the building. Substantial energy savings and improved occupant experience result from the implementation of the tiered resource allocation system as compared to a base case scenario with limited sensing and control capabilities. Energy savings result from each of the system capabilities, daylight compensation, occupant preference incorporation, and individual occupancy sensing with daylight compensation providing the most significant savings. The degradation in occupant performance for restricted energy use is compared for a typical uniform dimming response to an energy use curtailment request and for the same energy use fraction reduction spread throughout the system according to the utility-based resource allocation system. The results show improvement in the building-wide occupant experience during an energy restriction period when the energy is allocated based on the zone utility curves. The simulation study demonstrates the advantages of implementing a distributed sensing and actuation system capable of constructing tailored lighting scenes as appropriate and of allocating energy resources based on the ability of the zones to use those resources to improve occupant conditions.

Chapter 7.

SUMMARY, EVALUATION, AND FUTURE EXTENSIONS

This thesis focuses on developing and assessing a new method of optimizing energy use in commercial office spaces. A prototype is developed to demonstrate the implementability of the sensing and control system and a simulation tool is created to evaluate the potential real-world performance. This chapter provides a summary of the project followed by an evaluation of the results from the prototype development and building simulation. The chapter then concludes with a discussion of future developments and directions for this project.

7.1 SUMMARY

A new type of energy resource allocation system for commercial office building lighting is designed and evaluated. The system is hierarchical with distributed sensors and actuators at the bottom level, zone managers at the intermediate level, and a building server at the top level. The bottom level sensors capture an array of spatially distributed information about the building spaces, including light level, light level preference, and occupancy. The actuators individually set the dimming levels in individual lamps to set tailored lighting scenarios. The zone managers are responsible for coordinating groups of sensors and actuators in geographically defined subsections of the building such as a large shared office space or set of smaller offices. The zone managers gather the sensing information from the sensors and use it to determine the optimal light settings for the lamps in the zone and to create a zone utility curve which defines the relationship between energy use and improvement in meeting occupant preferences. The building server level communicates with all zone managers to assign building-wide energy resources based on maximizing building-wide utility using the zone utility curves. The utility curves are determined using mathematical programming techniques and the building-wide allocation of additional energy units is performed using a greedy algorithm and subsequent allocations are conducted by allowing the zones to trade energy units based on associated incremental increases or decreases in utility.

In addition to minimizing the use of energy and maximizing the experience for the building occupants, the system is structured to provide useful information to the building operators. Because the system relies on a characterization of energy use with respect to the improvement in meeting demands and therefore an improvement in the occupant experience, the direct tradeoff between energy use curtailment and performance degradation is provided.

A demonstration of the physical implementability of the system is conducted through the development of a laboratory-scale demonstration. A wireless sensing, actuation, and processing unit is designed to serve the distributed sensing and actuation functions of the system and to fulfill the role of the zone manager. Two sets of three wireless units are used to define two separate zones. These units are each connected to a light sensor and a dimmable fluorescent ballast to complete the sensing and actuation functions. One unit within each zone serves the additional functions of the zone manager. A laptop is connected to a wireless modem for use as the building server. The setup is arranged in a room with exterior windows and is allowed to run over portions of the day that provide large variations in natural light availability. Data gathered from this experiment show the ability of the system to track set preference levels, compensate for natural light availability, and process data in a distributed manner. This demonstration system is designed on a wireless platform to show that a functional system can operate over wireless networks, facilitating the use of this type of system in reconfigurable spaces or for retrofitting applications. While the prototype system specifically demonstrates the viability on this one type of system, the same resource allocation and distributed sensing and actuation system could be implemented on a variety of other types of systems. For example, the development of the Digital Addressable Lighting Interface (DALI) protocol by the International Electrotechnical Commission allows specially designed wired ballasts to be individually addressed and controlled which facilitates the same kind of tailored lighting scenario available with individually addressed wireless ballasts. Sensors linked to desktop computers and user interfaces designed to capture occupant preferences communicating via the internet could also supply the distributed sensing information necessary for this system to function. The resource allocation system is designed to be flexible enough for implementation on a variety of hardware platforms as building systems develop.

In order to assess the performance of the resource allocation system in a full building deployment, a building model is developed. The modeled building is a three story, 15,000 square foot commercial office building with a mix of private offices, conference rooms, and open-plan shared spaces. The simulation incorporates a stochastic occupancy model, randomly selected light preference levels for each occupant, and a year-long suite of varying natural light conditions. Based on these inputs, artificial light settings are calculated at five minute time intervals for the entire year and the performance of the fully implemented system, including occupancy sensing, light level sensing, and occupant preference incorporation, is compared to the performance of a system with each feature removed individually and a system with no advanced features. The results of the simulation indicate the resource allocation system has the potential to both increase occupant performance and significantly decrease energy use.

7.2 System Evaluation

The results of the laboratory prototype development and building simulation indicate that a tiered system allocating energy resources in accordance with the utility curves defined by collective subgroups of building occupants can perform well at both reducing energy and meeting occupant demand. This section focuses first on a discussion of the degree of effectiveness at meeting these goals and subsequently on identifying additional benefits available with this system.

7.2.1 System Effectiveness

The effectiveness of the sensing and control system is evaluated by the degree to which it is able to reduce energy use and match occupant demand. From the simulation results, an energy savings of 67.1 percent is realized with the implementation of the control system over the base case. This represents a significant cost savings to the building operators as well as a reduced energy load for the local utility. The results show the most significant energy use reduction arising from daylight compensation with additional savings from both preference incorporation and occupancy sensing. The occupancy sensing results are representative only of transitioning from a centralized occupancy sensor for all rooms to individual sensors at each desk. Therefore the contribution to total energy savings due to the installation of occupancy sensors for a building with no occupancy sensing would be more significant.

The energy savings demonstrated by the simulation study are specific to the design and location of the building model. The availability of natural light and the distribution of natural light in the building throughout the work day are dependent on the location and orientation of the building. The modeled building is located in Elizabeth City, North Carolina and is designed with a long east-west axis. The site is selected because of its temperate climate and mix of sunny and cloudy days to give an indication of performance in a variable environment. As shown by the comparison to relocating the building in El Paso, Texas, a building located in a sunnier environment experiences greater daylighting savings. By contrast, a building in a predominantly cloudy city such as Seattle, Washington would have less daylight availability and thus would show reduced energy savings. The building orientation is selected in accordance with current general design recommendations for efficient building design to maximize interior daylighting. The façade and interior design of the building further dictate the energy savings potential of the system. Buildings with narrow external offices, small façade window surface area, and high internal walls and partitions restrict daylight penetration within a building and thereby reduce the energy savings potential for daylight compensation. The modeled building has a mix of small offices and large open-plan spaces to provide a balanced result.

In addition to the demonstration of significant energy savings, the simulation shows the ability of the system to match the preferences of the occupants. As compared to the case of no control implementation, the error in matching the desired preferences of the occupants, adjusted for overabundance of natural light, is reduced by an order of magnitude with the use of the fully implemented control system. Allowing the individuals to specify preferences more than halves the error in meeting demand as compared to assigning a standard light level preference. These results indicate that the system performs well at addressing the two potentially competing goals of reducing energy use and providing improved service to the occupants.

The finding of improved performance is extended in the demand response scenarios. By allocating the reduced energy units in accordance with the zone utility curves and the mathematical programming optimization light settings selection, the error in meeting demand is reduced by approximately 20 percent for the dimming fractions evaluated as compared to execution of a universal dimming fraction throughout the building. This result indicates that the utility-based resource allocation system is well equipped to allocate limited resources to where they are best utilized.

7.2.2 System Advantages

In addition to reducing energy use under normal operating conditions, tracking occupant demand, and minimizing the impact of energy use curtailment, the tiered resource allocation system provides several auxiliary benefits.

In addition to enabling the system to better match the actual preferences of the occupants, allowing the occupants to select their preferred light setting gives them a sense of control over their space. Individualizing the control empowers each occupant to participate without the concern for the disruption of others which can be a cause of inaction in shared spaces (Wyon, 2000). A sense of control turns the occupants into active participants in the system and can contribute to avoiding a common issue of new building technology installation where occupant dissatisfaction leads to the subversion of the system or complete removal.

The zone level of the system is designed with a brief automated setup sequence which determines the relationships between light sources and locations of importance. The information recorded forms the basis of all future decision-making. As long as the sensors remain in fixed locations, this empirical information constitutes the entirety of the information required for the system to function appropriately; neither are models of the space nor are estimates of natural light distribution required. The automated sequence can also be rerun at intervals to account for degradation in the physical lighting system over time, following any redesign of the space to accommodate a new layout, or for the addition or subtraction of a sensor or lamp. The acquisition of the information regarding the relationship between light sources and sensors can readily deduce problem areas within a zone. As each sensor independently measures the light received from each lamp, if the sum over all lamps of the sensor measurements does not reach or exceed the minimum standard for that type of space, the design of the room is flawed. The system will still function and provide the best match to the light level specified for that sensor, but may not be able to meet the request and the zone may benefit from a different layout.

Monitoring the zone-by-zone energy usage over time can yield information about the efficiency of the layouts themselves. In a scenario where two zones have similar natural light availability and yet one zone consistently requires more energy to maintain the same level of performance, a building manager viewing the data may be able to recommend changes to the layout or utilization of the spaces to increase the overall energy efficiency. Disparate energy requirements for similar zones may not necessarily point to a problematic location, possibly the individuals in one zone simply prefer higher levels of light or are present more consistently at their desks, but the comparison can indicate places to start looking for improvements.

The provision of explicit tradeoff curves between energy use and performance enables building managers to see the direct performance cost of energy use curtailment and use this information to minimize operating costs when the improvement is minimal. Bounds on performance can be set relative to the current pricing or availability of energy resources provided the utility company is capable of providing this information. This knowledge of tradeoffs assures more than that usage curtailment is conducted in a minimally impactful manner, but also allows curtailment requests to be fulfilled only to the level that a minimal performance standard is maintained. Where other systems may allow a building manager to set energy use restrictions, the system proposed in this thesis enables the assignment of conditions to these settings to include minimum performance standards. The utility-based resource allocation system facilitates a new conceptualization of energy use, not simply as electricity usage as light output, but as electricity usage in terms of a provided service level.

7.3 FUTURE EXTENSIONS

The simulation and laboratory prototype discussed in this thesis are developed to a degree that demonstrates the feasibility of implementing the system and the potential benefits associated with it. In order to realize the full potential of the system and to ready it for real building installation, additional steps are necessary. The availability of the tiered resource allocation system for lighting has implications for new building and retrofit designs and the concept behind the design can be expanded to a wider realm of building systems. This section discusses the necessary developments, implications of installation of this system, and possible extensions of the concepts.

7.3.1 SYSTEM DEVELOPMENT

The laboratory prototype system and the building simulation indicate that the tiered resource allocation system brings with it many advantages. However, for such a system to be deployed in a real building, further development would be necessary.

The necessary developments can be characterized as hardware, software, or system design improvements.

With respect to hardware development, the prototype system is a basic platform suitable to a laboratory scale system. A more robust hardware system, designed to minimize energy use and cost while ensuring reliability would be necessary. Separating the board tasks into individual units specializing in single tasks would help in driving down per-unit cost and facilitate the use of task-specific software. The sensing units require additional sensing capabilities. The laboratory setup does not include occupancy sensors or analog input for occupant preferred dimming level as no physical occupants are located in the space. The prototype units include multiple sensing channels to allow for the expansion to additional sensors, but the interface to those sensors has not been executed. Determining the best positioning of the sensors within an office and which particular sensors work best for different types of office conditions is an issue that needs to be addressed. With the development of occupant interfaces on personal desktop computers for building management purposes, the occupancy and light level preference data could be gleaned in new ways without requiring additional sensor hardware. A light sensor connected to the computer could complete the sensing requirements of the system, thereby eliminating the need for a wireless sensing component to the system where the necessary information could be transmitted via the building computer network. For the sensor network in buildings where this technology is unavailable and for controlling the individual ballasts, wireless radio selection is crucial to the continued performance of the system in a building installation. The radios require appropriate indoor communication range for the site and may require multi-hop or mesh networking abilities to increase range and ensure communication reliability.

On the software side, a user interface for the building manager that works with building management systems needs to be developed. The installation of this new system makes available different types of information to the building manager but this information needs to be communicated in a meaningful way and the building manager needs the ability to define the parameters that determine the performance of the system at all levels. For this system to be used in a building, the lighting control system needs to be able to coexist with the building management systems responsible for managing the other building subsystems. Developing the system according to the ASHRAE BACnet® protocol would allow it to be used with a wide array of existing systems. Alternatively the system could be developed under a proprietary protocol for use in specific high-end building management systems produced by individual companies.

In order to transition this system from a laboratory prototype system to an installed, fully-operational system, incorporation of some additional features would be appropriate. In its current implementation, the simulated system operates on a time-based updating schedule. The light settings update every five minutes and the building allocation updates every thirty minutes. However, a switch to an event-driven updating system with a maximum frequency of changes may provide less distraction and better response for the occupants. For example, the arrival of a previously absent occupant may necessitate a quick lighting scheme update so that the new person is not left in the dark waiting until the next timed update. At the building level, reallocation could be triggered by either a zone request for additional energy resources or from a power company request for electricity usage curtailment. It may also be useful for some buildings to be able to set priorities on meeting demand for certain building segments. At the building-wide level, the zone utility curves can be scaled to give priority to some zones over others, but the user interface to obtain these preferences from the building manager would need to be developed. At the zone level, the mathematical programming formulations that set the optimal lighting scenarios incorporate weighting factors with respect to these priorities, but the ability to select priority sensor locations would require either an additional setting on the sensing units or backend access through the building management system for the building manager to set them directly. Additionally for sensor locations not assigned to specific individuals, such as hallway spaces, target light level settings can be selected using the same type of dimming interface used by occupants at their desks, but it may be preferable to allow internal preset target preferences for these types of units.

7.3.2 Impact on Building Design

The lighting system is designed for use in any type of commercial office building and may also have applicability to other types of commercial buildings and some industrial or manufacturing facilities. The preference and occupancy sensing features of the system contribute to significant occupant performance improvements and some energy savings, but the most significant energy savings are achieved through daylight compensation. The movement in the building design community to incorporate more daylight in new buildings due to effects on mood and occupant health has increased daylight availability in buildings, but only through implementing compensatory dimming are energy savings realized. Researchers are focused on developing new daylighting techniques such as light pipes (Kim & Kim, 2010), solar concentrators with optical fibers (Wang et al., 2010), and electroactive polymers in façade glazing (Krietemeyer et al., 2011) and the availability of an intelligent lighting control system encourages the incorporation of these new technologies in building design by creating operating cost savings that can help offset the initial investments in building design upgrades. Tailored lighting scenarios and incorporation of occupant-defined target preferences ensure that occupants maintain influence on the condition of their workspace and allow for compensation of new natural light distributions throughout the space caused by the new technologies.

Moving toward a control system that allows individual occupants to select their own preferences encourages the use of open-plan office spaces by giving even those in a shared space control over their local environment. Because interior walls and partitions block the penetration of daylight from windows, open plan floors are better equipped to make use of available daylight and reduce energy use. Typical lighting design practice in office spaces is largely concerned with evenness of light level across the space. However, due to the tailored lighting scene capability of the system, an open plan space designed with a high degree of lighting flexibility is advantageous to the system performance. Therefore in designing a room for installation of this system, consideration of the flexibility of the lighting system should be considered in addition to the ability of the installed lighting to provide at least a minimum light level across the space. For example, installing a higher quantity of smaller light fixtures will allow the system to better meet the preference targets of the occupants.

7.3.3 INTEGRATION WITH OTHER SYSTEMS AND BUILDINGS

The ultimate goal of this type of system is the integration of other building systems and potentially the extension to inter-building energy trading. Because the basis of the system is defining the explicit tradeoff between energy use and occupant performance, it is well suited to interaction with other building systems. If similar curves for other building systems requiring electricity resources are developed, in particular cooling and ventilation systems, these systems could trade energy resources according to their respective utility curves. Combining multiple systems into a single resource allocation scheme enables building-wide prioritization of resources and gives the building more latitude in demand response participation.

Given the availability of real-time electricity pricing information, the tiered resource allocation system could also incorporate operating costs into the decision-making process at the top level. For example, when electricity is at a premium, a lower standard of performance may be acceptable than when it is inexpensive. Alternatively a permissible cost versus performance level tradeoff could be defined by the building manager to define when additional costs are warranted. In this way operating cost can be reduced beyond the savings available from standard system performance.

With performance information characterized by energy requirements and monetary value, buildings could potentially enter into cooperative energy reduction agreements where building-wide energy units are traded between buildings based on utility. There are companies which act as intermediaries in selling watts of energy in the form of usage decreases to energy providers by cutting the energy use of their client companies and paying the companies for conforming to the reduction. Inter-building trading would allow building owners to enter into demand response agreements with energy companies to restrict energy use during necessary times while preserving the ability to buy additional resources if necessary from other buildings that have entered into similar agreements but have less need for all of their resources at the specific time. Reciprocally, energy units allowed within the use reduction agreement that are not of large utility value to the building could be sold to other buildings in a manner analogous to a cap-and-trade scenario. Knowing the value of the energy units facilitates informed decision-making with regard to responding to a demand response event with knowledge of the monetary and performance consequences of all courses of action.

REFERENCES

ASCE (2005). *American Society of Civil Engineers: Report card for America's infrastructure.* Annual Report No. 2005. ASCE: Washington, D.C.

Aleklett, K., Höök, M., Jakobsson, K., Lardelli, M., Snowden, S., & Söderbergh, B. (2010). The peak of the oil age-analyzing the world oil production reference scenario in world energy outlook 2008. *Energy Policy, 38*(3), 1398-1414.

Andrews, C., & Krogmann, U. (2009a). Technology diffusion and energy intensity in U.S. commercial buildings. *Energy Policy*, *37*(2), 541-553.

Andrews, C., & Krogmann, U. (2009b). Explaining the adoption of energy-efficient technologies in U.S. commercial buildings. *Energy and Buildings, 41*(3), 287-294.

Atif, M. R., & Galasiu, A. D. (2003). Energy performance of daylight-linked automatic lighting control systems in large atrium spaces: Report on two field-monitored case studies. *Energy and Buildings, 35*(5), 441-461.

Bourgeois, D., Hand, J., Macdonald, I., & Reinhart, C. (2004). Adding sub-hourly occupancy prediction, occupancy-sensing control and manual environmental control to ESP-r. *Proceedings of the ESIM 2004 Conference,* Vancouver, Canada. 1-8.

CIBSE. (1994). *Code for interior lighting* . Chartered Institution of Building Services Engineers: London, U.K.

CIE. (2011). *Selected colorimetric tables.* Commission Internationale de L'Eclairage: Vienna, Austria. Retrieved July 11, 2011, from http://www.cie.co.at/index.php/LEFTMENUE/index.php?i_ca_id=298

Creyts, J., Derkach, A., Nyquist, S., Ostrowski, K., & Stephenson, J. (2007). *Reducing U.S. greenhouse gas emissions: How much at what cost?* Executive Report. McKinsey & Company. Retrieved July 11, 2011, from http://ww1.mckinsey.com/clientservice/ccsi/pdf/US_ghg_final_report.pdf

Dantzig, G. B. (1951). Maximization of a linear function of variables subject to linear inequalities. In T. C. Koopmans (Ed.), *Activity analysis of production and allocation* (pp. 339-359). Wiley: New York, New York.

Doulos, L., Tsangrassoulis, A., & Topalis, F. (2008). The role of spectral response of photosensors in daylight responsive systems. *Energy and Buildings, 40*(4), 588-599.

EIA (2008). *2003 commercial buildings energy consumption survey (CBECS) detailed tables.* Energy Information Administration, United States Department of Energy. Retrieved July 25, 2011, from http://www.eia.doe.gov/emeu/cbecs/cbecs2003/

EIA (2009). *International energy statistics.* Energy Information Administration, United States Department of Energy. Retrieved July 26, 2011, from http://www.eia.gov/cfapps/ipdbproject/

EIA (2010). *Annual energy review 2009: Energy consumption by sector 1949-2009.* Energy Information Administration, United States Department of Energy. Retrieved July 25, 2011, from http://www.eia.gov/totalenergy/data/annual/#consumption

EIA (2011a). *Annual energy outlook 2011: with projections to 2035.* Annual Report No. DOE/EIA-0383(2011). United States Department of Energy, Energy Information Administration: Washington, D.C.

EIA (2011b). *Renewable & alternative fuels.* Energy Information Administration, United States Department of Energy. Retrieved July 26, 2011, from http://www.eia.gov/renewable/annual/preliminary/

Guillemin, A., & Morel, N. (2001). An innovative lighting controller integrated in a self-adaptive building control system. *Energy & Buildings, 33*(5), 477-487.

IESNA. (2000). *The IESNA lighting handbook* (9th Edition ed.). Illuminating Engineering Society of North America: New York, New York.

Jang, W. S., Healy, W. M., & Skibniewski, M. J. (2008). Wireless sensor networks as part of a web-based building environmental monitoring system. *Automation in Construction, 17*, 729-736.

Jennings, J. D., Rubinstein, F. M., DiBartolomeo, D., & Blanc, S. L. (2000). Comparison of control options in private offices in an advanced lighting controls testbed. *Journal of the Illuminating Engineering Society, 29*(2), 39-60.

Kastner, W., Neugschwandtner, G., Soucek, S., & Newman, H. M. (2005). Communication systems for building automation and control. *Proceedings of the IEEE, 93*(6), 1178-1203.

Kavalec, C., & Gorin, T. (2009). *California energy demand 2010-2020 adopoted forecast.* Technical Report No. CEC-200-2009-012-CMF. California Energy Commission: Sacramento, California.

Kensek, K., Noble, D., Schiler, M., & Narang, V. (2000). Controlling and monitoring of building systems. *Proceedings of the Eighth International Conference on Computing in Civil and Building Engineering,* Stanford, California. 596-603.

Kim, J. T., & Kim, G. (2010). Overview and new developments in optical daylighting systems for building a healthy indoor environment. *Building and Environment, 45*(2), 256-269.

Kintner-Meyer, M., & Conant, R. (2005). Opportunities of wireless sensors and controls for building operation. *Energy Engineering, 102*(5), 27-48.

Krietemeyer, E. A., Smith, S. I., & Dyson, A. H. (2011). Dynamic window daylighting systems: Electropolymeric technology for solar responsive building envelopes. *Proceedings of SPIE Electroactive Polymer Actuators and Device,* San Diego, California. *7976*(79763) A1-A15.

Kurian, C., Aithal, R., Bhat, J., & George, V. (2008). Robust control and optimisation of energy consumption in daylight—artificial light integrated schemes. *Lighting Research and Technology, 40*(1), 7-24.

Li, D. H. W., & Lam, J. C. (2001). Evaluation of lighting performance in office buildings with daylighting controls. *Energy and Buildings, 33*(8), 793-803.

Li, D. H. W., Lam, T. N. T., & Wong, S. L. (2006). Lighting and energy performance for an office using high frequency dimming controls. *Energy Conversion and Management, 47*(9-10), 1133-1145.

Littlefair, P. J., & Motin, A. (2001). Lighting controls in areas with innovative daylighting systems: A study of sensor type. *Lighting Research and Technology, 33*(1), 59-73.

Love, J. (1995). Field performance of daylighting systems with photoelectric controls. *3rd European Conference on Energy-Efficient Lighting*, Newcastle-upon-Tyne, United Kingdom. 75-82.

Lynch, J. P. (2002). *Decentralization of wireless monitoring and control technologies for smart civil structures.* Ph.D. Thesis, Department of Civil and Environmental Engineering, Stanford University. Stanford, California, USA.

Mahdavi, A., & Pröglhöf, C. (2008). User-system interaction models in the context of building automation. In A. Zarli, & R. Scherer (Eds.), *eWork and eBusiness in architecture, engineering, and construction: ECPPM 2008*. Taylor & Francis: London. 389-396.

Maniccia, D., Rutledge, B., Rea, M. S., & Morrow, W. (1999). Occupant use of manual lighting controls in private offices. *Journal of the Illuminating Engineering Society, 28*(2), 42-56.

Menzel, K., Pesch, D., O'Flynn, B., Keane, M., & O'Mathuna, C. (2008). Towards a wireless sensor platform for energy efficient building operation. *Tsinghua Science & Technology, 13*, 381-386.

Ne'eman, E., Sweitzer, G., & Vine, E. (1984). Office worker response to lighting and daylighting issues in workspace environments: A pilot survey. *Energy and Buildings*, *6*(2), 159-171.

Newsham, G. R. (1994). Manual control of window blinds and electric lighting: Implications for comfort and energy consumption. *Indoor and Built Environment, 3*(3), 135-144.

Newsham, G. R., Mahdavi, A., & Beausoleil-Morrison, I. (1995). LIGHTSWITCH: A stochastic model for predicting office lighting energy consumption. *Proceedings of the 3rd European Conference on Energy-Efficient Lighting: Right Light Three,* 59-66.

Newsham, G. R., Veitch, J. A., Arsenault, C., & Duval, C. (2004). Effect of dimming control on office worker satisfaction and performance. *IESNA Annual Conference Proceedings*, Tampa, Florida. 19-41.

Newsham, G. R., Aries, M., & Mancini, S. F.,G. (2008). Individual control of electric lighting in a daylit space. *Lighting Research & Technology, 40*(1), 25-41.

Newsham, G. R., Brand, J., Donnelly, C. L., Veitch, J. A., Aries, M., & Charles, K. E. (2009). *Linking indoor environment conditions to organizational productivity: A field study.* Technical Report No. NRCC-49714. National Research Council Canada Institute for Research in Construction: Ottawa, Canada.

Newsham, G. R., Mancini, S., Veitch, J., Marchand, R., Lei, W., Charles, K., & Arsenault, C. (2009). Control strategies for lighting and ventilation in offices:

Effects on energy and occupants. *Intelligent Buildings International, 1*(2), 101-121.

Newsham, G. R., & Birt, B. (2010). *Demand-responsive lighting: A field study.* Technical Report No. NRCC-53233. National Research Council Canada, Institute for Research in Construction: Ottawa, Canada.

NREL. (2004). *Reference solar spectral irradiance: Air mass 1.5.* ASTM G173-03. National Renewable Energy Laboratory. Retrieved July 11, 2011, from http://rredc.nrel.gov/solar/spectra/am1.5/

NREL. (2010). *Cooperative networks for renewable resource measurements (CONFRRM).* National Renewable Energy Laboratory. Retrieved 07/11, 2011, from http://rredc.nrel.gov/solar/new_data/confrrm/

Perez, R., Ineichen, P., Seals, R., Michalsky, J., & Stewart, R. (1990). Modeling daylight availability and irradiance components from direct and global irradiance. *Solar Energy, 44*(5), 271-289.

Perez-Lombard, L., Ortiz, J., & Pout, C. (2008). A review on buildings energy consumption information. *Energy and Buildings, 40*(3), 394-398.

PIER. (2007). *Demand response as a system reliability resource--1st and 2nd year test report.* Interim Report No. CEC-500-2007-092. Public Interest Energy Research Program, California Energy Commission: Sacramento, California.

Reinhart, C. F. (2004). Lightswitch-2002: A model for manual and automated control of electric lighting and blinds. *Solar Energy, 77*(1), 15-28.

Roche, L. (2002). Summertime performance of an automated lighting and blinds control system. *Lighting Research & Technology, 34*(1), 11-25.

Roulet, C. (2006). Perceived health and comfort in relation to energy use and building characteristics. *Building Research and Information, 34*(5), 467-474.

Rubinstein, F. M., Jennings, J., Avery, D., & Blanc, S. (1998). *Preliminary results from an advanced lighting controls testbed.* Technical Report No. LBNL-41633. Lawrence Berkeley National Laboratory: Berkeley, California.

Rubinstein, F. M., & Kiliccote, S. (2007). *Demand responsive lighting: A scoping study.* Technical Report No. LBNL-62226. Lawrence Berkeley National Laboratory: Berkeley, California.

Sezgen, O. (2000). Interactions between lighting and space conditioning energy use in U.S. commercial buildings. *Energy, 25*(8), 793-805.

Sierra, E., García-Martínez, R., Hossian, A., Britos, P., & Balbuena, E. (2006). Providing intelligent user-adapted control strategies in building environments. *Research in Computing Science Journal, 19*, 235-241.

Sierra, E., Hossian, A., Britos, P., Rodriguez, D., & Garcia Martinez, R. (2007). Fuzzy control for improving energy management within indoor building environments. *Proceedings of the Electronics, Robotics, and Automotive Mechanics Conference (CERMA)*, Cuernavaca, Mexico. 412-416.

Singhvi, V., Krause, A., Guestrin, C., Garrett Jr, J. H., & Matthews, H. S. (2005). Intelligent light control using sensor networks. *Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems,* San Diego, California, 218-229.

Stein, B., Reynolds, J. S., Grondzik, W. T., & Kwok, A. K. (2006). Illumination. In *Mechanical and electrical equipment for buildings 10th ed.* Wiley: Hoboken, New Jersey. 517-662.

Sueyoshi, T., & Gopalakrishna, R. T. (2008). Why did the California electricity crisis occur?: A numerical analysis using a multiagent intelligent simulator. *IEEE Transactions on Systems, Man and Cybernetics. 38*(6), 779-790.

Sylvania. (2000). *Spectral power distributions of SYLVANIA fluorescent lamps.* Technical Information Bulletin No. FAQ0041-0800. Osram Sylvania: Danvers, M.A.

TAOS. (2007). *Light-to-voltage converters: TSL12s, TSL13s, TSL14s.* Datasheet No. TAOS051E. Texas Advanced Optoelectronic Solutions Inc: Plano, Texas.

Tregenza, P., Romaya, S., Dawe, S., Heap, L., & Tuck, B. (1974). Consistency and variation in preferences for office lighting. *Lighting Research and Technology, 6*(4), 205.

Veitch, J. A., & Newsham, G. R. (2000). Preferred luminous conditions in open-plan offices: Research and practice recommendations. *Lighting Research and Technology, 32*(4), 199-212.

Wang, C., Abdul-Rahman, H., & Rao, S. (2010). Daylighting can be fluorescent: Development of a fiber solar concentrator and test for its indoor illumination. *Energy and Buildings, 42*(5), 717-727.

Wang, Y. (2007). *Wireless sensing and decentralized control for civil structures: Theory and implementation.* Ph.D. Thesis, Department of Civil and Environmental Engineering, Stanford University. Stanford, California, USA.

Wolfe, P. (1959). The Simplex method for quadratic programming. *Econometrica: Journal of the Econometric Society, 27*(3), 382-398.

Wyon, D. P. (2000). Individual control at each workplace: The means and the potential benefits. In *Creating the Productive Workplace.* Taylor & Francis: New York, New York. 192-206.