

Energy-Efficient and Comfortable Buildings through Multivariate Integrated Control (ECoMIC)

Award number: **DE-EE0003978**

Recipient: **Philips Electronics North America Corporation**

PI: **Dagnachew Birru**

Dagnachew.Birru@philips.com

Reporting type: **Final Report**

Reporting Date: **October 28, 2013**

Acknowledgment: "This material is based upon work supported by the Department of Energy under Award Number DE-EE0003798."

Disclaimer: "This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof."

Executive Summary

This project aims to develop an integrated control solution for enhanced energy efficiency and user comfort in commercial buildings. The developed technology is a zone-based control framework that minimizes energy usage while maintaining occupants' visual and thermal comfort through control of electric lights, motorized venetian blinds and thermostats. The control framework is designed following a modular, scalable and flexible architecture to facilitate easy integration with exiting building management systems.

The control framework contains two key algorithms: 1) the *lighting load balancing algorithm* and 2) the *thermostat control algorithm*. The *lighting load balancing algorithm* adopts a model-based closed-loop control approach to determine the optimal electric light and venetian blind settings. It is formulated into an optimization problem with minimizing lighting-related energy consumptions as the objective and delivering adequate task light and preventing daylight glare as the constraints. The *thermostat control algorithm* is based on a well-established thermal comfort model and formulated as a root-finding problem to dynamically determine the optimal thermostat setpoint for both energy savings and improved thermal comfort.

To address building-wide scalability, a system architecture was developed for the zone-based control technology. Three levels of services are defined in the architecture: external services, facility level services and zone level services. The zone-level service includes the control algorithms described above as well as the corresponding interfaces, profiles, sensors and actuators to realize the zone controller. The facility level services connect to the zones through a backbone network, handle supervisory level information and controls, and thus facilitate building-wide scalability. The external services provide communication capability to entities outside of the building for grid interaction and remote access.

Various aspects of the developed control technology were evaluated and verified through both simulations and testbed implementations. Simulations coupling a DOE medium office reference building in EnergyPlus building simulation software and a prototype controller in Matlab were performed. During summer time in a mixed-humid climate zone, the simulations revealed reductions of 27% and 42% in electric lighting load and cooling load, respectively, when compared to an advanced base case with daylight dimming and blinds automatically tilted to block direct sun. Two single-room testbeds were established. The testbed at Philips Lighting business building (Rosemont, IL) was designed for quantifying energy performance of integrated controls. This particular implementation achieved 40% and 79% savings on lighting and HVAC energy, respectively, compared to a relatively simple base case operated on predefined schedules. While the resulting energy savings was very encouraging, it should be noted that there may be several caveats associated with it. 1) The test was run during late spring and early summer, and the savings numbers might not be directly used to extrapolate the annual energy savings. 2) Due to the needs for separate control and metering of the small-scale demonstrator within a large building, the HVAC system, hence the corresponding savings, did not represent a typical energy code-compliant design. 3) The light level in the control case was regulated at a particular setpoint, which was lower than then the full-on light level in the base case, and the savings resulted from tuning down the light level to the setpoint was not attributable to the contribution of the developed technology. The

testbed at the Lawrence Berkeley National Laboratory (Berkeley, CA) specifically focused on glare control integration, and has demonstrated the feasibility and capability of the glare detection and prevention technique. While the short one-month test in this testbed provided a functional indication of the developed technology, and it would require at least a full solstice-to-solstice cycle to ruinously quantify the performance, which was not possible within the project timeframe.

There are certain limitations inherited from the operational assumptions, which could potentially affect the effectiveness and applicability of the developed control technologies. The system takes a typical ceiling-mounting approach for the photosensor locations, and therefore, the control performance relies on proper commissioning or the built-in intelligence of the photosensor for pertinent task light level estimations. For spaces where daylight penetration diminishes significantly deeper into the zone, certain modification to the control algorithms is required to accommodate multiple lighting control subzones and the corresponding sensors for providing a more uniform light level across the entire zone. Integrated control of visual and thermal comfort requires the lighting control zone and thermal control zone to coincide with each other. In other words, the area illuminated by a lighting circuit needs to be the same area served by the thermostat. Thus, the original zoning will potentially constrain the applicability of this technology in retrofitting projects.

This project demonstrated the technical feasibility of a zone-based integrated control technology. From the simulation results and testbed implementations, up to 60% lighting energy savings in daylit areas relative to a “no-controls” case can easily be achieved. A 20% reduction of whole building energy consumption is also attainable. In the aspect of occupant comfort, the testbed demonstrated the ability to maintain specified light level on the workplane while promptly mitigate daylight glare 90% of the time. The control system also managed to maintain the thermal environment at a comfortable level 90% of the time. The aspect of system scalability was guaranteed by the system architecture design, based on which the testbeds were instantiated. Analysis on the aspect of economic benefit has yielded an about 6-year payback time for a medium-sized building, including the installation of all hardware and software, such as motorized blinds and LED luminaires. The payback time can be significantly reduced if part of the hardware is already in place for retrofitting projects. It needs to be noted that since the payback analysis was partly based on the testbed performance results, it is constrained by the caveats associated with the testbed implementations. The main uncertainty lies in the contribution from the space conditioning energy savings as it was non-trivial to realistically configure a room-size HVAC system for directly extrapolating whole-building HVAC energy savings. It is recommended to further evaluate the developed technology at a larger scale, where the lighting and HVAC energy consumption can be realistically measured at the building level, to more rigorously quantify the performance potentials.

Table of Contents

1	Introduction	1
1.1	Project objectives.....	1
1.2	Scope of work and approach	1
1.3	Scope of the report.....	2
2	Control objectives	3
2.1	Energy efficiency	3
2.1.1	Electric load.....	3
2.1.2	Heat gains and cooling/heating loads.....	3
2.2	Visual performance and comfort	3
2.2.1	Task illumination	3
2.2.2	Daylight glare	4
2.3	Thermal comfort	4
3	System architecture	6
3.1	System level overview.....	6
3.2	External services specification	7
3.3	Facility services specification	7
3.4	Backbone network specification.....	9
3.5	Zone services specification	10
3.5.1	Backbone network interface.....	10
3.5.2	Inputs	11
3.5.3	ECoMIC algorithm	12
3.5.4	Profiles	13
3.5.5	Schedules	13
3.5.6	Actuators.....	14
4	Multivariate integrated controls – ECoMIC algorithm.....	15
4.1	Framework	15
4.2	Control systems and variables	15
4.3	Modeling (w.r.t. control variables)	16
4.3.1	Task illumination	17
4.3.2	Electric lighting load.....	19
4.3.3	Heat gains and cooling/heating load	20

4.3.4	Glare.....	23
4.3.5	Thermal comfort	26
4.4	Lighting load balancing algorithm.....	27
4.5	Thermostat control algorithm.....	28
4.6	User preference learning	29
4.6.1	User preference learning strategy	29
4.6.2	Preference learning algorithm	29
5	Simulation-based evaluation	32
5.1	System behavior evaluation using step response simulation.....	32
5.2	System behavior evaluation using simulations based on real measurement	33
5.3	Energy performance simulation using building energy modeling tools	36
5.4	Benefit of multivariate integrated control over the most advanced lighting control technologies.....	37
6	Testbed implementation and evaluation.....	41
6.1	Test sites and testing aspects	41
6.1.1	Rosemont	41
6.1.2	LBNL FLEXLAB.....	41
6.2	Hardware	42
6.2.1	Sensors	42
6.2.2	Control hardware	45
6.3	Software	47
6.3.1	Rosemont site	47
6.3.2	LBNL FLEXLAB site	49
6.4	Configuration	50
6.4.1	Rosemont	50
6.4.2	LBNL FLEXLAB	51
6.4.3	Commissioning.....	52
6.5	Performance	53
6.5.1	Energy performance	54
6.5.2	Visual performance and comfort	58
6.5.3	Thermal comfort performance	61
6.6	Generalizability and limitations	63

7	Success assessment	67
7.1	Energy savings.....	67
7.2	Occupant comfort.....	68
7.3	System scalability.....	68
7.4	Economic benefits.....	69
8	Conclusion.....	71
	Reference.....	73
Appendix A	Glare characterization, detection and estimation	A-1
Appendix B	Blind deployment and blind tilt angle for glare controls	B-1
Appendix C	General commissioning procedure for glare control.....	C-1
Appendix D	Return on investment calculations	D-1

1 Introduction

1.1 Project objectives

The principal objective of the Energy-Efficient and Comfortable Buildings through Multivariate Integrated Control (ECoMIC) project is to develop technologies that reduce energy consumption in new and existing commercial buildings while improving comfort. Specifically, the project aims at developing integrated energy optimization control solutions for electric light, daylight and local space conditioning (HVAC) by exploiting the interdependencies among the multiple variables involved. Occupant satisfaction considerations are also incorporated into the comprehensive solution in key aspects of visual performance, visual comfort and thermal comfort. In addition, a reconfigurable architecture for improved communication efficiency is developed as a carrier of the zone-based integrated control technologies to address reliability and scalability challenges in building-wide connectivity. The primary application space of the control solutions is office buildings; however, the same technologies may be applied to other types of commercial buildings where energy and occupant comfort are of concerns.

1.2 Scope of work and approach

This project was executed in three research stages, which were in line with the three administrative phases.

Phase 1 was the technology development stage. We formulated the integrated control problem and mathematically modeled factors that affect energy usage, visual comfort and thermal comfort under the control framework. Part of the efforts was also devoted to develop techniques for glare detection and prevention as well as user preference learning. The related works are summarized in Section 0. In the meantime, a system architecture addressing interoperability and information flow was developed for the zone-based sensing, actuating and control framework to easily scale up to a building-wide service. The system architecture is described in Section 3.

Phase 2 was the technology verification stage. The concepts developed in *Phase 1* were simulated, evaluated, improved and refined. We utilized a variety of modeling and simulation tools, including EnergyPlus, Radiance, Matlab, Building Controls Virtual Test Bed, etc., to assess the behavior and performance of various aspects of the developed technologies. The simulation results served as the basis of several improvements and adjustments for further enhancing the developed technologies. A description of the various simulation instances and the corresponding results is summarized in Section 5.

Phase 3 was the technology demonstration stage. Concept demonstrators were established at two test sites to physically evaluate the performance of the control technologies. We used as much commercial off-the-shelf components as possible while customizing and prototyping components as needed. The testbeds are instantiated following the developed system architecture, and the controllers were realized both as an embedded system and a software controller. The resulting energy savings, visual performance, visual comfort and thermal comfort were quantified, and further enhancements were made during the process. The activities on the testbed implementations and the corresponding analyses are summarized in Section 6.

1.3 Scope of the report

Various technical reports have been generated as deliverables along the course of this project, and most of the technology development was also narrated in the quarterly research project progress reports. In an effort to avoid an overwhelmingly lengthy final report, this document aims to be very concise while containing enough technical details to stand alone on its own. Documents describing a particular technical aspect of this project in full length can be found in previous deliverables and reports or in the appendices if not previously reported. In the meantime, materials from the later part of Phase 3 that have not been covered in previous documents are described in detail herein.

The control objectives and the corresponding considerations are first introduced in Section 2. The development of the building-wide system architecture for a scalable zone-based control technology is described in Section 3. Section 4 zooms in on the technical details of the multivariate integrated control algorithms residing in the zone service of the system architecture. Verifications and performance evaluations of the control algorithms through simulation studies are summarized in Section 5 followed by section 6 documenting the system performance on testbed implementations. Section 7 discusses the successfullness of this project, and Section 8 sums up the report with a conclusion.

2 Control objectives

2.1 Energy efficiency

In general, energy efficiency is the goal of efforts to reduce the amount of energy required to provide products and services. The objective of energy efficiency is to deliver desired comfort levels with the least possible amount of energy. Cutting electricity consumption of the lighting systems and reducing space cooling/heating load of the HVAC systems are means of minimizing energy usage under the scope of this particular project.

2.1.1 Electric load

Electric load in this project mainly refers to the electricity consumed by the lighting system. Various control strategies have been developed for reducing electric lighting load by turning off lights when not needed (occupancy sensing), reducing the light level to leverage illumination from natural light (daylight harvesting), tuning down the overall lighting to the recommended level (institutional light level tuning), etc. Built on the existing strategies, this project intends to incorporate this particular consideration with other energy and comfort aspects in the scope of lighting and daylighting management.

2.1.2 Heat gains and cooling/heating loads

The heat gains and the corresponding cooling/heating loads considered in this project are those directly related to lighting and daylighting controls. They include heat gain generated from the electric lighting systems and solar heat gain admitted into the controlled space through window and shading systems. The heat gains will eventually become cooling loads that need to be removed by the HVAC system in order to maintain a comfortable thermal environment during cooling seasons. On the other hand, the heat gains may play a role in relieving the HVAC heating load during heating seasons. Therefore, controlling the amount of heat gain will have a direct impact on the HVAC energy consumption.

Maintaining a comfortable thermal environment through temperature control is another interest of the project. The temperature in a space directly affects the HVAC energy demand. An unnecessarily low or high temperature setting will result in excessive cooling or heating energy consumption.

2.2 Visual performance and comfort

The two vision-related factors considered in this project are visual performance and visual comfort. Adequate task illuminance prevents eye strain and discomfort to ensure visual performance over the duration of the entire work hour. A key to maintaining visual comfort is to minimize the risk of disturbing visual sensation due to improper contrast introduced by excessive amount of daylight.

2.2.1 Task illumination

Task illumination in offices is often provided by a combined contribution of electric light and daylight, and the control criterion is to provide adequate and reasonably uniform illumination. The electric lighting system is designed to perform this function during night-time hours or in zones without daylight. Daylight is non-uniform. During daylit hours the control system dims the electric lights so that the least daylit area is still lit to the target level.

In the scope of this project, task illumination is measured using ceiling-mounted photosensors. It is difficult to accurately measure the daylight contribution. When the electric light is dimmed by the estimated daylight level at the target location, inadequate overall light could be an issue due to overestimation of the daylight component. An alternative procedure is to dim the electric light by a fraction of the estimated daylight level so as to avoid under lighting. Most people prefer more light, so this alternate procedure will usually produce increased user satisfaction, at a cost of slightly less savings than is possible with full dimming.

2.2.2 Daylight glare

Daylight glare is a major architectural concern. Vertical illuminance at the eye has been identified to have a reasonable correlation with comfort for large area glare sources when high intensity small area sources are excluded. Venetian blind attenuates incoming daylight and, hence, can be used to avoid glare. The first glare control objective is to control blind deployment and tilt to exclude direct sun from the work environment. This constraint may be relaxed if the sun is sufficiently dim (clouds, other obstructions, or very low solar altitude) and will not constitute a serious glare source. The second glare control objective is to further control blind deployment and tilt to limit the vertical illuminance from the window (sky and reflected sun light) to below a user determined target level. There could be a default value for this target level based on certain empirical model; however, the sensitivity to glare varies widely from person to person, so the ability of the user to modify the target level is an important design feature.

2.3 Thermal comfort

Thermal sensation is a subjective feeling perceived by an individual in response to the environment to which one is exposed. Thermal comfort is the condition of mind that expresses satisfaction with the thermal environment [1]. Therefore, both thermal sensation and comfort varies from person to person both physiologically and psychologically. One of the basic ideas of heating, ventilating and air conditioning (HVAC) engineering is to deliver an environment that is considered thermally comfortable by the occupants. Many chamber tests and field studies have been devoted to identifying thermal comfort conditions, and standards, such as ASHRAE Standard 55P and ISO 7730, have been established for HVAC design practice and references.

Table 1 PMV scale.

PMV	Perceived sensation
-3	Cold
-2	Cool
-1	Slightly cool
0	Neutral
+1	Slightly warm
+2	Warm
+3	Hot

The most recognized thermal comfort metric is Fanger's predict mean vote (PMV) and the associated predicted percent dissatisfied (PPD) value, which has been adopted in all major standards. PMV is a 7-

point index scaled to fit the votes of a large number of human subjects exposed to different thermal conditions in a controlled climate chamber. The PMV values and sensations they represent are shown in Table 1. The PPD value complements PMV and indicates the percentage of people that will vote the thermal condition unsatisfactory given a PMV number. The relationship between PMV and PPD is shown in Figure 1. This project adopts PMV when defining and evaluating thermal comfort.




Figure 1 Relationship between PMV and PPD.

3 System architecture

Having identified the scope of this control framework and objectives, one key aspect of this project was to develop a scalable and reconfigurable system architecture that facilitate implementation of the control algorithms, provide communication reliability and enable building-wide scalability.

This section summarizes the specification of the ECoMIC system architecture. The components and device interfaces that perform core functions in the ECoMIC architecture are described in the subsections below. The core requirements are abstracted to allow flexibility in implementation while formalizing the information exchange between components in a way that fosters interoperability. The architecture leverages the existing protocols and technologies that are suitable for ECoMIC implementation and well-received by building controls community; thus, making it suitable for implementation over legacy building control platforms. The architecture specification focuses on enabling easy implementation of novel concepts specific to multivariate integrated control.

3.1 System level overview

The system can be decomposed into three layers: external services, facility level services and zones as shown in Figure 2. External services are those that reside remote from the facility and provide useful information to a subscriber facility typically over the Internet (e.g. a weather service). Facility level services are services that typically serve the entire area supported by the architecture (e.g. scheduling service). A zone is defined as a space serviced by a single instantiation of the ECoMIC algorithm. Typically the facility will be partitioned into multiple (preferably non-overlapping) zones, each controlled by a dedicated ECoMIC zone controller. Zones are the most granular level of control in the architecture, aggregating sensor inputs, user preferences, facility manager overrides and schedules to drive actuator output (e.g. adapt lights, blinds, shades and thermostat for user comfort and energy savings). External services typically provide useful information to facility level services which would pass the relevant information to zones. External services rarely provide information further down the hierarchy to the zones directly without an intermediary service. The information flows through a variety of networks which are all implementation dependent and do not necessarily share a conceptual model. The facility level services connect to the individual zones via a backbone network that runs through the facility.




Figure 2 System level functional view.

3.2 External services specification

Several external services that may be of specific interest to the ECoMIC ecosystem are described herein. There may be additional external services connected to the system which are handled through mechanisms provided by existing facility level services or the backbone network.

Demand Response Service

Demand response (DR) refers to events initiated to reduce the peak energy demand in response to external message from a power utility intended to relieve the stress on a electric grid. Efforts to standardize these services are in progress which will enable ECoMIC systems to exploit this service to optimize energy utilization. The execution of DR should be automated by subscribing to a standardized automatic DR messaging service over Internet, such as OpenADR [2].

Remote Monitoring Service

Graphical visualization of system state is very useful for assessing the health of the system and diagnosing system issues. World-wide web-based solutions are an obvious technology direction for universal access and monitoring. A typical implementation of web-based access would employ a web-server running on an application server. The web-server would query the database to fetch the system state information which would be displayed to end-user on a web-browser. Furthermore, this service will also allow queries from Energy Information Systems (EIS) for historical time series data as well as from Energy Dashboards for displaying operation and energy usage conditions with real-time data.

3.3 Facility services specification

A facility level service is defined as a service that resides inside the facility and provides information to components and devices within the same facility. Typically a facility level service component exists in only one physical component and is accessible via the backbone network by the individual zones within the facility. Figure 3 shows the facility level functional components and their connection with other services within the overall architecture. Each component is described below.




Figure 3 Facility level functional components.

Building management system

A building automation system (BAS) is a facility level service that controls, monitors and optimizes a wide range of facility components. These systems provide functions that overlap with the ECoMIC system framework, but typically include a much larger range of systems including HVAC, security, access control, safety, fire alarming and electrical power control. The ECoMIC architecture aims to leverage existing systems in order to facilitate wider adoption, and therefore, existing BAS will be the primary driver for a system implementer to decide on which network protocol to bind the concepts in this architecture against.

Facility management

The facility management functional block provides a user interface for facility managers to interact with the BAS. This component provides monitoring and control of all system components in the building. It enables modifications to valid parameter ranges associated with a given profile in the ECoMIC algorithm. In the ECoMIC architecture the facility management system connects directly to the BAS and does not use the backbone network to communicate directly with the individual zones. In this regard, the BAS acts as a gateway to the backbone network.

Demand response interface

The ECoMIC system should be able to participate in DR events to temporarily reduce energy demand in the facility. The DR strategies employed are beyond the scope of the architecture, but the architecture needs to facilitate implementation of such strategies. A sample implementation based on OpenADR is present in Figure 4. A client interface connects via an protocol (e.g. HTTP or XMPP) to the external demand response service. Messages received from the client interface are acted upon by the demand response strategies. These strategies are configured via a demand response commissioning component, which is implementation dependent. The result of the strategy is converted to setpoints and transmitted over the backbone network to the individual zones.




Figure 4 Demand response interface utilizing OpenADR.

Data storage

System performance and operational data can be exploited to optimize and fine tune the system variables and algorithms. Systems logs are also useful for system state monitoring, diagnosis, fault detection, usage pattern, energy savings, trending, benchmarking and maintenance. To address this requirement, the ECoMIC architecture emphasizes easy access to information. Many existing protocols

support event notifications (e.g. COV in BACnet) and data aggregations should be utilized when feasible. In addition, zone controllers should be configured to provide an appropriate level of log information during the commissioning phase.

Remote access

The Remote Access component coordinates data transfer between the local data storage component and the application server, EIS, or Energy Dashboard. This component utilizes an implementation-specific protocol to transfer and filter information about the system. It may communicate over an IP network to the data storage component.

User controls

Personalized, easy to access controls can enhance user comfort, convenience and satisfaction, while improving system performance. Personalized control interfaces can be provided as a desktop GUI program or web browser based GUI. Users are associated with devices which control their environment during the commissioning phase. The User Controls component routes inputs from the user to the appropriate zone controller, typically over the backbone network in an implementation dependent manner.

Commissioning

Commissioning includes partitioning spaces into zones, configuring zone controllers, linking sensor and actuators to zone controllers and users, calibrating sensors and configuring timers, setpoints and other operational parameters. Overlaying logical control architecture over physical network architecture is performed during the commissioning phase which includes mapping the physical address of a device to a logical address and, in many cases, to a location for a quick identification by the maintenance crew.

The zone controller is not required to provide a direct interface to sensors at this level, but the zone controller does provide access to setpoints, limits, deadbands, etc. Individual zones are composed of logical groupings of sensors and actuators that are identified during the commissioning phase. These mappings are stored in non-volatile memory and are accessible to a zone controller. ECoMIC algorithms utilize these mappings to adapt the user environment.

3.4 Backbone network specification

The ECoMIC architecture is designed to be backbone network agnostic to simplify integration with legacy installations. A suitable backbone network has a physical layer that supports bidirectional communication with low-latency and high date rates. Multiple protocols may run atop the physical layer, possibly in a layered fashion. Preferably, the networking protocol provides a mechanism to address individual devices and groups of devices to minimize bandwidth requirements.

Many existing physical layers and protocols meet these requirements. Existing building control protocols, such as BACnet/MS-TP over RS-485, Modbus, BACnet/IP, oBIX or LonWorks, are well suited for this role. Recently IP over Ethernet/Wi-Fi had emerged as a preferred solution for building control networks due to its high-speed, high bandwidth, low-cost, abundant supply, proven technology, wide scale

deployment, well established user and developer community and easy integration with existing IT infrastructure.

3.5 Zone services specification

A zone in the ECoMIC architecture is defined as an area serviced by a single instantiation of the ECoMIC algorithm. An individual zone is comprised of multiple logical groups of sensors and actuators that control a subset of the zone (e.g. a single office in a zone that comprises ten offices). The ECoMIC algorithm optimizes energy usage and comfort for the set of logical groups that make up the zone. Figure 5 shows the zone functional components, and a brief description for each component is provided in the subsection below. The testbed implementations in Section 6 demonstrate instantiations of the specification described herein.




Figure 5 Zone functional components.

3.5.1 Backbone network interface

The backbone interface serves as a physical and protocol interface from the zone controller to other devices in the system via the backbone network. The backbone interface is specific to the type of backbone network. ECoMIC zone controller objects and properties are located at the backbone interface; these properties control the behavior of the ECoMIC algorithm.

Demand response objects

The demand response objects and their associated properties control the load shedding behavior and reporting of the ECoMIC algorithm. Load shedding interacts with the active *Profile*, as described in the section 3.5.4.

Algorithm control objects

The algorithm control objects and their associated properties control the behavior of the ECoMIC algorithm. These objects typically specify the range of legal values, reporting intervals and override behaviors. These properties are specified to keep the visual and thermal comfort environment within the desired range.

Table 2 Algorithm control properties.

Property	Units
Minimum/maximum light level	Percentage
Minimum/maximum temperature setpoint	Celsius
Minimum/maximum slat angle	Radians
Maximum/maximum length of blind/shade	Percentage
Maximum lighting fade rate	Percentage/Second
Maximum slat rotation fade rate	Radians/Second
Maximum shading extension rate	Percentage/Second

Global information

The global information objects and properties provide information about the wider facility environment which the ECoMIC algorithm can use to compute actuator state values.

Table 3 Global properties.

Property	Units
Outdoor temperature measurement	Celsius
Wind speed measurement	Meters/Second
Wind direction measurement	Radians
Direct solar irradiance measurement	Watts/Meter ²
Diffuse solar irradiance measurement	Watts/Meter ²
Exterior vertical illuminance measurement	Lux
Outdoor humidity measurement	Percentage

Discovery

If the backbone protocol selected for implementation supports the concept of device discovery, the backbone interface should implement the corresponding protocol to reduce the effort of system commissioning.

Input and actuator interfaces

All the sensors and actuators within a zone are logically connected to the zone controller in the ECoMIC architecture. The backbone network may be used to transmit sensor or actuator information. The device specific information should be routed to or from the appropriate generic device interface.

3.5.2 Inputs

An ECoMIC zone controller is logically connected to the sensors and local user interface elements in the associated zone. The ECoMIC algorithm uses sensor and user inputs to make actuator setpoint decisions.

Sensor inputs

The sensor interface translates the implementation specific sensor inputs to abstract sensor values suitable for input to the ECoMIC algorithm. The algorithm supports a wide variety of sensor types. The sensor interface may also combine multiple sensor types into an abstract device which may provide input for the algorithm.

The sensor interface is a software API which provides normalized values for sensor measurements, such as temperature, humidity, light level, glare, occupancy state, etc. The sensor interface may connect to multiple protocols or physical layers to accommodate some or all of these sensor types. A given physical interface may be shared with actuators as well as sensors, though the actuator information is routed through the appropriate actuator interface. Each sensor has a separate set of properties which are accessible to the algorithm as illustrated in Table 4.

Table 4 Example of sensor properties.

Property	Units
Illuminance measured value	Lux
Indoor temperature measured value	Celsius
Relative humidity measured value	Percentage
Occupancy	Boolean
Operative temperature measured value	Celsius

UI inputs

The zone controller allows an end-user to control the behavior of the system. The behavior of the system is altered through a user interface that allows users to select scenes, change setpoints, trigger manual override or adjust the actuators (*e.g.* dim the lights, open the blinds). The physical interface and protocols used by the zone controller are implementation dependent.

The UI interface is a software API which abstracts over the physical implementation of the UI. Each sensor has a separate set of properties which are accessible to the algorithm as illustrated in Table 4.

Table 5 User interface properties.

Property	Units
Lights on	Boolean
Current profile	Integer
Light fade rate	Percentage/Second
Lighting level	Percentage
Slat angle	Radians
Slat fade rate	Radians/Second
Blind/shade height	Percentage
Temperature set point	Celsius

3.5.3 ECoMIC algorithm

The ECoMIC algorithm component is the core decision making engine of the architecture. It is responsible for reading sensor values from the sensor interface, computing optimal actuator settings based on the current user profile for each logical group and passing those to the various output interfaces. The algorithm relies on configuration values read from the backbone interface to adjust targets and response times. Details of the developed ECoMIC algorithm is described in section 0.

3.5.4 Profiles

The ECoMIC zone controller supports the concept of profiles describing the desired environment for the individual logical groups in the zone. The **Zone Controller Profiles View**, as shown in Figure 6, consists of four parts: reference inputs derived from sensors, a profile describing the current accepted ranges for the reference inputs for a group, the ECoMIC algorithm which controls the actuator states, and the physical actuators. The goal of the algorithm is to maintain the inputs within the ranges specified in the profile while minimizing overall energy usage and maximizing user comfort.




Figure 6 Zone controller profiles view.

A Profile is a set of the minimum and maximum acceptable values for all input signals, including, but not limited to, light level, temperature, humidity and glare. A Zone Controller has a single active profile at any time. The current profile may be changed by a scheduled event, a remote profile change command, a local sensor input (e.g. change in occupancy state) or a local user interface. Profiles are constructed during the commissioning process to cover the use cases of a particular group within a zone.

In addition to profiles, the zone controller supports overriding individual actuator states. When an actuator state is overridden the state determined by the zone controller is replaced by an externally defined state. This is typically used to allow end-users to manually set the state of an individual system in the zone. In some installations it may be acceptable to limit the range of direct user overrides during commissioning. The ECoMIC algorithm may use the override information to internally tune its behavior, so as to avoid setting an actuator state to a value that will be overridden (see Section 4.6). The lifetime of an override is an implementation dependent decision.

3.5.5 Schedules

The ECoMIC zone controller supports scheduling of profiles to enable automatic activation of the most suitable profile based on the time (e.g. weekend profile or an after-hours profile). A schedule is defined by a start date and time, an end date, recurrence rules, exceptions and the profile to activate. Schedules can be used to conserve energy by dynamically triggering energy efficient profiles suitable for the nature of underlying activities and associated space conditioning requirements (e.g. an after-hours energy conserving profile for cleaning crew or security personnel).

3.5.6 Actuators

The ECoMIC zone controller logically connects to the actuators which drive the physical systems that control the temperature, lighting and shading in a zone. As with the zone sensors, the architecture provides a generic interface to support the widest possible range of actuators.

Thermostat interface

The thermostat interface translates from abstract temperature setpoint commands to implementation specific setpoint commands for a given thermostat. It is a software API, which provides a normalized temperature setpoint object. Table 6 shows the thermostat property accessible to the ECoMIC algorithm.

Table 6 Thermostat properties.

Property	Units
Temperature setpoint	Celsius

Shading interface

The shading interface translates from an abstract slat angle or shade height command to implementation specific setpoint commands for a shading system. It is a software API which provides a normalized representation of slat angle and shade height. An implementation will convert this value to a specific command over a specific protocol to drive an attached shading control system. The properties in Table 7 are accessible to the ECoMIC algorithm.

Table 7 Shading properties.

Property	Units
Slat angle	Radians
Blind height	Percentage
Slat fade rate	Radians/Second

Lighting interface

The lighting interface translates from abstract lighting setpoints to implementation specific commands for an attached lighting system. It is a software API which provides a normalized representation of the typical parameters found in an office lighting system. An implementation will convert these values to specific commands over a specific protocol to drive an attached lighting control system. The possible properties that are accessible to the ECoMIC algorithm are summarized in Table 8.

Table 8 Lighting properties.

Property	Units
Set fade rate	Percentage/Second
Set direct level	Percentage
Lights On	Boolean
Set color temperature	Kelvin

4 Multivariate integrated controls – ECoMIC algorithm

4.1 Framework

The control framework adopted a closed-loop control strategy, where sensor feedbacks are used for making control decisions. Figure 7 shows the operational diagram of the control system. The center ***MIC Controller*** block contains the main required components and their connection to the peripheral blocks. The decision-making engine is the ***Energy Balance Model*** located within the MIC Controller block, which hosts the ***Lighting Load Balancing*** algorithm for determining the optimal electric light level, blind height and slat angle as well as the ***Thermostat Setpoint Selection*** algorithm for determining the optimal thermostat temperature setpoints. An overview of the two algorithms is provided below.

The ***Zone*** block represents the physical space equipped with electric lighting and venetian blind systems as well as a thermostat. The ***Zone Sensing Infrastructure*** block is composed of sensors that monitor various environmental conditions in the controlled zone and feed the information back to the controller. In addition, the ***Global/External Variables*** are also measured and/or calculated to inform the controller for making the optimal decision. The models embedded in the controller (described in the next section) entail the type of information that needs to be provided by the interior and exterior sensing infrastructure. Furthermore, the control framework is designed to be flexible enough to accept controls on the supervisory level in the future, such as demand response signals, etc., as the ***Supervisory Signals*** block signifies in the figure.




Figure 7 Operational diagram of the control framework.

4.2 Control systems and variables

Three systems that are directly related to occupants' visual and thermal environment are controlled under this framework:

- Electric lighting system,
- Window shading system, particularly venetian blinds,
- Thermostat.

Dimmable electric lighting system is considered to provide maximal flexibility and granularity for illuminating the task. This closed-loop control framework is not restricted to any particular light source or luminous intensity distribution of a lighting system. Motorized venetian blind is the focused shading device for regulating daylight and the corresponding solar heat gain. The uniqueness of venetian blinds, in comparison with other operable shading devices, is that they provide two degrees of freedom, shade coverage and occlusion, for more flexible daylighting control while retaining certain level of view to the outside. The blinds can be either interior or exterior to the building fenestration, and the control algorithm is capable of working with either type of blind. Thermostat is the direct link to the HVAC system in this control framework. The control of thermostat setpoint(s) explicitly addresses considerations on HVAC energy usage and occupant thermal comfort.

Four control variables naturally follow the selection of the three control systems:

- Electric light dimming level,
- Venetian blind height,
- Venetian blind slat angle,
- Thermostat setpoint.

Figure 8 illustrates the interrelationship among the control variables, the controlled systems, energy and comfort.




Figure 8 Interrelationship between systems, control variables, energy and comfort.

4.3 Modeling (w.r.t. control variables)

In order to describe the interrelationship among the control variables, energy and comfort, a model-based approach was adopted such that each block in Figure 8 can be mathematically linked together for control purposes. In particular, analytical models of lighting electric load, electric lighting heat gain, fenestration solar heat gain, task light level, glare, and thermal comfort were required.

4.3.1 Task illumination

Adequate light level at task is one of the major factors contributing to visual comfort. The overall task illuminance is the combined contribution from electric light and daylight, each of which can be described by separate models and summed together.

4.3.1.1 Electric light contribution

The contribution from electric light to horizontal illuminance on the workplane (i_E) is modeled as a linear function of lighting output level as shown in equation (1), where x is the light output level normalized between 0 and 100%, and a represent the maximum achievable task illuminance by electric light only.

$$i_E(k) = ax(k) \quad (1)$$

Though some dimmable lights are non-linear around the lower and upper limits of their dimming curves, this simple model represents most of the operational points of a lighting system. This abstract model is sufficient to define most dimmable electric lighting system, independent of light source and luminaire. The model can easily be updated to represent a lighting system with specific dimming characteristics.

4.3.1.2 Daylight contribution

A model describing how daylight enters the controlled space from the fenestration system and falls onto the workplane has been developed. The model can be conceptually expressed as equation (2). The function f_D contains all the calculations in the model. The time-dependent variables θ, z, E_x^b and E_x^d represent the blind slat angle, blind height, exterior direct illuminance and exterior diffuse illuminance, respectively. The time-invariant variable ζ carries information about the room dimension, surface reflectances, task location, etc.

$$i_D(k) = f_D(\theta(k), z(k), E_x^b(k), E_x^d(k), \zeta) \quad (2)$$




Figure 9 Concept of estimating task light level resulting from blind height control.

The two control variables that affect task illuminance from daylight are blind height and blind slat angle. Different blind height positions are accounted for by dividing the window into two portions – a shaded portion and an unshaded portion as shown in Figure 9. Blind height governs the division of the shaded and unshaded parts of the window, and blind slat angle determines the effective transmittance of the

shaded part. The calculation of effective transmittances is adopted from [3]. The main inputs to this model are the exterior direct and diffuse illuminances (acquired by sensors). The model also require room dimension, surface reflectances and task location to determine the amount of daylight reaching the workplane.

This daylight model estimates the overall daylight level on the task plane from three sources of contributions: 1) direct sun, 2) diffuse daylight directly from the fenestration area, and 3) interreflected daylight from the room surfaces as depicted in Figure 10.




Figure 10 Estimating daylight level on a reference point from three contributions: 1) direct sun, 2) diffuse daylight from the fenestration area, and 3) interreflected daylight.

Contribution from direct sun is a geometric calculation, and any point that falls within the sunlit patch is assumed to receive the same illuminance from transmitted direct sunlight. This part is, however, not critical to the control framework since direct sun is always considered as disturbing glare and should be blocked from the task area. *Configuration factor* is used to estimate the contribution from diffuse daylight directly from the fenestration area. This method characterizes the attenuation of daylight along the depth of the room as well as the side-to-side variations. Contribution from interreflected daylight is estimated using the *split-flux method*, which splits the incoming light flux into two parts, one goes up into the ceiling cavity and the other goes down into the floor cavity. The interreflection results from the reflected light being bounced back-and-forth between surfaces and eventually falling onto the workplane. This method assumes the interreflected daylight is evenly distributed across the entire space at the same height.




Figure 11 Task illuminance from the three sources of daylight contribution and the overall task illuminance with respect to (a) different task (reference point) locations in the room at fixed blind height and slat angle; (b) different blind heights at fixed slat angle and task location; (c) different slat angle at fixed blind height and task location.

Figure 11 shows the magnitude of the three sources of daylight contribution under a certain sky and daylight condition (direct and diffuse illuminances on the façade) as well as the total task illuminance, which is the summation of the three contributions. In the plots, blind height is presented in percentage of deployment, *i.e.* 0% means the blind is completely retracted and 100% means the blind is fully extended. Also, 0° slat angle corresponds to horizontal positions for the blind slats. Notice that since the contribution from direct sun is undesirable, the example in Figure 11 is a sky condition where direct sun is not visible from the window, *i.e.* the illuminance contribution from direct sun is zero. Under the same sky condition, the task illuminance is simultaneously affected by blind height, blind slat angle and task (reference point) location; therefore, each plot in Figure 11 demonstrates the task illuminance with respect to the change in one factor while keeping the other two constant for easy visualization.

Figure 12 illustrates the task daylight illuminance at a particular point in the room with respect to different blind slat angles at three different blind height positions calculated from the developed model. It is intuitive that the task illuminance is independent of blind slat angle when the blind is completely retracted (red curve). Also notice that the impact of slat angle diminishes with the retraction of the blind, and thus the green curve (50% lowered blind) in Figure 12 is less steep than the blue curve (completely deployed blind). This shows the interplay between blind height and blind slat angle.




Figure 12 Task daylight illuminance vs. blind slat angle at different blind height.

4.3.2 Electric lighting load

Electric lighting load is one of the energy indexes to be minimized in the control algorithm. It is assumed to be linearly proportional to electric light output level as illustrated in Figure 13, where x is the light output level between 0 and 100%, i_H is the fraction of the maximum power consumption of the lighting system when full on, and i_L is the fraction of the minimum power consumption when turned off. The total electric lighting load is then the fraction of the maximum power consumption corresponding to the light level multiplied by lighting power density (W , in W/ft^2 or W/m^2) and the total floor area (A_f , in ft^2 or m^2) of the zone as shown in equation (3).

$$E(k) = WA_f(cx(k) + i_L) \quad (3)$$




Figure 13 Linear relationship between light output level and electric lighting load.

This model is independent from any particular light source, and provides flexibility for the same control framework to work with any lighting system. A specific lighting system will dictate the exact values of i_H and i_L , hence c , in equation (3), and W can be determined from the designed lighting layout or field survey.

4.3.3 Heat gains and cooling/heating load

Lighting related cooling load primarily comes from two sources: electric lighting heat gain and fenestration solar heat gain. These are modeled as follows.

4.3.3.1 Electric lighting heat gain

Electric lighting heat gain is generated from the lighting hardware and eventually becomes a cooling load, which is the other energy index this control algorithm aims to minimize. Electric lighting heat gain is directly tied to electric lighting load as the power consumed by the lighting system will end up being dissipated in the form of heat. Therefore, the electric lighting heat gain model in the following equation takes exactly the same form as that in equation (3).

$$q_L(k) = WA_f(cx(k) + i_L) \quad (4)$$

Notice that both lighting electric load and electric lighting heat gain can potentially be modeled in more sophisticated forms taking into account any nonlinearity, discontinuity (between minimum dimming level and off), etc. However, it does not seem to bring out a significant benefit for the control algorithm at this stage.

4.3.3.2 Fenestration solar heat gain

Fenestration solar heat gain is another lighting-related heat gain that eventually becomes a cooling load to be minimized¹. Fenestration solar heat gain is determined by the amount of solar radiation being admitted into the room through the windows and venetian blinds. The fenestration solar heat gain model adopts the work of Wright et al. [3]-[7], and this section only provides qualitative descriptions.

¹ Minimizing cooling load is the strategy to minimize lighting-related HVAC energy consumption in cooling seasons. When in heating season, cooling load may help relieve the load on the heating system, and hence reduce the HVAC energy consumption.

Equation (5) illustrates the general form of the model and all the modeling complexities are included in the function f_s . The variables θ, z, I^b and I^d are the blind slat angle, blind height, direct solar irradiance, diffuse solar irradiance, respectively. ψ contains all the time-varying environment variables, including solar profile angle, wind speed, wind direction, indoor and outdoor temperatures, etc. The time-invariant variable η carries information including window glazing optical and solar properties, slat properties, and other window-related properties.

$$q_s(k) = f_s(\theta(k), z(k), I^b(k), I^d(k), \psi(k), \eta) \quad (5)$$

The two inputs to this model are direct and diffuse solar irradiances received on the exterior of the fenestration system. The portion of the solar radiation that is allowed to enter the space is determined by the effective solar and longwave properties of the fenestration system. Different blind height positions are accounted for by dividing the window in the same way as shown in Figure 9. The transmittances of the unshaded part of the fenestration system are constants and are governed by the properties of the window glazings as well as the filled gas in the case of a multi-layer window. In addition to the determining factors for the unshaded window, the transmission of the shaded part is also governed by the effective transmittances resulting from a particular blind slat angle and solar position as shown in Figure 14. In this model, the direct and diffuse solar irradiances are transmitted through or reflected from the blind in three separate mechanisms: direct irradiance transmitted/reflected as a direct irradiance, direct irradiance scattered by the blind and transmitted/reflected as diffuse irradiance, and diffuse irradiance transmitted/reflected as diffuse irradiance. The effective direct-to-direct, direct-to-diffuse and diffuse-to-diffuse transmittances are designated as τ_{bb} , τ_{bd} and τ_{dd} , respectively, and the effective reflectances are designated as ρ_{bb} , ρ_{bd} and ρ_{dd} . This model has been verified against ISO 15099 [8] as described in the original papers [9].




Figure 14 Effective properties (transmittances and reflectances) of a venetian blind with respect to (a) different solar profile angle and (b) different slat angle.

Figure 15 (a) illustrates solar heat flux (solar heat gain of a unit area) with respect to different slat angles at four solar positions for a complete fenestration system with completely deployed blind under a particular combination of direct and diffuse solar irradiances. Figure 15 (b) shows how the solar heat gain varies with blind slat angle at three different blind height under another solar radiation condition.

Notice that the impact of blind slat angle on solar heat gain diminishes with the retraction of the blind as the covered area decreases. Therefore, the green curve is less steeper than the blue curve, and the red curve is independent of slat angle when the blind is completely retracted.




Figure 15 (a) fenestration solar heat flux vs. slat angle; (b) fenestration solar heat gain vs. blind slat angle at different blind height.

4.3.3.3 Cooling/heating load conversion

The heat gains from electric lighting and solar radiation will eventually become cooling load of the HVAC system. Part of the instantaneous heat gain is convective and will show up as cooling load immediately. The other part of the heat gain, the radiant heat gain, will be absorbed by the surfaces, structures, furniture, etc., in the room and re-radiated as cooling load at a later time. The convective-radiant split of the instantaneous heat gain depends on the characteristic of the source, and the time constant for the radiant heat gain to reappear as cooling load is determined by the material and properties of the objects that absorb it.

The heat gain to cooling load conversion is modeled using a first-order difference equation in this control framework as shown in equation (6), where Q is the cooling load and g represent the instantaneous heat gain. In the context of cooling load estimation, equation (6) is usually referred to as the *room transfer function* and is documented in earlier versions of ASHRAE Handbook of Fundamentals. It is obvious from the equation that the current cooling load depends on the current heat gain as well as previous heat gain and cooling load. The coefficients, w_1 , v_0 and v_1 carries information of the building thermal mass, the convective-radiant split, etc. Consequently, cooling loads from electric lighting heat gain and fenestration solar heat gain are characterized by two different sets of coefficients.

$$Q(k) = w_1 \cdot Q(k-1) + v_0 \cdot g(k) + v_1 \cdot g(k-1) \quad (6)$$

The time series in Figure 16 illustrates how instantaneous heat gain is converted into cooling load over the duration of a summer week. Notice that the peak magnitude of the instantaneous heat gain is much larger than the cooling load, but the areas under the two curves will be roughly the same if integrated. Therefore, if blindly supplying cooling power solely in response to the instantaneous heat gain, it is very likely to result in over cooling as well as waste of energy.




Figure 16 Time series of cooling load and instantaneous heat gain.

4.3.4 Glare

The glare control strategy consists of two parts. The first part is to block direct sun penetration, which is primarily a geometry problem. The second part is to limit vertical illuminance from the window. An empirical model was derived from simulations that correlates glare at a point in the space to the vertical illuminance measurement from a physically feasible sensor location on the back wall. Glare modeling and estimation is summarized in this section, and detailed characterization and derivation can be found in Appendix A and Appendix B.

Direct sun control

Direct sun control becomes necessary when the sun is directly visible from the window. The location of the sun in terms of altitude and azimuth can be computed as a function of the time of day. A sensing device has been prototyped, as will be discussed in Section 6.2.1, for detecting the illuminance from direct sun on the window. Alternatively, vertical global illuminance on the window could provide a reasonable indication if a good threshold value is used to discern the disturbing direct sun condition.

The strategy for deploying the blind requires that the sun azimuth is less than 90 degrees from the window normal (sun in the plane of the window), and the illuminance from direct sun on the window exceeds a target value, D . Under variable sky conditions, the illuminance may drop below the target value (sun behind a cloud), and then abruptly rise again. To avoid excessive activity of the blind, the blind is not retracted until the vertical global illuminance on the window falls below a threshold, R , which is a condition free of glare from either direct sun or strong diffuse daylight on the window. In practice, additional time delays may be set for blind deployment and retraction to further constraint excessive blind motion.

Partial deployment may be useful if the window extends below the work area, or if the worker does not directly face the window and does not see the direct sun except possibly at certain times of the day. The level of partial deployment will need to be carefully designed depending upon the window size, location and the workstation arrangement to prevent undesired penetration of direct sun.

Once the blinds are deployed, the control system needs to determine the blind tilt angle that is needed to actually block the sun view. The blinds can either be tilted in a negative direction with the inside of the blind slats lower than the outside, so that light is directed to the floor, or in a positive direction so that light is directed towards the ceiling. We limit blind tilt to the positive direction from 0 (horizontal slats) to +90 degrees. Negative blind tilts are not as useful at blocking direct sun. For all solar altitudes above 0 degrees (sun on the horizon), negative blind tilts that blocked the direct sun admitted much less light than positive blind tilts that blocked the sun. A further strike against the use of negative blind angles is that they resulted in a much poorer correlation between actual glare and glare sensor readings, which makes it more difficult to control glare.

The blind angle needed to block the direct sun can be estimated by assuming that the blind slats are flat. Real blinds are generally concave down, which reduces the total transmitted light at any given blind angle, but normal blind curvatures do not affect the cut-off angle for sources such as the sun that are at or above the horizon line. The blind tilt angle, θ , needed to block the direct sun, depends upon the ratio, r , between the width of the slat to the spacing between the slats, and the sun profile angle, ϕ , which in turn is a function of solar altitude (β) and azimuth (α).

$$\begin{aligned}\phi &= \tan^{-1} \left(\frac{\tan \beta}{\cos \alpha} \right) \\ &= \tan^{-1} \left(\frac{1 - r \cdot \sin \theta}{r \cdot \cos \theta} \right)\end{aligned}\quad (7)$$

In the case where $r = 1$, the second equation can be solved for θ directly, giving:

$$\theta = 90 - 2 \cdot \phi \quad (8)$$

If $r > 1$, this equation will always overestimate the blind tilt needed to block the direct sun. If tighter control is desired, in order to get the maximum amount of daylight, the tilt angle can be approximated to within a maximum error of just over 1 degree by rational fractions:

$$\theta \approx \frac{(C \cdot x + B) \cdot x}{1 + A \cdot x} \quad (9)$$

Where $x = (\tan^{-1}(1/r) - \phi)/\tan^{-1}(1/r)$, and A , B , and C are fitted constants that depend on r as shown in Table 9.

Table 9 Constants and the maximum blind tilt angle for various r value.

r	A	B	C	max angle
1.1	-0.5102	77.797	-45.774	65.38
1.2	-0.5620	67.355	-42.633	56.44
1.3	-0.6084	59.582	-39.889	50.28
1.4	-0.6665	53.362	-38.161	45.58

The constants in the above table were chosen with the constraint that the fitted values never underestimate the tilt angle needed to block the direct sun.

Limit vertical illuminance from window

If the illuminance at the eye from the window exceeds an observer's glare sensitivity level, glare can occur even if the view of the sun is blocked, or the sun is out of the plane of the window. Figure 17 shows that the illuminance that correlated to a subjective impression of window brightness on the borderline between noticeable and disturbing glare varied over a 10:1 range, with 40% of the subjects reporting that glare was at or below this level at 2,000 lux. A study on blind behavior found that about half the subjects pulled the blinds at about 2,000 lux. [10]. Therefore, 2,000 lux is chosen as the default target glare threshold, but given the wide range of responses shown in studies, allows this target level to be user modified. Also notice that the 2,000 lux default value is derived assuming the occupant facing directly at the window, and therefore, it should be adjusted accordingly if the occupant sits at an angle to the window.




Figure 17 Comfort probability vs. vertical illuminance.

When the direct sun component is excluded, the vertical illuminance facing the window drops in a well-defined manner as a function of the distance from the window wall. The general features of this trend are that the illuminance drops follows an inverse quadratic function of distance from the window wall. The coefficients of the inverse quadratic are relatively insensitive to the shape of the room, if the distance from the window wall is expressed as a cavity ratio, CR , where:

$$CR = \frac{5 \times \text{distance} \times (\text{height} + \text{width})}{\text{height} \times \text{width}} \quad (10)$$

The coefficients of this function, $f(CR)$, are insensitive to sky conditions, blind tilt, room surface reflectances and room shape. The function does, however, seem to be somewhat sensitive to the presence of large workstations. Changing the coefficients to give a more rapid fall-off in the portion of the room in front of the workstations reduced the root-mean-square deviation of the fit from 23% to 14% based on two large room simulations.

In addition to the drop-off of illumination as a function of the distance from the window, there is also a drop-off as a function of the distance from the center-line of the room to the side wall. The illuminance distribution from side-to-side is within a few percent until close to the side wall, where we found a drop

of 40%. We do not have a full model of this effect, but we have seen it in both our small and large room simulations. The most extreme deviations from the basic distance fit are due to this drop-off at the edges.

Our fits were based on a fixed monitoring sensor position facing the window at six feet height on the back wall, and centered between the two side walls. The accuracy of the fit is insensitive to raising the sensor to seven feet, or moving closer to the side wall, as long as it was further than three feet from the side wall. The default fit can be used directly with the rear monitoring sensor to get an estimate accurate to about a factor of two for the vertical illuminance for any other given location in the space. Equation (11) shows the fit, where E is the glare illuminance, M_1 is the illuminance on the glare sensor, and k is a constant. Including the side-to-side effect substantially improves the accuracy, but still leaves the trend errors in the large rooms.

$$E = M_1 \cdot k \cdot f(CR) \quad (11)$$

Adjusting the fit coefficients on a case by case basis is impractical, and irrelevant due to the large differences in glare sensitivity that are expected. The important issue is that when we did fit the trends optimally, the residual error was on the order of 10%. This means that when the occupant adjusts the target glare level to their preferred value, the estimate for what the value actually is may be off by a factor of two, but the closed-loop system can always control the illuminance to on the order of 10% of the desired level. The relative error is small, even though the error in the estimate of the absolute value may be large.

For the ECoMIC algorithm, glare is maintained at or below the target level by adjusting blind tilt in discrete steps. When window glare reaches the target level, the blind tilt angle is increased to reduce the glare illuminance. To prevent the blinds from hunting between the two tilt angles the control algorithm has to estimate the glare illuminance expected if blind tilt is decreased by one step, and compare it to a target level. Based on our simulation study, the illumination on the glare sensor, M_1 , at an arbitrary slat angle (θ) fits the following simple formula:

$$M_1 = M_{10} \cdot \exp((4 \cdot z - 1) \cdot \cos \theta) \quad (12)$$

where M_{10} is the M_1 illuminance level when the blind tilt angle (θ) is 0° , and z is the blind deployment ratio.

4.3.5 Thermal comfort

Fanger's thermal comfort model and the corresponding Predicted Mean Vote (PMV) rating, which is on a scale between -3 and 3, were adopted in this control framework as the main link to thermostat control. This particular model is based on the steady-state thermal load of human bodies, and PMV correlates the body thermal load to thermal sensation. The model can be expressed in the following form with six factors, T_a , T_{MRT} , v , RH , met and clo , representing air temperature, mean radiant temperature, air flow rate, humidity, activity level and clothing insulation, respectively. This has been documented in detail by the original developer in [11].

$$PMV = f_T(T_a, T_{MRT}, v, RH, met, clo) \quad (13)$$

A PMV value of 0 corresponds to a neutral thermal sensation. Positive and negative PMV values represent warm and cool sensations, respectively, and the farther away from 0, the stronger the sensation. A 0 PMV value correlates to 95% satisfaction of the general population². A PMV value within ± 0.5 , ± 0.7 and ± 1.0 corresponds to 90%, 85% and 80% satisfaction rate, respectively.

Based on equation (13), the same thermal sensation, *i.e.* PMV value, can be achieved by different combinations of the six determining factors. In other words, the change in one or more factors can be offset by other factors to retain the same PMV value. This is one of the main concepts exploited in the ECoMIC control framework for thermostat setpoint control. Figure 18 shows the relationship between air temperature (T_a) and mean radiant temperature (T_{MRT}) for the same comfort level when the other four factors are fixed. Figure 18 suggests that the thermostat setpoint, hence the air temperature, can be adjusted along the direction of the orange arrow for the same comfort level in response to changes in mean radiant temperature. In the meantime, since the comfort conditions within the paired satisfaction boundary lines are all considered acceptable, the thermostat setpoint may be pushed towards the boundaries along the blue arrow to relieve the demand on HVAC system for more energy savings. This is the second concept for thermostat control in the control framework.




Figure 18 Air temperature vs. mean radiant temperature for the same thermal comfort level.

4.4 Lighting load balancing algorithm

The lighting load balancing algorithm is formulated as an optimization problem to minimize the lighting and lighting-related HVAC energy demand while maintaining the specified task illuminance and avoiding daylight glare. The formulation is expressed in equation (14).

² PMV is an empirical rating that correlates body thermal load to thermal sensation, which is a subjective perception. Some people may not feel thermally neutral even at zero body thermal load. It is, therefore, impossible to create a thermal environment to satisfy everyone, and zero PMV value only corresponds to 95% satisfaction.

$$\begin{aligned}
\min \quad & E(k) + m \cdot (Q_L(k) + w \cdot Q_S(k)) \\
\text{subject to} \quad & \mathbf{e}_{min} \leq \mathbf{e}(k) \leq \mathbf{e}_{max} \\
& \mathbf{g}(k) \leq \mathbf{g}_{thd} \\
& \mathbf{x}_{min} \leq \mathbf{x}(k) \leq \mathbf{x}_{max}
\end{aligned} \tag{14}$$

Energy demand minimization is posed as the objective function in equation (14), where $E(k)$ represents the lighting electric load at the time step k , $Q_L(k)$ and $Q_S(k)$ account for HVAC load from electric lighting and solar heat gain, and the factor m equates the HVAC load to electric load. In addition, w is the season indicator in response to different HVAC operating modes. In cooling mode, solar heat gain will eventually become cooling load that has to be removed by the HVAC system; on the contrary, solar heat gain will relieve the heating load of the HVAC system in heating mode. The comfort aspects are addressed as the constraints in the optimization problem. The second line in equation (14) is meant for regulating comfort factors $\mathbf{e}(k)$, specifically the task light level, within a specified range between \mathbf{e}_{min} and \mathbf{e}_{max} . The third line in (14) ensures the comfort factors $\mathbf{g}(k)$, the glare condition in particular, do not exceed a given threshold \mathbf{g}_{thd} . The last row in (14) defines the physical and practical ranges of movement ($\mathbf{x}_{min}, \mathbf{x}_{max}$) for all the control variables $\mathbf{x}(k)$. Notice that $\mathbf{e}, \mathbf{g}, \mathbf{x}$ and the corresponding limits can be vectors if there are multiple variables in them. It should become obvious that this formulation is instantiated by plugging the models described in the previous section.

There are at least two ways to determine the value of the season indicator, and both are included in the operational diagram in Figure 7. The most straightforward method is to obtain the HVAC operating mode directly from the HVAC system or the energy management and control system (EMCS) that controls it if such connection can be established. The other method is to infer the HVAC operating mode by comparing the outdoor (T_o) and indoor supply air (T_s) temperatures as the sample rules below illustrates.

IF $T_o < 60^{\circ}F$ AND $T_s > 70^{\circ}F$	THEN HVAC mode = HEAT
IF $T_o > 75^{\circ}F$ AND $T_s < 65^{\circ}F$	THEN HVAC mode = COOL
IF $70^{\circ}F < T_o < 75^{\circ}F$ AND $65^{\circ}F < T_s < 70^{\circ}F$	THEN HVAC mode = SWING

4.5 Thermostat control algorithm

The thermostat control algorithm is based on Fanger's thermal comfort model and the Predicted Mean Vote (PMV) derived from it as described in section 4.3.5.

The central idea of thermostat setpoint control is to exploit the following two facts.

- The change in one thermal comfort factor can be offset by other comfort factors to achieve the same comfort level.
- All thermal comfort factors together determine a range, within which the comfort condition is considered satisfactory.

Optimal thermal comfort control is achieved by finding the highest/lowest thermostat setpoint in cooling/heating season that can retain the specified comfort level while resulting in the least demand on the HVAC system. The formulation is shown in equation (15), where T_a is the indoor air temperature, and hence the thermostat setpoint, and κ_{min} and κ_{max} represent the range for the level of thermal

satisfaction, which can be a nominal value or a level specified by the user. Notice that the *PMV* is also a function of other five thermal comfort factors, which are held constant in the equation.

$$\begin{aligned} \arg \min_{T_a} PMV(T_a) &:= \{T_a \mid \kappa_{min} \leq PMV(T_a) \leq \kappa_{max}\} && \text{for heating} \\ \arg \max_{T_a} PMV(T_a) &:= \{T_a \mid \kappa_{min} \leq PMV(T_a) \leq \kappa_{max}\} && \text{for cooling} \end{aligned} \quad (15)$$

4.6 User preference learning

User preference learning is in a different scope from the control algorithm depicted in Figure 7. In the ECoMIC architecture, this functionality is meant for dynamically tuning the policy in the ***Profile*** as described in Section 3.5.4, which provides the related information to the ***User Preferences*** block in Figure 7 as one of the supervisory inputs to the control algorithm.

4.6.1 User preference learning strategy

The specific technology developed is called a ***user preference learning machine***, which connects to the zone controller and interacts with users as the conceptual drawing in Figure 19 shows. The automatic learning process includes adjusting different lighting aspects at time intervals chosen stochastically according to a prescribed distribution to solicit user's feedback, which may be in various forms. User preference is recorded in a database for different conditions such as time, day, week, weather, sensor readings, geographical measures, user IDs, task IDs, etc. Upon completion of the learning process, the machine will provide a model of the user preference described in terms of a policy and recorded in a database.




Figure 19 Schematic representation of the interaction between the learning machine, the controller and the user.

4.6.2 Preference learning algorithm

Figure 20 shows the flow chart of the learning algorithm. The user preference learning machine rely on a “reward” mechanism to acquire users’ preference through three user-machine interactions: adjustment interaction, mutual interaction and user interaction. A reward is granted if a user responds to an “state” created at a certain time of day, light level, blind position, temperature setting, etc., as satisfactory.




Figure 20 User preference learning algorithm flow chart.

Adjustment Interaction happens through a Gamma process. At each occurrence, the machine issues an action without a chance of being eligible for reward except possibly through user interaction.

Mutual Interaction happens through a Gamma process. At each occurrence, the machine issues an action and the user must report satisfaction/dissatisfaction within a given time frame.

User Interaction happens through user voluntarily satisfaction/dissatisfaction reporting. User can do so at any time. Upon such interaction, the initial zero reward reported in adjustment interaction immediately preceding this interaction is updated.

The learning record contains a set (L) of observed learning triplets in the form of (s, a, R) , where s is the state including time of day, light level, blind position, temperature setpoint, etc., in an instance; a is the action comprising the new light level, blind position, temperature setpoint, etc.; and R is the reward. In particular, for a given state-action pair, we have:

$$L(s, a) = \{R_1, R_2 \dots\}, \quad (16)$$

which are the observed rewards over the learning history. The state-action policy for each state is updated by an epsilon-greedy reinforcement method to a distribution of stat-actions that satisfies users with maximum probability given a particular state, i.e. $\mathbb{P}(a|s)$.

Having sufficient learning records refers to the number of learning triplets added to the learning records since the last policy update. For a typical value of Gamma distribution, sufficient learning records will be collected over the course of a day and so, policy update may happen on a daily basis. Figure 21 illustrates the technology in learning a 500 lux preferred task illuminance over a course of 3 days.




Figure 21 Example of learning a 500 lux task illuminance preference over a course of 3 days (left to right). Top plots: the probability density function of the preference. Bottom plots: learned task illuminance preferences throughout a day.

5 Simulation-based evaluation

Different aspects of the control framework have been studied and evaluated through simulations. With the complexity of the control algorithm as well as the huge combinations of the simulation factors, it was not possible to simulate all aspects and assess the performance with a single simulation instance. The inherent limitations of the simulation tools also constrained the type of simulation-based performance evaluations. Therefore, several different simulation-based investigation were conducted and documented in this section, each serves a specific purpose that will be clearly stated in following subsections.

5.1 System behavior evaluation using step response simulation

Steady state response simulation assumes all exterior conditions stay the same throughout the duration of the simulation and study the final state of the control system. It is the most suitable method for examining the basic behavioral characteristics of the control algorithm. Notice that the time steps in steady state response simulations do not correspond to any specific physical time interval, and the actual duration of each time step can usually be freely adjusted at will.

Figure 22 shows the steady state response of the control algorithm for a target task light level of 500 lux. The electric light level and blind slat angle were dynamically controlled while the blind was assumed to be fully deployed at all time. This simulation started with arbitrary electric light level and slat angle, 50% electric light level and 45° slat angle to be specific. As can be observed in Figure 22 (a), electric light level and slat angle converge to the optimal setting over time. Figure 22 (b) shows the corresponding changes in the contributions of electric light (blue curve) and daylight (green curve). Even though the lighting contributions from the two sources kept on changing during the transition period, the overall task light level (red curve) was always maintained at the target 500 lux. The total energy demand, including electric lighting load and lighting-related cooling load, is plotted in Figure 22 (c). The starting point with high energy demand was obviously not an optimal point due to high energy demand, and the algorithm managed to bring the demand to the minimum possible value. Notice that although the dip around the 90th time step seems to have the minimal energy demand, it is not an achievable steady state minimum since the cooling load was still building up over time.




Figure 22 Steady state response with an arbitrary starting point. (a) Electric light level and blind slat angle; (b) maintained light levels; (c) energy demand.

The steady state response of the control algorithm when incorporating glare control is illustrated in Figure 23. The target task light level is still 500 lux, and the threshold glare is specified as 1500 lux of

vertical illuminance at observer's eye. Two exterior lighting conditions were considered and the resulting responses are shown in Figure 23 (a) and (b), respectively. The exterior lighting condition in the first case did not cause any glare concerns since the vertical illuminance at view point (red curve in Figure 23 (a) bottom plot) was only 750 lux while the task illuminance (blue curve) was regulated at 500 lux. The glare illuminance was predicted from the vertical illuminance at the back wall sensor location (green curve) based on equation (11). The electric light (blue curve in Figure 23(a) top plot) was turned off at steady state and the slat angle (green curve) was set to provide the 500 lux task illuminance from daylight. In the second scenario, glare did become a concern and was capped at 1500 lux (red curve in Figure 23 (b) bottom plot) at steady state while maintaining 500 lux task illuminance (blue curve). The blind slat (green curve in Figure 23(b) top plot) was set to block excessive daylight for the 1500 lux vertical illuminance at view point, and, as a result, electric light (blue curve) had to remain on at a certain level in order for the task illuminance to reach the 500 lux target. Notice that these two scenarios were created solely for demonstration and testing purposes. The correlation between horizontal and vertical illuminance was not rigorously modeled, and therefore, the simulated conditions may not be realistically representative.




Figure 23 Steady state response for task lighting and glare control. (a) Vertical daylight illuminance was not high enough to trigger glare avoidance; (b) high vertical daylight illuminance activated glare avoidance.

5.2 System behavior evaluation using simulations based on real measurement

The control algorithm was simulated with a set of real measurements collected from an earlier study. The data contain actual solar and sky conditions that the control algorithm will encounter in reality and, thus, are best for verifying its expected behavior and performance. This set of simulations assumed that the electric light and blind slat angle are dynamically controlled while the blind is fully deployed at all time. Two representative cases are shown in Figure 24 and Figure 25 for a typical clear sky condition and a partly cloudy sky condition, respectively.




Figure 24 Simulation using measurements from a sunny day with a clear sky condition; (a) exterior solar condition; (b) maintained light levels; (c) energy demand; (d) light level and slat angle.




Figure 25 Simulation using measurements from a partly cloudy day with an unstable sky condition; (a) exterior solar condition; (b) maintained light levels; (c) energy demand; (d) light level and slat angle.

Figure 24 (a) and Figure 25 (a) illustrate the exterior solar and daylight conditions that were fed into the control algorithms. Notice that the direct irradiance (red curve) in the plots is the direct normal measurement, and the global irradiance (green curve) is the horizontal global measurement. The clean and smooth curves in Figure 24 (a) signifies the characteristics of a typical stable clear sky condition, while the high-frequency fluctuation in the curves in Figure 25 (a) reflects the unstable sky condition on a partly cloudy day. In spite of the drastically different sky conditions, the red curves in Figure 24 (b) and Figure 25 (b) shows that the overall task illuminances were maintained very close to 500 lux during the day in both cases. The corresponding movements of electric light level (blue curves) and blind slat angle (green curves) are shown in Figure 24 (d) and Figure 25 (d). The resulting electric lighting load (red curves) and lighting-related cooling demand (cyan curves) are plotted in Figure 24 (c) and Figure 25 (c).

In an attempt to demonstrate and compare the behavior of the control algorithm in cooling and heating seasons, the algorithm was slightly modified and applied to the same set of data as that used in Figure 24. In this case, the acceptable task light level is widened to the range between 450 and 550 lux, and w in equation (14) was manually set to +1 and -1 to force the algorithm to run in cooling and heating mode, respectively. When the algorithm was run in cooling mode, solar heat gain was blocked with more occluded blind, and the overall task light level is kept at the lower limit of the target range as shown in Figure 26 (a). When in heating season, Figure 26 (c) shows that the overall task illuminance was regulated at the upper bound of the target range during business hours for admitting more solar heat gain to relieve HVAC system heating load. Different operating modes also resulted in a difference of about 1000W in solar heat gain comparing Figure 26 (b) and (d).




Figure 26 Simulation of the control algorithm runs in different mode. Cooling seasons: (a) maintained light levels; (b) solar heat gain. Heating season: (c) maintained light levels; (d) solar heat gain.

5.3 Energy performance simulation using building energy modeling tools

This particular simulation exercise was closest to reality as a representative building model was used with EnergyPlus simulation, which permits detailed evaluation of energy performance. In particular, the simulation focused on the south zone on the mid floor of the three-story DOE reference medium office building model. EnergyPlus can only implement deterministic control sequences due to its schedule-based nature, and it was impossible to implement the dynamic feedback control algorithm directly. Therefore, the ECoMIC algorithm was realized in Matlab and co-simulated with EnergyPlus through the Building Controls Virtual Test Bed (BCVTB). In this co-simulation framework, the exterior and interior conditions simulated by EnergyPlus, including exterior window wall solar irradiances and daylight illuminances, task daylight illuminance, air temperature, etc., were transmitted to the control algorithm in Matlab as its inputs. The control algorithm then calculated and sent the control decisions, including electric light level, blind slat angle and thermostat setpoints, back to EnergyPlus for calculations in the subsequent time step.

The complete control algorithm except blind height control and glare control were implemented in this simulation. In other words, electric light level, blind slat angle and thermostat setpoints were controlled dynamically in response to the lighting and thermal conditions. The blind was assumed fully deployed at all time since blind height control was not directly supported in EnergyPlus. Detailed lighting simulation for glare control was not possible either as EnergyPlus could not perform any vertical illuminance calculation. The goal was to maintain task light level at 500 lux and maintain thermal comfort within the 90% satisfaction level (PMV within ± 0.5).

In addition to the controlled case with the developed control algorithm, a non-trivial base case was also constructed and simulated in EnergyPlus. The base case was implemented with daylight dimming for 500 lux task illuminance, and the blind slats were set to automatically rotate to the cut-off angle to block direct sun. The thermostat heating and cooling setpoints were fixed at 21°C (70°F) and 24°C (75°F), respectively.

Figure 27 shows the result of a five-day simulation from August 2nd to 6th. Figure 27 (a) plots the exterior solar and daylight conditions as a reference to the weather and sky condition. The 500 lux target task illuminance is well-maintained throughout the simulation period as the red curve in Figure 27 (b) indicates. The same plot also shows the decomposed contributions from electric light and daylight that made up the overall task light level. Figure 27 (c) illustrates how the thermostat heating (red curve) and cooling (cyan curve) setpoints are adjusted in the controlled case. The blue and green curves in the same plot offer a comparison of the thermal comfort conditions between the controlled case and the base case. The PMV values in the controlled case were maintained within ± 0.5 as specified, corresponding to 90% satisfaction. The PMV values in the base case, on the other hand, were mostly negative values using the standard setpoints, which means that the environment was maintained more towards the cool side. This implied not only a lower satisfaction rate on thermal comfort but also more cooling energy consumption, which can easily be verified by comparing the green (base case cooling load) and blue (controlled case cooling load) curves in Figure 27 (d). Compared to the non-trivial base case, the ECoMIC algorithm yielded a 27% reduction in electric lighting load. The average cooling load reduction is 42% with a 56% reduction in peak cooling demand over the simulated period.




Figure 27 Five-day EnergyPlus simulation during summer time. (a) exterior solar condition; (b) maintained light levels; (c) thermostat setpoints and thermal comfort; (d) cooling and heating rate.

5.4 Benefit of multivariate integrated control over the most advanced lighting control technologies

This study aimed to extract the additional benefit of the ECoMIC framework over the most advanced lighting and daylighting control techniques. EnergyPlus served as the main energy simulation tool in this study. However, instead of implementing and simulating the entire ECoMIC algorithm (due to the

aforementioned reasons of complexity and huge simulation combinations), a surrogate approach was developed that allowed us to conduct this investigation in a more efficient way.

Modeling of the most advanced lighting and daylighting control technique was based on a system developed prior to this project, which has been deployed in an Army base as part of a demonstration project sponsored by DOD³. This system integrated a dimmable lighting system, a motorized blind system, photosensors and occupancy sensors. The control objective was to offset as much electric light with daylight as possible while restricting excessive daylight. For example, when the setpoint of task lighting was set at 500 lux, the system would try to provide this light level by opening the blind and only use electric to top off if daylight itself was not sufficient. On the other hand, the blind slats would rotate towards the closed position if the incoming daylight was exceeding the setpoint.

The limitation of the above control technique is that it does not explicitly consider the thermal aspect. The solar heat gain associated with the admitted daylight will become additional cooling load to the HVAC system. As the lighting power goes down due to the advance in light source, luminaire optics, etc., it might actually be more efficient overall to use electric light as the main source of illumination under certain circumstances.




Figure 28 Steady-state total (lighting and HVAC) energy consumption under various lighting power density.

The factors that may affect the tradeoff between lighting energy and lighting-related HVAC energy usage include lighting power, window efficiency (e.g. SHGC, U-factor), shade location, and window size. Figure 28 shows a simulated steady-state energy consumption combining lighting electric load and lighting-related HVAC load under various lighting power density (LPD) at a certain time instance. As the blind closes (slat angle increases from 0° towards 90°), the overall energy consumption continues to decrease. At about 54° slat angle, the task light level reaches the 500 lux setpoint, which is provided using solely daylight. This is the point where the advanced control technique will settle to. For cases with high LPD,

³ The occupancy sensing aspect was not included in this particular study since both control techniques should have identical response to occupancy status. Also, due to the difficulty of simulating immediate blind height positions in EnergyPlus (or any other energy simulation tools), the blinds were assumed to have only two status: completely retracted and fully extended. This is true for the current advanced control technique; however, the ECoMIC algorithm does allow immediate blind height position in practice, which should bring more benefit both in terms of energy and user satisfaction. In addition, the energy benefit of controlling thermostat was not included in the simulation exercise as the comparison would not be fair since the advanced control technique did not have the same capability.

this point indeed results in minimal energy consumption as shown in the plot. However, for cases with lower LPD, the energy usage continuous to decrease as the blind keeps on closing, and the minimum occurs at 90° slat angle when the blind is completely closed. In these cases, electric light is used to supplement daylight for the 500 lux task illuminance. And the ECoMIC algorithm should be able to settle to this point for optimal energy usage.

Based on this observation, we developed a surrogate approach for studying the additional benefit of the ECoMIC framework. The simulations were again based on the DOE reference medium office building model. Cases with different combinations of the aforementioned factors were simulated, where lighting power was expressed as LPD, window size was represented in the form of window-to-wall ratio (WWR). In addition, three climate types were considered: hot-humid, mixed-humid and cold-dry. The diagram in Figure 29 illustrates details considered for each factor. In each combination, two scenarios were simulated: one with the blind always closed during daytime (designated as S1) and one with the advanced control technique (designated as S2). In other words, S1 simulates the condition on the leftmost end of the x-axis in Figure 28, and S2 simulates the condition on the gray vertical line in the plot. One can easily observe from Figure 28 that the lowest total energy point always occur either at the leftmost side (S1) or on the gray vertical line (S2), and ideally, the ECoMIC algorithm should be able to switch between S1 and S2 whichever is more advantageous at the time. With this approach, the additional benefit of ECoMIC would show up in instances where S1 results in less energy consumption than S2.




Figure 29 Combinations of factors considered in the simulations.

The simulation results were compared for each zone in the building. Since all the zones in the building model were served by a single HVAC system, it was impossible to attribute the exact HVAC consumption to each zone. Therefore, we analyzed the sensible cooling and heating load for each zone instead of the actual energy consumption. The differences between S1 and S2 in lighting energy usage, sensible cooling and heating loads were calculated, and a special **Ratio** was defined as follows.

$$Ratio = \frac{\text{Cooling/heating load reduced/provided by S1}}{\text{Lighting energy saved by S2}}$$

How the cooling and heating load actually translates into energy consumption, of course, depends on the efficiency of the HVAC system. However, if the **Ratio** is large enough, it will be a clear indication that S1 has better energy performance than S2. Figure 30 shows the distribution of the **Ratio** for the four perimeter zones on the middle floor with a particular combination of the considered factors. In all four zones, **Ratios** greater than 3 can be observed. And specifically in East and West zones, there are many instances where the **Ratio** is greater than 4, which implies that S1 is performing better than S2 as long as the HVAC efficiency is reasonable. In other words, the ECoMIC algorithm will outperform the most advanced control technique (S2) under those circumstances.




Figure 30 Distribution of the **Ratio** for the four perimeter zones on the middle floor.

In addition to the particular example discussed above, similar observations were made on the results for other combinations of the considered factors and climate types. This simulation exercise clearly demonstrated that there were rooms for better energy performance on top of the current most advanced lighting control techniques, and the ECoMIC algorithm was designed to achieve that.

6 Testbed implementation and evaluation

6.1 Test sites and testing aspects

Two sites were selected for establishing testbeds to demonstrate the developed control technologies. The pictures of the testbeds on each site are shown in Figure 31. Since the two sites had very different characteristics, they were configured to evaluate different aspects of interests as described below. Nonetheless, both testbeds were implemented following the architecture described in Section 3.5



Figure 31 (a) Test site at Philips Lighting, Rosemont, IL; (b) test site at the FLEXLAB at LBNL.

6.1.1 Rosemont

The first test site was located at the Philips Lighting Electronics building, a large-size 8-story office building in Rosemont, IL. This Midwest climate requires cooling for roughly half of the year and heating for the other half. The testbed was established in a room on the 4th floor on the east side of the building with east-facing windows. Due to the orientation of the testbed, the space would be exposed to direct sunlight during morning hours and have moderate level of daylight for the rest of a day. The room is 19.5' wide, 15' deep and 8.5' high with four existing 2'-by-2' dimmable 2-lamp recessed parabolic luminaires. The HVAC in this particular room was isolated from rest of the building by closing valves and dampers, and a set of dedicated portable heater and air conditioner were put in place such that the HVAC system can be independently controlled and metered. Since this room was previously used for storage before converted into the testbed, there was not an actual occupant assigned to work in this space. The control objectives were set identical to those in the simulations performed in section 5.3:

- Maintain the task light at a constant level (500 lux);
- Maintain a comfortable thermal environment;
- Minimize the related energy consumption.

6.1.2 LBNL FLEXLAB

The second test site was located at the Lawrence Berkeley National Laboratory in Berkeley, CA, and was part of the FLEXLAB (the Facility for Low-Energy eXperiments in buildings) facility. The testbed was established in a private office in the perimeter of a southwest-facing building. It was not possible to isolate HVAC of this testbed from rest of the building, and therefore, independently metering HVAC

energy could not be realized. Since the room did not have any architectural exterior shading, it would receive abundant natural light, including strong diffuse daylight throughout a day as well as direct sunlight in the afternoon. This particular characteristic made it an ideal candidate for assessing glare control of the developed technologies, which could not be studied with the simulations in section 5. In addition, as the southwest-facing window provides a view to the San Francisco bay, preserving the view as much as possible also became an objective for the performance evaluation. As a result, glare served as the main governing factor for this testbed. The control objectives were set as follows:

- Maintain the task light above a specified minimum level;
- Prevent glare condition from exceeding a specified threshold;
- Minimize the related energy consumption.

6.2 Hardware

The hardware components, including commercial off-the-shelf products and prototypes, used in the testbed implementation at both sites are described in detail in this section.

6.2.1 Sensors

Photosensors

The horizontal (task illuminance) and vertical (glare) light levels were measured using commercial photosensors from Philips Dynalite (model DUS804C), which are commonly used for indoor light sensing in Dynalite control systems. The DUS804C is a low profile recessed flush mount 360° ceiling sensor that combines motion detection (PIR), infrared remote control reception and ambient light level detection in the one device. The sensors can be read using simple commands over Dynet (a protocol using 2-wire RS485). Only the photosensing functionality was used in the implementation.

In each testbed, one photosensor was mounted on the ceiling above the workplane for estimating task light level (task illuminance), and one photosensor was installed on the back wall about 6' above the ground as a glare sensor. All sensors were characterized using a lighting meter with and without daylight. Figure 32 (a) shows the picture of the mounted photosensors in the test space.

Temperature sensors

Indoor air temperature was measured using an off-the-shelf sensor by Sontay (model TT-911) placed along the center of the side wall. Operative temperature of the space was measured by an off-the-shelf globe sensor by Sontay (model TT-915) also mounted along the same side wall. Figure 32 (b) shows the two temperature sensors installed on one of the side walls.



Figure 32 (a) Photosensors; (b) temperatures (top: operative, bottom: air)

Outdoor environmental sensors

The exterior conditions were measured in real time on or close to the façade of the testbed by an exterior sensor suite and an anemometer. The photometers and pyranometers for measuring sky illuminance and irradiance, respectively, were hosted on a custom-designed weather-proof PVC enclosure. Figure 33 (a) shows the enclosure prototype that was designed and built in-house at the Philips Research Center. An off-the-shelf PVC junction box was chosen as the housing box with further modifications. A custom-made metal bracket held a photometer-pyranometer pair on each of the four sides of the box. Furthermore, the sensors were shaded by non-reflective fans made from Black Delrin Resin to improve measurement accuracy by blocking unwanted reflections from the building surfaces. This sensor arrangement allowed the algorithm to compute direct and diffuse components of solar radiations in real-time. In addition, a shielded thermistor was mounted on the enclosure to measure outdoor air temperature (right-hand side of the box behind the black fan in Figure 33 (a)).

Figure 33 (b) shows the anemometer used for measuring both wind speed and direction as part of the control sensor signals. This was a low-cost off-the-shelf product, and the output signals were wired into the sensor suite enclosure. To avoid practical wiring issues on getting the exterior sensor signals to the inside of the test space, a dedicated sensor board was designed to enable wireless communication using radio frequency. The sensor board was housed inside the enclosure to perform sampling, preprocessing and transmission of the sensor readings to the ECoMIC controller inside the testbed.




Figure 33 (a) exterior sensor enclosure hosting photosensors, pyranometers and temperature sensors; (b) anemometer; (c) exterior sensor enclosure mounted in place.

IP Camera

In addition to the control sensors, the Rosemont testbed was fitted with an IP-enabled camera that captures wide-angle images of the testbed every 60 seconds. The images show the light variations, blind movements, and glare perception. A fileserver was setup in the demo laptop to collect the images from the camera. These images were used to evaluate system behavior and debug issues. The images were also used to generate time-lapse videos to visualize and compare the system response between the *BaseCase* and *ECoMIC* cases.

Simulated occupancy sensor

As mentioned in Section 6.1.1, the Rosemont testbed was not occupied by an actual worker, and the energy benefit from occupancy sensing could not be evaluated. Therefore, a software occupancy model was implemented to simulate real-time occupancy sensing. For a fair performance comparison between the *BaseCase* and *ECoMIC* case, identical models were implemented in both cases. In other words, the occupancy status at any given moment would be exactly the same between *BaseCase* and *ECoMIC* case. A parameterized occupancy model was derived using actual occupancy data collected from an office building, and a minute-by-minute occupancy pattern was generated with 94% occupancy between the hours of 7am and 6pm as Figure 34 illustrates. For simplicity, the same occupancy dataset was used every day throughout the testing period. For the actual deployment of the control system, the model would simply be replaced by the motion detection capability of the Dynalite DUS804C sensor.




Figure 34 Simulated occupancy which is fed to the Rosemont testbed in real-time.

6.2.2 Control hardware

Venetian blinds and drivers

The motorized venetian blinds were manufactured in different length and width to fit the window areas in the testbeds. The blind slats are 50 mm wide and 42 mm apart, which allows a slight overlap between the slats to prevent leaking of daylight when the blind is completely closed. The deployment and occlusion of each blind was driven by a Somfy tubular motor integrated in the blind header. Four and two blinds were used to cover the entire window area at the Rosemont and LBNL testbed, respectively, and the motion of all the blinds at each site was synchronized.

The Dynalite DDNG485 network gateway modules, shown in Figure 35, functioned as blind drivers to drive the Somfy blind motors. This particular device was originally designed to enable the integration between a Philips Dynalite control system and third-party devices. In this implementation, the DDNG485 was repurposed as a “bridge” between Dynalite’s native DyNet protocol and the blind motor’s Somfy network with a customized firmware. The blind driver supported the following commands and queries along with several advanced commands for commissioning purposes.

- Set blind height (0 – 255 levels)
- Stop blind motion
- Set blind slat angle (0 – 180°)
- Request/Reply current blind height
- Request/Reply current blind slat angle



Figure 35 DDNG485 network communication gateway.

HVAC system (Rosemont testbed)

As there was no way to directly measure or disaggregate HVAC energy consumption of the testbed from the rest of the building HVAC system, an independent mini HVAC system was a logical choice in an attempt to quantify the impact of the control system on HVAC energy usage. The test space was first isolated from the building's HVAC system, and a portable heater and air conditioner were put in. The air conditioner provided a 10,000 BTU cooling capacity, and the fan speed was fixed at the maximum level while the on/off states were controlled by the thermostat. This configuration should provide a closer cycling simulation of a constant air volume cooling system in a large building. The 1500-watt heating capacity was provided by a 16-inch ceramic oscillating heater, which was also controlled by the same thermostat. The dual-setpoint thermostat was subsequently controlled by the ECoMIC controller board. Figure 36 shows the components of the HVAC system.



Figure 36 HVAC system at Rosemont site: (a) portable cooler (b) portable heater (c) Digital thermostat

Control board (Rosemont testbed)

ECoMIC controller board (ECB), shown in Figure 37, was the heart of Rosemont real-time demonstration. ECB implemented both *BaseCase* and *ECoMIC* algorithm to control electric lights, venetian blinds and the mini HVAC system in the test space. The switch between *BaseCase* and *ECoMIC* case was instructed by a supervisory program, and the switch always happened at midnight.




Figure 37 ECoMIC controller board (ECB).

The complete ECoMIC algorithm, mathematical models and optimization functions were reprogrammed in C language for the target controller from the original Matlab prototype. A modular software design approach was followed in the coding process so that the physical models and components can easily be

improved or replaced without affecting the rest of the real-time software. In addition to the ECoMIC algorithm, the controller was also programmed with *BaseCase*, where a set of operation schedules of lights, blinds and thermostat were defined during occupied and unoccupied periods.

The ECB hardware was realized leveraging a commercial Philips platform with simple modifications. The board houses a 32-bit microcontroller, 3 UARTs, EEPROM for non-volatile parameters and runs on 24V DC power supply. The custom modifications included the addition of 0-10V interfaces for light dimming control. The UART ports were used for interfacing with DyNet (photosensors), Somfy network (Blinds) and thermostat (HVAC). The controller had direct interface (RS485 and Analog I/O) to the indoor sensors. The exterior sensors were sampled by the ECB wirelessly using a dedicated sub-GHz radio.

Metering (Rosemont testbed)

A power meter was installed to measure the power consumption across three different circuits in the testbed. Individual power consumptions of lighting, HVAC and plug loads were read directly by a laptop and recorded in the central database. Plug loads include laptop, ECoMIC controller board power supply, blind motors and gateways, etc. Figure 38 is a snapshot of the power meter in operation showing the power consumption from the three channels.



Figure 38 Three phase digital power meter with external current transformer

6.3 Software

The development of software tools was site- and implementation-specific to ensure tight integration of the overall testbed operation.

6.3.1 Rosemont site

ECB software

The ECoMIC controller's software architecture, shown in Figure 39, is an instantiation of the ECoMIC zone service architecture discussed in Section 3.5, and comprises 5 basic blocks: ***sensing***, ***ECoMIC algorithm***, ***user interface***, ***host system*** and ***control drivers***. The design also makes data logging/retrieval easy for testbedding purposes. The sensing subsystem consists of basic drivers as well as a middle layer for receiving and interpreting light/temperature sensors as well as radio messages from exterior sensor board. The ***LightControl*** driver converts lighting commands from the algorithm (0-100%) into discrete dimming steps (0-255), which are consequently translated into 0-10V dimming signals by the on-board

circuitry. **BlindControl** driver performs translation of height/angle commands, broadcasts commands, reads blind states, etc. **HVACControl** driver is a RS485 driver that communicates heating/cooling setpoints and modes with the thermostat.




Figure 39 Software architecture of real-time ECoMIC controller implementation.

The **user interface** is provided in command line, which is used by the visual basic application tool (described later in this section) on the demo laptop. The interface lays out a comprehensive set of commands and formats for different actions. Notice that the user interface in this particular testbed implementation is meant to assist the researchers in the operation of the testbed rather than providing a means for occupants to interact with the control system. Other user interfaces for interacting with the control system were planned and the capabilities have been built in, but no app was developed since it was out of the scope of this implementation.

Sensor board software

A small sensor board, housed in the exterior sensor box enclosure, with a microcontroller (same family as the ECB) is used to sample, filter, process and transmit the sensor readings from the exterior sensor box over wireless radio frequency. The microcontroller on the sensor board polls the sensors constantly and keeps an image of the sensor readings in a table. When there is a radio message from the ECB requesting exterior sensor data, the most recent image is packed into a text stream and transmitted through radio frequency.

VBA Tool for demo laptop/PC

A simple visual basic application (VBA) was developed for managing control, monitoring and commissioning of the ECoMIC real-time controller. The tool has a *monitoring page*, shown in Figure 40 (a), for monitoring selected system parameters in real time, such as light level, temperature, occupancy, etc. It also reads the power meter, computes and displays energy consumption of the HVAC system, lighting system and plug loads. The *control page* of the VBA tool, shown in Figure 40 (b), provides the capability to command the ECB, which allows manual operations like setting system modes (*BaseCase/ECoMIC*), dim lights, move blinds and change temperature setpoints, etc. The commissioning page, shown in Figure 40 (c), supports commissioning of a new ECoMIC controller or changing existing

configurations. In order to reduce time and complexity of commissioning, each commissioning parameter is provided with a default value as a guidance.




Figure 40 EcoMIC VBA tool example: (a) monitoring page (b) control page (c) commissioning page.

Database and Visualization

A central database was established at the Philips Research Center using PostgreSQL, an open source database management system. The relational database system was designed to host several tables that contained real-time sensory and system information from different testbeds. Also, a web interface was developed to provides graphical display of current and historical information of the selected parameters. As an example, Figure 41 shows the visualization of the interior illuminance data from the Rosemont testbed on a selected period. The chart handle also provides a horizontal bar to zoom in on the data.



Figure 41 Example: Real-time visualization of testbed data.

6.3.2 LBNL FLEXLAB site

The LBNL testbed was established leveraging the FLEXLAB Cal-Bay acquisition and control infrastructure, which greatly reduced the need of accessory software tools. Capabilities, such as data base visualization, report, etc., provided by the Cal-Bay infrastructure are considered standard practice and are omitted herein.

Graphical user interface

Since the test space was an office constantly utilized by an occupant, a graphical user interface was developed for the user to interact with the control system. This GUI was written in Matlab and compiled as a standalone executable for running on a compact laptop inside the testbed. The landing page, Figure 42 (a) allows the user to choose to set their lighting preference or to bypass the controller and manually manipulate electric lights and blinds. The *Set Lighting Preference* button in Figure 42 (a) will lead to the preference setting page, Figure 42 (b), where task light level and glare tolerance level may be specified. When the *manual control* button in Figure 42 (a) is pressed, the controller bypass page, Figure 42 (c), will show up to allow separate manual control of electric light level, blind deployment and blind tilt.




Figure 42 User interface at the LBNL testbed. (a) landing screen; (b) preference setting; (c) manual control override.

Notice that a human subject test was not in the scope of this particular testbed implementation, and no user input was solicited or studied in any way through the GUI during the experiment period.

6.4 Configuration

6.4.1 Rosemont

The Rosemont testbed was developed with a commercial product in mind following the ECoMIC zone service architecture specified in Section 3.5. The system integrated the three components: 1) the ECoMIC controller board (ECB) implementing the ECoMIC algorithm, 2) the exterior sensor board collecting and transmitting the exterior sensor measurements, and 3) interior sensors and actuators.

The ECB is the center part of the testbed, where both *BaseCase* and *ECoMIC* case are implemented. Figure 43 shows the testbed configuration. The ECB interfaces with interior sensors, motorized venetian blinds, thermostat using serial communication (DyNet on RS485). An onboard 0-10V dimming circuit controls the lighting in the space. The ECB receives exterior sensor information from the exterior sensors over wireless digital radio link to the external sensor board.

A dedicated demo laptop controls and collects data from the ECB. The data include interior and exterior sensor information, algorithm outputs, actuator states, system modes and debugging information. In addition, the demo laptop reads the energy information directly from the power meter. The laptop then

sends all the data over TCP/IP to the database server remotely hosted at the Philips Research Center. An on-board Wi-Fi chip enables control over smartphone user interface for future enhancement.




Figure 43 Configuration of Rosemont testbed.

BaseCase mode

In the *BaseCase* mode, the lights were left ON at full output level during work hours between 7 am and 6 pm. The heating setpoint on the thermostat was fixed at 68°F. The thermostat cooling setpoint was fixed at 72°F for the first 21 days to emulate a business-as-usual setting and raised to 76°F for the rest 12 days to simulate a more aggressive setpoint for energy efficiency. The blinds were fully deployed with slats rotate to 30°, which blocks direct sunlight for the majority of time during this particular test period.

ECoMIC mode

In the *ECoMIC* mode, light level, blind slat angle and blind height were optimized by the ECoMIC algorithm and were allowed to move freely every 1 minute, 5 minutes and 15 minutes, respectively. These time delays were designed to avoid frequent movement of venetian blinds, which could be annoying in reality when occupants present. The cooling and heating setpoints were also updated by the ECoMIC algorithm every 15 minutes.

6.4.2 LBNL FLEXLAB

The testbed implementation at LBNL leveraged on the testing infrastructure in the FLEXLAB, which includes the Cal-Bay sensing and control system, and IntraStage database and reporting system. The Cal-

Bay system administers data acquisition and actuation for analog devices as well as accepting data from extraneous sources. The IntraStage system aggregates data recorded in the Cal-Bay system, stores them into a database, and provides a web-based user interface for querying and report generating purposes.




Figure 44 LBNL FLEXLAB implementation setup.

Some of the devices used in this implementation were digital devices, including blind drivers and light sensors, and thus were not directly supported by the Cal-Bay infrastructure. Therefore, an independent signal chain had to be created for data acquisition and actuation on these devices. Figure 44 shows the complete configuration of the LBNL testbed, and the following two software components were developed for bridging the ECoMIC controller, the Cal-Bay infrastructure and the independent signal chain of digital devices.

- The Cal-Bay interface: This interface is a gateway to the FLEXLAB Cal-Bay system as well as the devices on the digital channel. It serves three purposes for the control algorithm and GUI implemented in Matlab to 1) acquire readings from the analog sensors on the Cal-Bay system, 2) control the electric light level through the Cal-Bay system, 3) communicate with devices on the digital channel through the LabView interface.
- The LabView interface: This interface translates messages to/from the Cal-Bay interface into serial signals for communicating with the devices on the digital channel using the DyNet protocol, including venetian blind drivers and photosensors.

The ECoMIC controller in this testbed implementation was realized in Matlab. This software approach allows rapid prototyping and revision of the control algorithm with minimum interruption to the running test, which is one of the FLEXLAB's main features.

6.4.3 Commissioning

Most of the commissioning effort was dedicated to the assignment of the parameters required in the control algorithm. Assigning values to the parameters involves basic measurement, observations, hardware specification and sometimes best estimate. The parameters fall into the following categories:

- Site related parameters;
- Façade related parameters;
- Zone related parameters;
- Reference point or sensing point parameters;
- Fenestration/blind/glazing related parameters;
- Electric lighting system parameters;
- Room thermal characteristics parameters;
- Hardware limits.

The glare sensor was given special attention during the commissioning process, especially at the LBNL FLEXLAB testbed, and is detailed herein. The sensor had to be mounted in an unobstructed and reasonably centered location on the back wall. The ratio of daylight glare at the occupant's location to the glare sensor reading was analytically calculated by the algorithm using the glare model along with the above commissioning parameters. As an alternative, if the location where glare needs to be controlled is known, the ratio can be measured and directly inputted into the controller. To properly interpret the daylight glare from the glare sensor, the controller also needs the relationship between the electric light levels and the resultant illuminance value on the glare sensor. This is most easily done at night, but can be done during the day if the daylight values are subtracted out from the sensor readings. A general commissioning procedure for glare control including further details on the default values and logic behind the commissioning is provided in Appendix C.

In addition, there are also time delays in the controller that may be modified. These are the time delays to raise or lower the blinds after the target level has been reached; the time delays to increase or decrease blind tilt after reaching the target level, and the time delays to increase or decrease the electric lighting.

6.5 Performance

The following analyses are based on the data collected from both test sites during the period listed in Table 10.

Table 10 Data collection period.

Test site	Data collection period
Rosemont	4/26/2013 – 7/7/2013
LBNL FLEXLAB	6/7/2013 – 7/7/2013

The dataset from Rosemont testbed includes 33 *BaseCase* days and 39 *ECoMIC* days. During the entire test period, there was only one day where data collection failed due to technical problems, and the data from that day was excluded from the dataset. The entire dataset was utilized in the analyses unless specifically noted.

LBNL testbed was purposed for evaluating glare control performance, and therefore, no baseline case was administered. Data points during the occurrences of technical problems and interruptions were excluded from the analyses. Since both task light level and glare setpoints were allowed to change

during the test period, the analyses were performed on subsets of the dataset from comparable days. The data selection criteria will be specified in each analysis.

6.5.1 Energy performance

Energy performance analysis was performed using the dataset from the Rosemont testbed where energy usage was independently metered. Recall that the testbed was switched between the *BaseCase* and *ECoMIC* case every three days to mitigate potential bias from weather factors. In addition, an attempt was also made to normalize the HVAC energy usage against outdoor temperature; however, as Figure 45 suggests, no strong correlation between HVAC energy and outdoor temperature was found through this exercise. This was likely due to the relatively short test period as well as the thermal connection between this small room and the adjacent rooms. Therefore, no further weather normalization was performed on the dataset.




Figure 45 HVAC energy vs. outdoor temperature. Histogram: daily average of HVAC energy consumption; dotted lines: daily average of outdoor temperature.

The resulting energy savings are shown in Figure 46 (a). The *ECoMIC* case achieved about 68% overall savings in daily average energy consumption over the *BaseCase*. The HVAC daily energy consumption, primarily cooling energy, was reduced by 79%. Lighting energy consumption was also reduced by 40%. The plug load in the *ECoMIC* case was 1% higher than that in the *BaseCase*. The average daily energy usages in each category for both the *ECoMIC* and *BaseCase* are shown in the histogram in Figure 46 (b).




Figure 46 Energy performance. (a) Overall energy savings in each category; (b) average daily energy consumptions.

HVAC Energy

The HVAC savings are primarily due to dynamic thermostat setpoint controls. Meanwhile, cooling demand dominated the HVAC energy consumption in both cases while heating was seldom activated during the test period from late spring to early summer.

Figure 47 (a) compares thermostat cooling setpoints for the *BaseCase* and *ECoMIC* case. Recall that the cooling setpoint for the *BaseCase* was fixed at 72°F for the first 21 days of the test period to emulate a typical practice in commercial buildings. The setpoint was later increased to 76°F for the remaining 12 days to represent the recommendation of a more aggressive efficiency strategy. Therefore, the distribution of the *BaseCase* setpoint concentrated at both 72°F and 76°F. However, no significant change in the daily HVAC energy consumption was observed before (10.17 kWh) and after (10.23 kWh) this change in the *BaseCase* (also see Figure 45).




Figure 47 Distribution of (a) cooling setpoint and (b) indoor temperature measurement during work hours.

The cooling setpoint in the *ECoMIC* case, on the other hand, were shifted towards the warmer end of the spectrum. The air conditioner was consequently cycled on much less frequently, which was the main contributing factor to the high energy savings. It worth noting that even though the cooling setpoints in the *ECoMIC* case were pushed much to right of the *BaseCase* setpoints, it did not mean the actual temperature was regulated right at the setpoints; rather, the air temperature was allowed to float within a much wider range without having to cycle on the air conditioning system. The resulting distribution of air temperature measurement is shown in Figure 47 (b).

The configuration of the HVAC system might also contribute to the high energy savings. As described in Section 6.2.2, the independent mini HVAC system exclusively served the testbed and was not affected by the cooling/heating demand from adjacent zones in the same building. It might also have created a different pattern of air circulation as well as stratification. In this sense, the energy performance of such system is not representative to a typical HVAC system in commercial buildings.

Lighting control optimization may also be a contributor to the HVAC energy savings. As will be described next, lighting in the *ECoMIC* case consumed only about 60% energy compared to the *BaseCase*. The full-on power consumption of the lighting system, as in the *BaseCase*, is 360W, most of which eventually became heat gain. The 40% savings in lighting energy also implies a lighting heat gain reduction of about 150W in the *ECoMIC* case. In the meantime, the blind in the *ECoMIC* case would be closed more to block excessive daylight for the specified task light level. When daylight is at the strongest, the corresponding solar heat gain also peaks; therefore, closing the blind also helped blocking solar heat gain to some degree, which in turn reduced the HVAC energy usage.

Lighting energy

The majority of savings in lighting energy came from better utilization and management of daylight. The savings contribution from occupancy sensing is less than 5% due to the simulated high-density (94%) occupancy pattern as described in section 6.2.1 as well as Figure 34. This implies that lighting energy savings potential will be even higher in most cases as a typical office worker is expected to be absent from the office more often than 6% of the work hours for meetings, lunch, breaks, etc. The median of dimming levels issued by the algorithm throughout the measurement period was 48%. Another interesting finding is that the exterior daylight condition in the *ECoMIC* case was on average darker than that in the *BaseCase* as can be observed from Figure 48. This is also an indication that higher savings is possible if the exterior daylight conditions were made comparable in both *ECoMIC* and *BaseCase*.




Figure 48 Lighting energy vs. exterior light level. Histogram: daily average of lighting energy consumption; dotted lines: daily average of exterior vertical illuminance on the façade.

Another consideration is the designed light level in the testbed. In the *BaseCase*, the electric light was *full-on* during business hours, illuminating the workplane at the designed light level, which is higher than the 500 lux setpoint in the *ECoMIC* case. In other words, a small part of the lighting energy savings came from tuning the light down to the level needed for the task. Since the lighting preference, hence the task light level setpoint, differs from individual to individual, this part of savings will usually be an uncertainty. On the other hand, one could also argue that the lighting energy savings is not an apple-to-apple comparison since the *BaseCase* has a higher light level to begin with. In this case, the actual energy savings will slightly diminish when the maximum light level in both cases are normalized.

Also notice that the test was run during late spring and early summer, so some amount of daylight was always available during the entire work period (6AM-7PM) for offsetting the use of electric light in the *ECoMIC* case. This will not be the case for other months. In other words, lighting energy consumption in the *ECoMIC* case should be higher during months when there is no daylight in early morning and late evening; hence, the energy savings will be lightly lower if projected to the entire year.

Plug load energy

The plug load includes energy used to power the *ECoMIC* controller board, blind motors, blind drivers, a demo laptop and an IP camera. Compared to zero blind movements in the *BaseCase*, it does make sense that the *ECoMIC* case had slightly higher plug load. However, the 1% difference in plug load is well within the noise level and hence is deemed insignificant.

6.5.2 Visual performance and comfort

6.5.2.1 Task illuminance

The task illuminance setpoint was set at 500 lux for the *ECoMIC* mode over the entire test period at the Rosemont testbed, and the actual workplane illuminance was allowed to float within a small range to avoid any unnecessary fluctuation. 500 lux illuminance is the typical light level for most office tasks recommended by the Illuminating Engineering Society (IES) and other authorities. The distribution of the resulting workplane illuminance is shown in Figure 49. More than 85% of the time the workplane illuminance was maintained between 400 and 500 lux in the *ECoMIC* case. That the distribution skewing towards the left of the 500 lux setpoint within the allowable floating range demonstrates that the *ECoMIC* algorithm attempted to utilize less electric light and block unnecessary daylight, and hence heat gain, for higher energy savings. Notice that the instances where task illuminance went below 400 lux in the *ECoMIC* case correspond to unoccupied durations where electric lights were turned off.




Figure 49 Distribution of task illuminance at the Rosemont testbed.

The task illuminance variations in the *BaseCase* were only governed by the intensity of daylight since there was no blind movement regulating it. The 700 lux bar includes all data points with task illuminance above 700 lux, and the more-than-20% occurrence in the *BaseCase* implies a higher possibility of glare.

Another perspective of visual performance delivered by the *ECoMIC* algorithm with respect to *BaseCase* is presented in Figure 50. Each small rectangle in the plots represents the average of task illuminance during a particular hour (y-axis) on a particular day (x-axis). The average task illuminance over the entire test period was uniform between 10 AM and 6PM in the *ECoMIC* case. During the same hours in *BaseCase* days, task illuminance was not maintained uniformly, and, in many cases, implied high glare conditions (see Figure 50 (a)). For east-facing windows, direct sun is expected to occur during morning hours, and hence higher illuminance or glare. The algorithm was able to predict and minimize glare as observed for most of the days in Figure 50 (b). The few instances where the averages of hourly illuminance level in the *ECoMIC* case were slightly higher than expected can be attributed to the slightly slow response of the algorithm in curtailing vertical illuminance. As a result, the hourly averages still remain high. This is an area of further investigation and improvement.




Figure 50 Interior work plane illuminance hourly averages (a) BaseCase (b) ECoMIC case.

A subset of the dataset from the LBNL testbed where task illuminance setpoint was 400 lux (also with a reasonable deadband) was selected to examine the resulting workplane illuminance. Notice that the task illuminance setpoint at this testbed effectively served as the lower bound of the workplane illuminance while the upper bound was implicitly governed by the glare threshold setpoint. In other words, instead of regulating the task illuminance within a narrow band of the setpoint like the Rosemont testbed, the workplane illuminance was allowed to drift much higher so long as glare threshold was not triggered. This strategy exploited the wide adaptive range of human eyes under daylight to preserve more outside view through the window.

The distribution of workplane illuminance at the LBNL testbed is shown in Figure 51. The task illuminance was regulated above 300 lux almost all the time. Notice that many occurrences fell into the bin of 300 lux due to the allowed deadband surrounding the setpoint and were actually close to 400 lux. There was no occurrence above 800 lux, beyond which daylight glare was likely to occur causing glare prevention control to come into play. This aspect will be discussed next.




Figure 51 Distribution of task illuminance at the LBNL testbed.

6.5.2.2 Glare prevention

The same set of data in the task illuminance analysis from the LBNL testbed was used to assess the ability of preventing glare. The glare threshold setpoint was 4000 lux (recall that the task light level setpoint was 400 lux, and task illuminance was allowed to float beyond 400 lux as long as no glare was detected). Notice that this glare setpoint is higher than the 2000 lux default described in Section 4.3.4 by a factor of two due to the occupant's typical seating orientation in this office space. The 2000 lux default setpoint assumes that a person sits directly facing the window. In this particular office configuration, the seating orientation was 90° from the window, and, thus, 4000 lux was determined as a reasonable setpoint. In any case, the controller is able to regulate glare at any user-specified setpoint, and the actual setpoint value should not be a factor of the glare prevention performance.

The distribution of the glare condition throughout the selected period is shown in Figure 52. About 95% of the time the controller was able to keep glare below the 4000 lux threshold. And the fact that more than 25% of the time glare was capped right at 4000 lux implies that the vertical illuminance was indeed actively controlled for glare rather than the daylight level being low to start with.

The time series plot in Figure 53 illustrates several instances where blinds were closed to avoid glare as marked by the blue shaded areas. In addition to closing the blind slats to avoid glare, the green shaded area demonstrates the reverse action where the blind slats were opened more when the controller predicted that doing so would not cause glare. Modulating blind slats to prevent glare from direct sun is shown in the purple shaded area. Recall that the testbed had a southwest-facing window, and thus direct sun glare occurred at late afternoon. The same plot also shows an instance (red shaded area) that belongs to the remaining 5% of the time in Figure 52 where the controller was slower in response to glare. This is certainly an important area for further investigation and improvement. Also notice that

there are many more electric light and blind actions in Figure 53, which could be the combined results of the optimization over multiple considered factors in the ECoMIC algorithm, and attributing each of them to a single factor is not possible.




Figure 52 Distribution of glare at the LBNL testbed.




Figure 53 Illustration of active glare controls. Blue shaded areas: blind tilted closer to prevent daylight glare; purple shaded area: blind closed and opened to prevent glare from direct sun; green shaded area: blind tilted more open when no glare condition was predicted; red shaded area: controller was slower in response to glare.

6.5.3 Thermal comfort performance

Since the thermostat setpoints at the Rosemont testbed were dynamically determined and adjusted in the ECoMIC case, a thermal comfort analysis was performed to evaluate its value. Fanger's thermal comfort model, namely, predicted mean vote (PMV) was employed to compare the thermal comfort condition between the *ECoMIC* and *BaseCase*. Recall that the ECoMIC algorithm was set to maintain PMV within ± 0.7 , corresponding to a greater-than-85% satisfaction.

Figure 54 shows the distribution of daily averaged PMV values. As the histogram reveals, close to 90% of the time the PMV value was maintained within the targeted ± 0.7 in the *ECoMIC* case. The PMV in the *BaseCase*, on the other hand, was maintained in the same range for only about 70% of the time. It can also be seen in Figure 54 that the distribution of PMV for the *ECoMIC* case was shifted towards to the warmer spectrum (positive values) of the PMV scale compared to the *BaseCase*. This illustrates that the *ECoMIC* algorithm attempted to generate more energy savings by allowing higher indoor air temperature during cooling seasons as long as an acceptable thermal comfort condition was retained. The reverse behavior is expected for heating seasons where the *ECoMIC* algorithm would shift the distribution towards the cooler spectrum (negative values) of the PMV scale while still maintaining an acceptable comfort level.




Figure 54 Distribution of the daily-averaged PMV.

The control and analysis of thermal comfort, including the ± 0.7 PMV range, are largely based on the empirical PMV scale; recall that there was no actual occupant in the Rosemont testbed. This testbed implementation demonstrated the capability of the *ECoMIC* algorithm to regulate thermal comfort within a specified range in the most energy-efficient way using PMV ± 0.7 as a nominal range. As thermal sensation differs from person to person, acceptable comfort range will need to be adjusted based on user input in reality, both in terms of the width and the end points of the range. These adjustments could potentially result in a diminished energy savings depending on each individual's thermal comfort preference. Other comfort factors that could not be explicitly captured by the controller, including activity level, clothing insulation, humidity and air flow, will also dictate the acceptable indoor temperature, and hence the possible range for thermostat controls. Therefore, the corresponding energy savings will vary in practice.

Another illustration of thermal comfort delivered by the control algorithm is shown in Figure 55, where the small rectangles represent the hourly averages of the PMV values for every *BaseCase* and *ECoMIC* day. As described before, thermostat setpoints in the *BaseCase* were fixed throughout the test period, generating non-uniform thermal sensations across multiple days and sometimes within one day. The PMV variation in the *BaseCase* fluctuated between slightly warm to very cool (Figure 55 (a)). On the other hand, the *ECoMIC* algorithm strived to maintain a narrow PMV range around the 'neutral'

sensation (PMV=0). There were few instances in early morning where PMV was on the ‘cool’ side. This was likely due to the fact that the ECoMIC algorithm only becomes active at 7am when the testbed was first switched to ECoMIC case every three days. Thus, there was a delay in bringing comfort level to within the preferred range.




Figure 55 Predicted mean vote (PMV) hourly averages (a) *BaseCase* (b) *ECoMIC* case.

6.6 Generalizability and limitations

The results from the testbed implementations in general show promising performance in both energy savings and occupant comfort. Even though they were not lab studies in a strictly controlled environment, these single-room, small-scale tests may still not represent a realistic field study. Recognizing this, limitations of generalizing the results and recommendations are discussed in this section.

Integrated controls

The ECoMIC algorithm models and minimizes lighting-related cooling loads. We acknowledge that while this zone-based control strategy will create minimized energy demand in the controlled space, it may not directly translate into a proportional reduction in building-wide HVAC energy consumption. For instance, on a HVAC system with terminal reheat, reducing the cooling demand in the perimeter zones may consequently activate the reheat coils, hence additional heating energy, if the cooling demand in the core zones remain the same or if the supply air temperature is not reset accordingly. As another example of zone control with variable air volume (VAV) boxes, ECoMIC thermostat setpoint optimization is directly proportional to the VAV box damper position. However, there is no guarantee that this volume reduction will correspond to the building-wide HVAC energy consumption unless all zones react similarly at a given time.

It will be very difficult, if possible at all, to account for this particular building-wide interdependency in the zone-level ECoMIC controller without any direct link to the central HVAC system. Therefore, certain communication with the HVAC system is recommended to provide supervisory level inputs to the ECoMIC controller to guarantee the correlation between zonal cooling/heating demand minimization and overall reduction in HVAC energy consumption. This link should easily be accommodated as one of the facility level services in the developed system architecture described in Section 3.

HVAC energy savings

The testbed implementation shows an energy saving that is much higher than expected. This positive result certainly demonstrates the savings potential of the developed control technologies. However, cautions should be used when attempting to generalize the results due to the following reasons: the particular configuration of the mini-scale nonstandard HVAC system, the test period and duration, and user's actual thermal comfort preference.

As mentioned in Section 6.5.1, the HVAC system exclusively serves the testbed and is isolated from the adjacent zones. Therefore, it may not represent a typical configuration of a centralized HVAC system in commercial buildings. Also, due to its small size and capacity, the cooling and heating mechanisms employed in the system are also significantly different from a HVAC system serving the entire building. These factors could contribute to overestimate or underestimate of the projected HVAC system savings in reality.

This testbed implementation was only run for a little more than two months during late spring and early summer, which is considered to be a relative short duration for energy monitoring. Cooling demands dominated the HVAC energy consumption during this period, so the energy performance in heating seasons is hard to project based on the collected data. There may also be a difference in energy savings if the system was exposed to the hottest months of the year instead of the early cooling season. Moreover, the period may be too short to allow meaningful weather normalization for a more pertinent comparison between the base case and the controlled case. In short, it is highly recommended to extend the test period to at least a year before drawing conclusions on the annual energy performance of the developed control technologies.

The majority of HVAC energy savings came from dynamic control of thermostat setpoints, and the thermal comfort reference point in this testbed implementation was based on a well-established empirical model. However, as thermal sensation is subjective and varies from person to person, it is important in practice to allow the user to adjust the comfort reference point. Each individual's ideal thermal comfort condition will have a significant implication on HVAC energy usage in practice. For example, if one prefers a cooler environment, *i.e.* lower cooling setpoint, the resulting cooling energy savings will diminish. Therefore, when projecting potential energy savings, the uncertainties introduced by the diversity of occupant preference need to be carefully considered. A human subject study will be very valuable to provide insight into the potential impact on this aspect.

Lighting energy savings

The lighting energy savings from the testbed implementation was mainly contributed by better daylight utilization and management. It is important to recognize that the test was run in late spring and early summer where daylight was always available during the entire work period for replacing at least some amount of electric light to illuminate the workplane. The electric lighting usage is expected to increase during other months with shorter daytime period. It is recommended to extend the test duration to a year or at least from solstice to solstice in order to establish a better understanding on the annual savings potential.

The orientation and area of the windows may have also played a non-negligible role in the resulting lighting energy savings. In particular, the energy savings in this case is from east-facing windows. The savings could be even higher for south-facing windows while significantly lower for north-facing windows. Furthermore, the building where the testbed resides employs a curtain wall design. In other words, the windows at our testbed are relatively large in area and are capable of admitting more daylight than those in buildings with a conventional fenestration design. These factors also need to be considered when generalizing the lighting energy savings.

It is also worth noting that the reported savings were achieved for an office that was occupied most of the time (with the occupancy model with 94% density). Occupancy sensors were known to generate about 30% electric lighting energy savings in general. Hence, same east-facing rooms with real occupants may generate more savings in lighting energy than reported from our testbed.

Visual performance

The ability to provide proper light level on the workplane relies on pertinent measurement of task illuminance. Like most daylight harvesting systems on the market, the control system estimates task illuminance from a ceiling-mounted photosensor. And therefore, the quality of the correlation between the photosensor measurement on the ceiling and the actual task illuminance directly affects the system performance since the system might be regulating the light at an ill-calibrated workplane light level. Although improving the estimation of task illuminance from a ceiling photosensor is not in the scope of this project, it is important to keep this in mind as it will inevitably be a major factor impacting occupants' satisfaction.

Glare detection

As the measured data reveal, the control system have demonstrated a reasonable performance in delivering a glare-free work environment. The development of the glare detection and prevention techniques has been predominately based on theoretical derivations using large sets of lighting simulations. Whether an occupant will really feel comfortable and satisfied with the glare control cannot be rigorously examined from this testbed implementation without any feedback and comment from actual users. Therefore, a human subject study is highly recommended as a future research direction to verify the effectiveness of the glare control technique.

7 Success assessment

7.1 Energy savings

The energy performance goals set out for this project are

- 60% lighting energy savings relative to a “no-controls” case in daylit areas;
- 20-35% savings of whole building energy consumption.

The preliminary analysis performed in phase 1 of this project indicated a 66% annual lighting energy savings in daylit zones with integrated lighting and daylighting controls compared to a “no-controls” base case. Along with a separate analysis on cooling and heating energy savings due to dynamic thermostat setpoint control, the building-wide annual energy savings is estimated to be 15%.

In phase 2 of the project, the ECoMIC algorithm was implemented and evaluated in an energy modeling and simulation environment. During a typical summer week in a mixed-humid climate zone, the average lighting energy and cooling load reductions were 27% and 42%, respectively, in the south perimeter zones compared to a non-trivial base case with lighting and daylighting controls (daylight dimming on electric lights and venetian blinds are always rotated to block direct sun). The numbers would increase to 87% and 51% for reductions in lighting energy and cooling load, respectively, if compared to a trivial “no-controls” base case operated on scheduled lighting without shading systems.

The ECoMIC algorithm was evaluated in a testbed implementation during the third phase of this project in a single zone daylit office space. A test from late spring to early summer for a little more than two months demonstrated a 68% overall energy savings compared to a representative base case (electric lights full on during business hours and blinds always deployed and slightly closed at all time). The reduction in electric lighting energy was 40%, and the majority of the savings resulted from better daylight management with small contributions from occupancy sensing. The energy savings on HVAC energy was 79%.

The 40% lighting energy savings from the phase 3 testbed implementation seems to fall between the two simulated lighting energy savings, namely 27% compared to an advanced base case and 87% compared to a “no-controls” case. Considering that the sophistication of the base case in the testbed implementation also sits in between the two simulated base cases, we believe that lighting energy savings in both the simulations and the testbed implementation will agree with each other if compared to the same baseline. The combined results from the simulation and testbed implementation exercises clearly indicate the feasibility of 60% lighting energy savings relative to a “no-controls” case in daylit areas.

The HVAC energy savings is higher in the testbed implementation than in simulations possibly due to the different HVAC system configuration. It may not be realistic to attempt to extrapolate the overall energy performance from the single-room testbed implementation to a whole building. However, the energy savings from the testbed demonstration as well as the simulations are strong evidences that the goal of 20% savings on whole building energy consumption is achievable with the ECoMIC framework.

7.2 Occupant comfort

User learning

The developed preference learning algorithm has demonstrated the capability to reliably model user preference. The learning accuracy of the preference models is dependent on many factors, such as the length of learning period, the resolution of the recorded data points, etc. Nonetheless, our simulations have repeatedly confirmed that with 5 days of learning and 1 hour learning record resolution, the algorithm is able to achieve more than 96% relative accuracy.

Visual performance and comfort

Task illuminance is the primary metric for assessing visual performance resulting from the control framework. The testbed implementation at both sites have shown that the system is able to maintain the task illuminance within an acceptable range of the light level setpoint on the workplane, especially in holding the minimum required light level.

Glare is the main metric for evaluating visual comfort. The system managed to detect glare and respond immediately by actuating the blind in the real-time testbed implementation, and the controlled space was maintained glare-free about 95% of the time. Furthermore, the system was also able to make correct prediction when the glare condition no longer existed and open the blind to allow better view and admit daylight in a timely manner. This marks the success in glare control. In the meantime, the remaining 5% of the time where glare illuminance was allowed to go beyond the threshold by as much as 10% also indicates rooms for further improvement.

Thermal comfort

We adopted Fanger's model to evaluate thermal comfort and satisfaction rating in terms of Predicted Mean Vote (PMV) and Predicted Percentage Dissatisfied (PPD), respectively. The simulations in Phase 2 of this project confirmed the system's ability to maintain comfort at a level well above 80% predicted satisfaction (less than 20% PPD). In the testbed implementation where the reality of HVAC system response came into play, the resulting thermal comfort condition was still maintained at targeted satisfaction level close to 90% of the time. Although thermal comfort preference may differ from person to person, the system should have no problem maintaining any occupant-specified thermal condition as long as a user interface is provided for this adjustment.

7.3 System scalability

The ECoMIC architecture allows for systems that scale from single offices to large buildings via a backbone network. The backbone network is not restricted to a single technology, instead it allows the system implementer to choose from any existing building control networking topology or to create a new network if desired. This flexibility is designed to address retrofit installations where an existing building control system is currently in place, as well as new construction, where creation of a new backbone network or leveraging an existing data network is possible. The freedom of specifying the network technology makes the system scale to the needs of the installation. The ECoMIC architecture

distributes the execution of the system across multiple components which may be added with minimal impact on the backbone network as all sensor and actuator interfaces are local components.

The ECoMIC architecture specifies that sensors and actuators connect directly to a Zone Controller component which executes the ECoMIC algorithm. Local message are not transmitted across the backbone network. This design decreases the latency an end-user will experience when interacting with the system. The actual system latency will be defined by the system implementer's selection of sensors, actuators, throughput of the communication link and hardware resources on which the ECoMIC algorithm is executed and must be chosen with consideration of acceptable latency.

The two testbed implementations were successfully instantiated following the ECoMIC architecture. While utilizing the same sensing and actuating hardware, they were realized on completely different controller platforms, backbone networks, etc. This demonstrates the flexibility of the proposed architecture. During the entire test period, the systems have never experienced any interruption caused by network downtime or latency.

7.4 Economic benefits

Economic benefits of the ECoMIC framework may be analyzed using a variety of techniques. We focus the analysis on the return on investment (ROI) as this is typically the most accepted metric in the decision-making process of a new construction or a retrofit project.

The analysis was performed on a DOE reference energy model for a post-1980 medium office building [12] with an 1.5:1 aspect ratio. Specifically, we assumed that the perimeter zones of the building were retrofitted with the control solutions, including the new dimmable luminaires and motorized venetian blinds. Instead of the original five-zone configuration on each floor in the model (one big zone in each perimeter and a core zone), the south and north zones on each floor were further divided into 6 zones, and the east and west zones were divided into 4 zones. The energy and demand charges were calculated through EnergyPlus simulations using the GS-2 rate for medium demand customers from the Virginia Electric and Power Company. The cost of the control solution was estimated based on the costs of the components used in the testbed implementation in Section 6 with an assumed discount rate for the mass production and purchasing. The labor hour for installing each component was also derived from our experience implementing the testbeds, and a reasonable labor rate was assumed. A detailed estimation can be found in Appendix D.

The DOE reference building model as is (complied with ASHRAE Standard 90.1-1989) is considered as the base case. There are two layers of savings on energy and demand. The first layer of savings is from the retrofit per se., which brings down the lighting power density and the corresponding HVAC energy usage due to the reduction in lighting heat gain. This savings is directly captured by the EnergyPlus simulations. The second layer of savings is introduced by the developed control solutions. The results from simulations and testbed implementation in Sections 5 and 6 serve as the basis of our estimations on energy intensity reductions. We conservatively assume that the demand savings are proportional to the energy savings, but in reality, peak demand reduction often contribute to much larger savings. The detailed estimation can be found in Appendix D.

Combining the energy savings from the EnergyPlus simulations and the estimated system costs yielded a payback time of 6.6 years. The payback period would of course be much shorter (less than 2 years) for new constructions if dimmable lights and motorized venetian blinds have already been specified as part of the original design. Also, the payback time in 2020 will be more advantageous as energy prices are expected to rise and the costs on the control systems will drop further due to higher demand.

Our approach is user centric. As a result, we anticipate increased user satisfaction, which will in turn lead to productivity and other benefits. It is rather difficult to quantify the productivity benefit as this entire area is largely understudied, and data on this is limited. However, some researchers have reported a productivity gain of 3-14% due to daylighting, view to outside and personal controls [13]. If we factor a rather small productivity improvement of just 0.25%, the ROI can be reduced to just a few months.

8 Conclusion

An integrated control framework has been developed and demonstrated during the course of this project. The zone-based control technology aims to minimize energy usage while improving occupants' visual and thermal comfort through the integrated control of electric lights, motorized venetian blinds and thermostats. Electric light level, venetian blind deployment, venetian blind tilt angle and thermostat setpoint are the four control variables. The mathematical models linking the control variables to the interrelated energy and comfort elements in the control framework were derived, simulated and fitted into the control algorithms.

Three main algorithms are included in the control framework: the lighting load balancing algorithm, the thermostat control algorithm and the preference learning algorithm. The lighting load balancing algorithm adopts a model-based closed-loop control strategy. The control problem is formulated into an optimization problem with minimizing lighting-related energy consumptions as the objective and delivering adequate task light for visual performance and preventing daylight glare for visual comfort as the constraints. The thermostat control algorithm is based on a well-established thermal comfort model and formulated as a root-finding problem to dynamically determine the optimal thermostat setpoint for both energy savings and improved thermal comfort. The preference learning algorithm employs a reinforcement machine learning technique to derive desired setpoints through users' interactions with the control system. The setpoints that can be learned using the algorithm include task light level setpoint, glare threshold setpoint and thermal comfort setpoint.

Accompanying the zone-based control framework, a system architecture has been developed to address scalability of the control technology. Three levels of services are defined in the architecture: external services, facility level services and zone level services. The control algorithms are included as part of the zone-level services along with the corresponding interfaces, profiles, sensors and actuators to realize the controller and perform physical sensing and actuation within a zone. On top of the zone level services are the facility level services that connect to the zones through a backbone network, provide supervisory level information and controls, and thus facilitate building-wide scalability. Beyond the facility level services, the external services provide communication capability to entities outside of the building. These services can foster capabilities such as demand response and remote monitoring.

Various aspects of the developed control technology has been evaluated and verified through both simulations and testbed implementations. A most comprehensive energy simulation for a period during summer time revealed a 27% savings in electric lighting load and 42% reduction on cooling load when compared to an advanced base case with daylight dimming and blinds always tilted to block direct sun. The testbed implementation resulted in 40% and 79% savings on lighting and HVAC energy consumptions, respectively, over a two month period in late spring and early summer compared to a base case operated on predefined time schedules. The testbed implementation also demonstrated a reasonable glare detection and prevention capability. Through the testbed exercises, areas for further improvement and enhancement were identified. The potential reasons of the high savings and the limitations of generalizing the results from the testbed implementation were investigated and discussed in details.

Judging from the criteria of energy savings, occupant comfort, system scalability and economic benefits, we confidently conclude that the project is a success. Based on the results from simulations and testbed implementations, 60% lighting energy savings can easily be achieved by the developed control technology in daylit areas relative to a “no-controls” case. A 20% reduction of whole building energy consumption is also attainable. In the aspect of occupant comfort, the testbed implementations demonstrated the ability to successfully maintain specified light level on the workplane while promptly controlling the blind tilt to mitigate daylight glare 90% of the time. The control system also managed to maintain the thermal environment at a comfortable level 90% of the time. The aspect of system scalability was guaranteed by the system architecture design, based on which the testbeds were instantiated. The systems have never experienced any interruption due to network downtime or latency over the entire test period. Analysis on the aspect of economic benefit showed a 6-year payback period for retrofitting projects, and a payback time of less than 2 years for new constructions if dimmable lights and motorized blinds already in place.

Reference

- [1] Standard 55-2010 -- Thermal Environmental Conditions for Human Occupancy, ASHRAE, 2010.
- [2] OpenADR Alliance, <http://www.openadr.org>.
- [3] N.A. Kotey, M.R. Collins, J.L. Wright and T. Jiang, "A Simplified Method for Calculating the Effective Solar Optical Properties of a Venetian Blind Layer for Building Energy Simulation," *J. Solar Energy Engineering*, vol. 131, no. 2, May 2009, pp 021002-1 – 021002-9.
- [4] D.S. Yahoda and J.L. Wright, "Methods for Calculating the Effective Longwave Radiative Properties of a Venetian Blind Layer," *ASHRAE Transactions*, vol. 110, no. 1, 2004, pp. 463-473.
- [5] J.L. Wright and N.A. Kotey, "Solar Absorption by Each Element in a Glazing/Shading Layer Array," *ASHRAE Transactions*, vol. 112, no. 2, 2006, pp. 3-12.
- [6] J.L. Wright, "Calculating Center-Class Performance Indices of Glazing Systems with Shading Devices," *ASHRAE Transactions*, vol. 114, no. 2, 2008, pp. 199-209.
- [7] J.L. Wright, "A Correlation to Quantify Convective Heat Transfer Between Vertical Window Glazings," *ASHRAE Transactions*, vol. 102, no. 1, 1996, pp. 940-946.
- [8] ISO 15099, Thermal Performance of Windows, Doors and Shading Devices – Detailed Calculations, 2003.
- [9] N.A. Kotey, C.S. Barnaby, J.L. Wright and M.R. Collins, "Solar Gain through Windows with Shading Devices: Simulation versus Measurement." *To appear in ASHRAE transactions*, vol. 115, no.2, 2009.
- [10] R.D. Clear, V. Inkarojrit and E.S. Lee, "Subject Response to Electrochromic Windows", *Energy and Buildings*, Vol. 38, No. 7, 2006, pp. 758-779.
- [11] P.O. Fanger, *Thermal Comfort: Analysis and Applications in Environmental Engineering*, McGraw-Hill, New York, 1973.
- [12] M. Deru, K. Field, D. Studer, K. Benne, B. Griffith, P. Torcellini, B. Liu, M. Halverson, D. Winiarski, M. Rosenberg, M. Yazdanian, J. Huang and D. Crawley, "U.S. Department of Energy Commercial Reference Building Models of the National building Stock", Technical Report, NREL/TP-5500-46861, 2011.
- [13] Boyce, P., "Reviews of technical reports on Daylight and Productivity", *Lighting Research Center*, RPI 2004.

Appendix A Glare characterization, detection and estimation

This is a summary of glare characterization and prediction produced at the end of Phase 2 of this project.

A.1 Theory

A.1.1 Prior work

In the Phase 1 of this project, it was hypothesized that under conditions where the sun does not directly enter the space, that one monitor sensor would be sufficient to capture the influence of outside conditions, and that illuminance would drop off in a relatively smooth fashion as a function of the distance to the window wall. No a priori assumptions were made regarding the functional form of the drop-off, but we were hoping that a relatively simple form, such as exponential, trigonometric and low-order polynomial, or simple combinations of these forms, would be sufficient to capture the trend.

Simulations were performed for a single test room, with variable sun and blind angle. Simulated vertical illuminances were calculated at 5 different depths in the room, but with most of the values at just 3 depths (20%, 50% and 80% of the distance to the back wall). In addition to the 12 vertical illuminances facing the window wall, we also had data for 6 horizontal illuminances, and two side-wall illuminances. The horizontal illuminances were at 3 different depths from the window wall. The simulated data covered 93 different sky and blind conditions where the sun was blocked, and an additional 8 conditions where it was not. The latter 8 conditions were not analyzed for the fit, as they did not meet the hypothesized limits for a good fit, and in fact showed significant deviations from the patterns shown under the conditions where the sun was blocked.

A plot of the simulated illuminance values versus the distance from the window showed that very similar trends for the three types of measurements described above. We found that the following relatively simple form captured over 98.6% of the variance of the simulated illuminance values:

$$E_s = E_m \cdot K_{orient} \cdot \exp((K_2 \cdot distance + K_1) \cdot distance) \quad (\text{A.1})$$

In this equation E_s is the simulated illuminance at one of the simulated sensors, E_m is the simulated illuminance for a monitoring sensor near or at the back of the room, K_{orient} is a value that depends on the orientation of the simulated sensor (facing the window, facing the side, or horizontal) and K_1 and K_2 are constants describing the fall-off of the illuminance as a function of the distance from the window. The monitor and orientation terms alone explain about 80% of the variance. Adding the distance terms reduces the residual variance by more than another factor of ten, to give the final R2 values of 98.6% with the monitoring sensor located at a 6 foot height on the center of the back wall.

E_s was fit to points at three locations side-to-side in the room, three blind angles (0°, 30°, and 60°, with positive angles indicating that the blinds redirect light to the ceiling), and nine different solar azimuths. We assumed that these would be less important factors than overall level, orientation, or distance, and the data analysis bore this out. In our limited data set, the effects of solar azimuth and the side-to-side location in the room were weak and somewhat entangled with each other. Azimuth was the stronger effect, but was still only gave $\pm 2\%$ differences in fitted values. The blind effect was significantly stronger, and made a difference of $\pm 14\%$. The illuminance levels varied by a factor 6 from the front to

the back of the room, and the illuminance at the rear monitor sensor varied by a factor of 360 over the range of solar altitudes (excluding direct sun), blind angles, and reflectances studied.

A.1.2 Current work

The main issues were validation and generalization of the overall illuminance level and distance effects. As an initial priority we decided to focus on vertical illuminance, as this determines the glare level which is the trigger for the blind control. The simulation runs were confined to a single room size, so the first issue was to determine how to generalize the fits to different room sizes and shapes. We expect that the constant relating the monitor value, E_m , to other locations in the room will depend on the distance from the window, as well as surface reflectances and room width. In addition, the values of equation (A.1) start to increase when distance from the window is over 35 feet, which is not physically possible, so we wanted to determine the distances over which equation (A.1) is valid.

A.1.2.1 Computational Methods

Two computational methods were used: (1) ray-tracing simulations with the Radiance lighting simulation software package and (2) approximations using form factors.

A.1.2.1.1 Radiance

Radiance is a stochastic ray-tracing tool that derives results from models of light, materials, and geometry. It is capable of accurately simulating realistic scenes, which form factor methods cannot. However, if one is (as we are) trying to derive closed-form functions that relate illumination and geometry, the only way to apply Radiance is to run many simulations and experiment with different functions; Radiance gives no help in choosing the functions.

A.1.2.1.2 Form Factors

Prior to the use of computers, approximate methods were developed that were based on form factors. A **form factor** represents the fraction of light leaving a uniform emitting Lambertian surface that is incident on a receiving surface. Analytical formulas for form factors have been derived for parallel and perpendicular rectangular surfaces. Information about the distribution of light on the receiving surface requires the use of **configuration factors**, which give the illuminance of a point on a receiving plane that is facing or is perpendicular to an emitting surface.

A.1.2.2 Finding a Functional Form for Illumination

The glare evaluation points were assumed to be at a height of 4 feet (1.22 meters), which is the approximate eye height of a seated adult. The side-to-side variation of the configuration factor along a fixed height is a maximum of 2:1 for the center versus the edge. This is solely due to the direct component. If we assume an average reflectance of about 50%, and equal direct and reflected emittances, then the side-to-side variation drops down to around 15% by the time the sensor is a meter or more in from the window wall. This variation implies a $\pm 2\%$ variation in the subjective rating as a function to the side-to-side location, and is therefore not a high priority for further modeling.

At the front of the room, the direct component, which is calculated using the form factor, dominates the illuminance on a plane parallel to the window. At the back of the room, the reflected component is expected to become larger. To get an idea of the magnitude of the reflectance effect we performed a

crude calculation for its magnitude by approximating the reference room as a simple sphere with a small port. This gives a 50:50 estimate for the direct and reflected total fluxes. In the simple sphere model the reflected flux gives a uniform illuminance on all the surfaces. The illuminance from this reflected light on the vertical plane facing the window can be computed from the form factors from the surfaces to the plane. For our reference room, this calculation gives a reflected to direct component ratio which ranges from 2 percent at the front of the room to 50 percent at the back. The actual ratio at the back of the room is probably considerably less than this, as the actual reflected component will be higher at the front of the room than at the back. If we assume that our 50 percent estimate is an upper bound to the true value, it suggests that the form factor is the dominant component, but it also suggests that the reflected component somewhat flattens the fall-off light that would be expected from the form factor alone. This led us to examine the possibility of using form factors, or possibly simple modifications of them, to model the simulated data for the illuminance distribution as a function of distance from the window. Using a simple split flux model, and an average reflectance of 50% for all surfaces, except the window wall at about 10%, leads to an estimate of about a 50:50 split for the reflected versus the direct component as the calculation location approaches the back of the room. If we assume that this is at least order of magnitude accurate, it suggests that the direct component is a major or dominant factor at all distances, and therefore should strongly constrain the functional form of the illuminance versus distance from the window wall.

Unfortunately the form factor is given by a moderately complex function of logarithms and inverse tangents involving the width, length, and height of the room, along with the reflectances, and it is not clear how this form might be modified to give an approximate measure of the light distribution in a space with non-zero reflectance. An approximation to the form factor that is simple enough that it is clear how to include the effects of reflectance would be useful. Our previous work showed that it was possible to get excellent fits with a fixed room size. The question going forward was to see if we could derive a more general function.

Our original fit was a function of the distance from the window, however, the IESNA handbook notes that for moderate to low cavity ratios¹, the form factor is relatively insensitive to the width to height ratio, and can be approximated as a function of the cavity ratio. In our case the cavity ratio equals $5 \cdot D \cdot (H + W)/(H \cdot W)$, where D is the distance from the window, and H and W are the height and width of the room. Going forward, we looked for other trends or bounds. We found two fairly obvious ones. For a very deep room, the form factor for the flux from the window wall to the back wall has to approach the form factor for a point source, which is simply the inverse-square law. If one flattens the room, bringing the back wall to the window, the form factor has to equal one. These two extremes in distance are the only special points for the form factor function, so we looked to see if there was any more information available at these limits. We found that at zero distance the first derivative of the form factor with respect to the cavity ratio is bounded and is equal to -0.2. This means that the form factor falls as a constant divided by the distance for distances approaching zero.

¹ "Cavity" = interior of room. The *cavity ratio* is proportional to the perimeter of the back wall of the room, divided by the area of the back wall of the room. When $W \gg H$, $CR \approx 5D/H$. When $W \approx H$, $CR \approx 10D/W$.

These three constraints can be met by an inverse quadratic, and we therefore hypothesized that an inverse quadratic function of the cavity ratio, with adjustable parameters, should be capable of approximating the form factor fairly accurately, especially over a bounded range of aspect and cavity ratios, and should provide a reasonable base for the functional form to be expected when the reflected component is added (see equation 3 in the beginning of the results section as an example). Figure A-1, below, shows a plot of form factors for rooms with a width/height ratio W/H of 1, 2, and 10. The first two ratios bound the small room simulations that we have performed. The cavity ratios were computed for distances from the window wall over the range from 0.031 to 3.1 times the height (from 0.25 to 25 feet assuming an 8 foot high room). The graph shows that the form factors are indeed insensitive to this range of width/height ratios. A single least-squares inverse quadratic fit (fit1) fits this data to within a maximum error of 7% (this deviation is small enough that the fit line is hard to see on the graph). However, as can be seen from the graph, the form factors do become different when the width/height ratio reaches 10, which is a reasonable value for a large room. The inverse quadratic concept still works, in that coefficients can be found that give a fit to within 2.1%, but the error approaches 40% if the coefficients are not changed.




Figure A-1: Form factors as a function of width to length ratio (W/H) and the cavity ratio. The form factors are shown as x's. Least-squares inverse quadratic functions to the data are shown as lines, but are hard to see as the fits are very close to the data.

As an aside, we noted that inverse quadratic with coefficients which are constrained to fit the theoretical limits described above fits better than might be expected, and has a maximum error of about 23% over the range of cavity ratios shown in Figure A-1. The error for the constrained fit for a width to height ratio of 10 is *lower* over than for the lower width to height ratios, at least over the range of distances plotted in Figure A-1. For larger width to height ratios the fit overestimates the form factor for medium size values of the cavity ratio.

The least-squares fits to the form factor show that fixed constants provide a good fit for small width to height ratios, but may need to be modified for wide and deep rooms, where $W \gg H$, and the cavity ratio exceeds values of about 10. The pattern of error with the theoretical fit suggests that it might be

possible to get a correction term for the coefficients as a function of cavity ratio, but it should be noted that the amount of error that this actually introduces into our calculation is not actually clear for the problem as a whole, as the above calculation does not include reflected light. We expect that reflected light will reduce the error, but it also means that the coefficients that are derived for the inverse quadratic form factor approximation are not likely to be the best coefficients for the problem as a whole.

We have looked at incorporating reflected light into the modeling, but have not yet made much progress, despite a fair amount of work. The goal here is to either have the model incorporate the effect of surface reflectances, or to have an idea of the limits to our accuracy if we ignore it. In particular, we would like to know if there are some sets of conditions in terms of width to height ratios or reflectances which can be reasonably expected to give significantly different results than those obtained fitting the simulated room conditions. Our modeling has provided some qualitative insights that supports the basic idea of the inverse quadratic approach described above, but at this time we cannot make any stronger statements. This work is described in the appendix.

What we can note is that as distance from the window plane to the observation plane is increased, the contribution to the illuminance on the observation plane from surface areas near the front of the room has to drop. In addition the illuminance on surfaces that are near the observation plane decreases with distance. This indicates that at least at some distance the absolute contribution to illuminance from the reflected component begins to drop. In fact, in the limit of large distances we can treat the walls simply as an expansion of the size of the source. In this limit the illuminance on the plane still follows an inverse-square law, but with a slightly bigger constant.

Conversely, at the front of the room, there is no reflected component, because there is no side surface area. The area available for reflection increases linearly with the distance from the window, so we expect that initially there is a linear increase in the reflected component. The combination of a linear increase near zero distance, followed by a inverse-square drop at large distances, is captured in the inverse-quadratic form by a reduction in both of the coefficients from what is expected from the form factor alone.

This result is in contradiction to the result one gets from a sphere method, where the reflected component is a constant. The later result would add a constant to the form factor, and thus would result in a significant departure from the general functional relationship of the form factor. The sphere model may provide a reasonable first approximation for the overall increase in illuminance due to reflectance, but we suggest that it is not adequate to describe the spatial distribution resulting from sidelighting. Instead, the qualitative argument supports the inverse quadratic form.

A.1.3 Summary

We have not had time to successfully extend the theory to handle reflectance. Initially we hoped to show that the illuminance decline with distance from a window should closely follow the form factor, and more specifically the inverse-quadratic approximation to the form factor. To date, we have developed qualitative arguments that show that the inverse quadratic form should fit the illuminance

distribution at the limits of both near and far distances from the window. In addition, we have shown that the form factor can be approximated moderately accurately by a simple function with constant coefficients for small room geometries. We do not have proof that the inverse quadratic will always provide good fits in the middle range of distances when both the direct and reflected components are included, although the bounding of the fits as described above makes the situation look hopeful. For a mix of small and large rooms, it appears that it may be possible to derive a width to height ratio correction that would continue to provide good fits to the form factor, and by extension to the room problem as a whole.

We do not currently have a good theoretical model for the effect of changes in room reflectance on the functional form of the drop-off at intermediate distances. We are hopeful that the inverse-quadratic form will continue to work, but we do not have a quantitative understanding of how the reflected component will affect the coefficients of the function.

We expect that future work on the theory to have a lower priority, as the work to date does confirm a basic form. If time permits, we will further examine the multiple bounce formulas to see if it is possible to get further insight into the proper form for approximate solutions.

A.1.4 Conclusion for the theory section

The modeling suggests that an inverse quadratic function of the cavity ratios, with fixed coefficients, should be suitable for modeling a wide range of rooms sizes, shapes and reflectances, as long the light enters the room in a relatively diffuse manner. Further numerical calculations should help provide better bounds on the effect of these factors. Further modeling may also help in evaluating the effect of partial window walls.

A.2 Results to date

A.2.1 Prior work

In the prior work, simulated data for a reference room, 8 feet high, by 10 feet wide, by 15 feet deep, and with near-standard reflectances: wall = 50%, ceiling = 84%, and floor = 16%, was fit to an empirical function of the depth from the window wall. There were 93 runs covering a limited range of sun and blind angles, and the three CIE sky types. Both vertical and horizontal illuminances were analyzed as a function of the illuminance on several monitoring sensors placed on the walls, or at the front on the window wall. The empirical function provided fits with a standard deviation of about 20%, which is sufficiently small to be useful. We also fit the ratio of the illuminance on a monitor sensor to external illuminance as a function of blind angle. This plot showed a relatively linear relationship, especially at high blind angles. This relationship means that control of the glare by varying blind angle should be relatively straight-forward, at least as long as blind angle is positive.

A.2.2 Current work

The current phase of work added another 228 runs. The new runs contain vertical side-wall facing illuminances in addition to the illuminance positions for the previous runs, however our initial analysis has just looked at the vertical window-wall illuminances, as these are the ones that are most likely to create glare. The new runs were distributed as follows:

# of runs	Purpose
9	Check changes to simulation algorithm
33	Vary wall reflectances (30% reflectance on 1 or more walls)
27	Restrict window to only a portion of the window wall
30	Change room shape (15,20,8; 20,10,8; and 20,10,10)
8	Effect of blind angle with overcast sky
24	Negative blind angles (-30° & -60°) with a clear sky
81	Positive blind angles - extended range of sun angles
16	Fixed 75° blind angle - partly cloudy sky at different sun angles

In the above table, if the sky type is not mentioned the runs were done with a clear sky, with a single check run done with an overcast sky.

We fit this new data, plus the original data as a function of an inverse quadratic, as described in the theory section. The equations are as follows:

$$E(CR) = \frac{E_m \cdot K_1 \cdot Q(CR_m)}{Q(CR)} \quad (A.2)$$

where E is the illuminance at the cavity ratio CR , E_m is the illuminance on the monitor which is located at the cavity ratio CR_m , and Q is the inverse quadratic function:

$$Q(x) = \frac{1}{1 + K_2 \cdot x + K_3 \cdot x^2} \quad (A.3)$$

Where $K_1 = 0.8843$, $K_2 = 0.04925$, and $K_3 = 0.01996$.

Ideally, the constant K_1 gives the ratio between the illuminance at the monitor sensor location (6 feet up, along the center-line from front to back) to the average illuminance across the width of the room at 4 feet high at the back of the room. The product K_1 times $Q(CR_m)$ gives the corresponding ratio for the front of the room. The constants K_2 and K_3 determine the shape of the function, and were empirically derived as the constants that give a best fit to the first two hundred runs (the difference between the best fit to the data set plotted in Figure A-3, which is described below, and the values listed was so minor that we didn't bother to change the fits or the plots).

Figure A-2, below, shows the data for these new runs combined with the vertical illuminance data for the old runs, as fit against an inverse quadratic function of distance times the illuminance of a monitor sensor mounted 6 feet above floor on the center of the back wall. The root-mean-square error of the fit is about 30%, but there are some very obvious bad points, with the worst point being a factor 17 higher than the estimated value.




Figure A-2: Measured (simulated) illuminances versus fitted vertical-window wall facing illuminances. Data for all 321 runs. The vast majority of the bad points, and all the very worst points, come from either the runs with negative blind angles, or the runs with the windows covering only part of the window wall. The two worst underestimates were points where the blind angle was negative, and the sun illuminated the front sensors in the room. The remaining significant underestimates are also for negative blind angles, but are due simply to the sky being significantly brighter than the blinds. The errors are again primarily at the front of the room.

The significant overestimates of the illuminance are for conditions where the windows only cover part of the window wall, and are again for sensor locations near the front of the room. In the standard room, the window wall had three windows that covered the width of the wall, and much of the height. In the partial window wall cases, only one of the three windows was installed. Sensors that face a wall near the front of the room, do not see much of the window, and have significantly lower illuminances than sensors to their sides, or even to sensors that are farther back.

Figure A-3, below, shows the fit that is achieved when the partial window wall, and negative blind angle conditions are removed from the data. The root-mean-square error was about 15%, with a worst-case overestimate of about 60% (log-error of 0.49), and worst-case underestimate of about 33% (log-error of -0.40). The worst underestimate was a zero-degree blind condition, and most of the other significant underestimate cases were also for this condition. The worst overestimate was for the extra wide room with the low ceiling, and once again most of the other significant overestimates were for this condition.

The problem of underestimation is an extension of the problem that was seen for negative blind angles. Visible bright sky is much reduced at zero degree blind angle, but is not totally eliminated. A slightly higher blind angle should give better results, although the current results are not that bad.




Figure A-3: Measured (simulated) illuminances versus fitted vertical-window wall facing illuminances. Data from 270 runs. Runs with negative blind angles or partial window walls were excluded.

In the theory section we noted that a inverse quadratic with fixed coefficients could provide a fairly good fit to the form factor over a range of width to height ratios, but noted that the best fit does depend on this ratio. We also noted that a real room with reflectance would not behave identically to the form factor. Figure A-4 plots the logarithmic error of the fits for the non-reference room shapes as a function of the cavity ratio. Two types of errors are seen. Two of the rooms show a tendency to overestimate the illuminance in the center of the room, which indicates that the shape coefficients of the standard fit are not ideal for these rooms. However the largest error in all of the fits is that there is an overall shift, or bias, in the calculated values relative to the measured values.




Figure A-4: Error in the reference fit for the three non-reference room shapes. The error has a definite shape as a function of the cavity ratio, but the main error seems to be an overall bias error.

The best fit coefficients for the three non-standard rooms are different. For the two deep rooms, the value of K_1 does not change much from that of the standard room, but the value for the wide room is significantly smaller (0.78). This appears to be an artifact of our choice of the monitoring locations in the simulations. For the most of the normal width rooms, the vertical illuminance was measured (simulated) along the center line, and ± 2.5 feet from the center line (front to back). In the wide room the values were measured at ± 7.5 feet and ± 2.5 from the center line. The monitor sensor is located

along the center line at the back of the room. Figure A-5 shows that there is a small, but statistically significant effect on the illuminance (shown as the error in the fit) as a function of the distance from the center line of the room. After correcting for the difference in the simulation locations, the value for K_1 for the wide room becomes 0.8714, which is not significantly different than the value for the reference fit.




Figure A-5: Error in the reference fit as a function of the distance from the center line from front to back in the room. The data set is the reduced set used in Figure A-3.

The side to side effect also provides an explanation for why K_1 is as low as it is. We can determine K_1' , the ratio of the monitoring sensor to the illuminance on the center line at the back of the room, by adjusting K_1 by the estimated center line value. Our best fit for the side to side effect is an inverse quadratic in the absolute value of the distance from the center line, d :

$$Q(d) = \frac{1}{H_1 + H_2 \cdot d + H_3 \cdot d^2} \quad (\text{A.4})$$

with fit constants:

$$\begin{aligned} H_1 &= 0.9521, \\ H_2 &= 0.0217, \text{ and} \\ H_3 &= 0.00315. \end{aligned}$$

Dividing through by H_1 gives $K_1' = 0.931$, which is a not totally unreasonable estimate for the effect of height on the vertical illuminance at the back of the room. Inclusion of $Q(d)$ in the reference fit reduces the root-mean-square error to 14%, and the worst case errors to 45% overestimate and 30% underestimate.

It is important to note that while $Q(d)$ provides a good fit to the mean trend of the side to side illuminances, it is just an empirical fit. We do not have a theoretical justification for this form. The inverse quadratic is bounded and well-behaved, but we strongly suspect that it will not provide good fits for rooms that are significantly wider than our current test rooms. We expect to have to revisit this analysis after we examine the large rooms in our next phase of the study.

Returning to the general question of the effect of room shape on the fit, we find that the constant K_3 was fairly stable for the non-standard and reference rooms, but K_2 is significantly higher in the non-standard rooms, especially for the wide room and the deep and tall room. The effect of a higher K_2 is to make the difference between the illuminance in the front of the room versus the illuminance in the back of the room larger, and to make the drop off with distance quicker than is predicted by the reference function. We currently do not have an explanation for the more rapid decline for these room shapes. In addition we do not have enough data to provide even an empirical fit as to the effect of shape on these constants. As a practical matter, the error appears tolerable over the range of rooms we have examined. Furthermore to a large extent it creates a bias error, which can be compensated for during commissioning. We expect to get considerably more information on the extent of the shape problem when we begin to analyze the wide room.

The effect of reflectance was examined by simulating the same sky conditions in rooms with different wall reflectances. Nine reference conditions, and 33 different test conditions were examined. The test conditions consisted of 9 runs with a 30% west wall reflectance, 8 runs with a 30% east wall reflectance, 7 runs with a 30% back wall reflectance, and 9 runs with all three wall at 30% reflectance. The changes in reflectance produced statistically significant changes in the absolute illuminance, but the effects were small, and, for the most part are accompanied by a similar change in the illuminance on the monitoring sensor. The maximum overall relative error was 3.4% for the condition with all three walls at 30% reflectance. This overall effect has little or no practical significance.

We noted in the theory section that we expected that a change in reflectance should affect the distribution of light as a function of the distance from the window. This effect is statistically significant. The effect is zero, or close to it, at the front of the room. The effect as a function of distance appears linear versus the cavity ratio (CR), which means that the ratio of the values in the reflectance test rooms, E_t , to the values in the reference room, E_r , fit the form $E_t = E_r \cdot (1 - k \cdot CR)$. The constant k ranged from 0.001 for the back wall being 30% reflectance, to 0.004 for the side walls at 30%, and 0.01 when all walls were at 30% reflectance. As noted earlier, we do not yet have a proper model for this effect. Fortunately, the data indicates that it is not a large effect, at least over the range of reflectances that we examined. We do hope to look at this problem further.

The last item of major interest for this phase of our work was how blind angle affects the overall level of glare. We have already noted that blind angle did not significantly affect the fit of the relative values of vertical illuminance versus distance from the window. However, this still leaves the question of how much the blind has to be tilted to control glare, how reproducible this is, and, at an even more basic level, whether the effect of blind angle on the absolute illuminance level is even monotonic versus the absolute deviation from the full open blind (0°).

To examine this question we fixed the sky type, and solar angles, and examined up to 8 different blind angles, with the maximum, or reference, angle being 75° . A total of 104 simulations were run with clear (96 runs) and overcast skies, 5 solar altitudes, and 3 solar azimuths (0° , 90° , and 180°). The result for each of the different combination of the sky type and solar angles was then scaled against the illuminance with the 75° blind as a reference. This normalization isolates the blind effect from the sky

effects. The results appeared independent of solar altitude, and were only weakly dependent upon sky type. Figure A-6 shows the logarithm of the ratios plotted against the blind angle.

One data point at 15° was very faintly higher than the corresponding point at 0° , but all remaining points showed monotonicity versus the absolute value of the blind angle. The one reversal is well within the noise level of the simulations, and thus it does appear that we have monotonicity.

The overall shape of the function is fit by a quadratic function in blind angle, but there is considerable scatter about the fit line. If the azimuth is not included in the fit, the standard deviation is 22%, but the maximum deviation is over a factor of 2. Including the azimuth reduces the standard deviation to $\pm 11\%$, but the maximum deviation is still a fairly large 37%. This result suggests that blind will require real-time feed-back to attain the proper glare level. We have unambiguous information on which direction to tilt the blind based on whether we are above or below the target glare level, but the magnitude of the tilt change required to reach that level has significant uncertainty to it.




Figure A-6 Effect of blind angle and solar azimuth on the illuminance at the monitoring position.

The vertical axis of Figure A-6 shows the logarithm of the ratio of the illuminance at the blind angle on the horizontal axis, to the illuminance at a 75° blind angle. The ratios were simulated over a range of solar altitudes for clear skies, plus one set of runs for overcast skies. The two dominant effects were blind angle and solar azimuth. Sky type (clear versus overcast) had a statistically significant influence on the blind angle effect, but it was only $\pm 10\%$ of the main effect, and thus has little practical significance. Solar altitude was not significant.

A.2.3 Conclusion

We started the current phase of work with an empirical model of the relative distribution of light as a function of the distance from a window wall. Although this model worked well over the range of conditions studied, only one room size was studied, and it was clear that the model would not work for rooms that were significantly different in shape than this reference room. There was some concern as to how well the model would work with a wider range of blind angles, and with different surface reflectances, and sun angles.

We have now derived a semi-empirical model for the relative light distribution as a function of distance from the window wall. This model is based on the assumption that the window has a uniform, approximately Lambertian distribution of light. The model is more reasonably bounded, and should be able to handle a wider range of room shapes. The simulations have confirmed that this model works tolerably well for positive blind angles, where the uniform Lambertian assumption is reasonable. The simulations also showed that it does not work well when the Lambertian assumption fails, such as when the sky is visible inside the room (blind angles less than zero). Simulation runs have confirmed that the model extends tolerably well to different room shapes. They have also confirmed that the model produces reasonable results when wall reflectances are reduced from 50% to 30%.

The semi-empirical model does not include corrections for reflectance, or the side-to-side variation. We do have a strictly empirical correction for the side-to-side variation. It is not an important factor over the limited size ranges we have examined, but we expect that it may become a problem with wider rooms. We will be simulating these rooms in the next phase, and will be attempting to a semi-empirical model for this dimension which does have appropriate bounds.

Our modeling to date has suggested that while the light distribution depends on reflectance, it may not be very sensitive to moderate changes in it. We found fairly minor differences in the distribution when wall reflectance was reduced from 50% to 30%.

A.3 Next Steps

The current semi-empirical model provides a fit which is as good as that of the previous model, while fitting a wider range of conditions. The model provides a reasonably accurate procedure for assessing the mean level of glare from a window based on the illuminance on a sensor at the back of the room. Equally important, we found that as long as direct sun is blocked, that the illuminance on the rear sensor varies essentially monotonically as the angle of the blinds departs from zero degrees (horizontal). This puts us only one step away from having a model that can take weather file data, and compute the glare level as a function of blind angle. The missing step is the calculation of the absolute illuminance at the rear monitor for a reference blind angle as a function of the exterior vertical illuminance. There are algorithms for the computation of the vertical irradiance as a function of the direct and diffuse irradiance levels, and there are efficacy estimates so that we should be able to estimate vertical illuminance from a weather file. We have the exterior values for our simulated runs, and the missing step of computing the interior illuminance from the exterior illuminance will be an immediate priority for the next phase of work.

Once we have glare levels as a function of exterior conditions, we will be in a position to begin to evaluate the possible energy savings and comfort control for an automatic blind/electric light control system.

A.4 Appendix

The standard procedure that is used to compute effective cavity reflectances and coefficients of utilization is to assume that all surfaces have a uniform Lambertian exitance. This leads to a set of simultaneous equations with form factors that is relatively easy to solve for a room in the shape of a box (3 surfaces - ceiling, floor, and walls). We know from calculations using configuration factors that the direct illuminance on surfaces perpendicular to the window drops by a factor of two at depth to width or height ratios of 0.25 to 0.5 (cavity ratio range of 1.25 to 5). Cavity reflectances and coefficients of utilization have been used as approximations for electric lighting design, where the geometry is such that the ratios are in this range, or lower. The ratios for side-lighting from daylight are in the range of one to two in the small room simulations, and can extend to larger values if a room is narrow and deep. Thus we do not have evidence that this procedure will have much accuracy for daylighting.

In fact, calculations for effective cavity reflectances suggest that this approach may not be accurate at these larger depth ratios. We computed cavity reflectances assuming that the back wall had zero reflectance. In this case the cavity reflectance should increase from zero to an asymptotic limit as the depth of the cavity is increased. The form factor approximation gives a cavity reflectance that reaches a maximum at a cavity ratio that is less than 10, and then begins to drop. The range of cavity ratios for which the effective cavity reflection approximation is used is generally less than 2.5. In short the form factor approach gives non-physical answers for the cavity reflectance at cavity ratio values only somewhat larger than those for which the approximation is generally used, but definitely smaller than the cavity ratios that can apply for side lighting (up to 30).

In light of the questions raised in the above analysis, we examined two other possible approaches for producing a hopefully relatively simple approximation for the fall-off of light in the room. The first approach we tried was to see if it was possible to relax the restriction of uniform luminance of the side surfaces. We found that it was possible to produce an analytically solvable system of equations if we assumed a linear relationship of luminance to distance from the window wall. Unfortunately, the resulting solution for illuminance as a function of distance from the window wall is fairly complex. Worse yet, when we tried calculating effective cavity reflectances as a check, we got results that were only marginally better than those of the simple form factor calculation. Although it is possible that we have made a mathematical error, our analysis of the formula suggests that the lack of improvement is real, and is due to the non-linear nature of the actual distribution of of luminance of the side surfaces.

The second approach that we examined was to see if we could compute the illuminance as a sum. As was noted earlier, a simple split flux model suggests that the direct component should be at least 50% of the illuminance on a plane. In addition, the first bounce should be about 50% of the remaining contribution, the second bounce should be about 50% of the contribution past the first bounce, and so on. Getting the direct component right gives a significant portion of the answer, and getting the first bounce should improve the answer significantly, and so on. We have the direct component, so we have

a major portion of the full form. Unfortunately, even the first bounce looks quite complex, so we have not been able to extend this analysis yet.

Appendix B Blind deployment and blind tilt angle for glare controls

A general recommendation for venetian blind deployment and tilt angle controls is provided based on the simulation-based investigation of glare estimation in Appendix A.

B.1 Introduction

We have fairly extensive information on the effect of blind tilt when the blinds are fully deployed. We have had very limited information on the effect of blind deployment. In the next to last set of runs, we examined the effect of blind deployment (1 = fully deployed, 0 is fully up) over a limited set of blind angles and sun azimuthal angles. The runs used a fixed solar altitude of 45°, and covered both clear and overcast skies. Simulations were performed for the large deep room and the reference small room.

Solar altitude was ignored as it has shown little effect on our fits in past runs. Azimuthal angles were varied in fairly large steps over the full range (0°, 30°, 60°, 90°, and for the small room 180°). Blind tilt was varied in 15° steps, but was limited to those angles where there was no direct sun, given the shade deployment and sun azimuth (blind angles ranged down to zero for the overcast sky).

Two sets of analyses were performed with the data. The first was to examine whether existing fits or fit form for glare illuminance were still valid when the blinds were only partially deployed. The second was to evaluate how blind deployment affected the illuminance on the glare monitor.

B.2 Window glare versus M_1

The reference fitting function for illuminance as a function of the distance from the window is the inverse quadratic as a function of the cavity ratio (CR):

$$M_1 = K_1 / ((K_2 * CR + K_3) * CR)$$

The reference value of K_1 depends upon the cavity ratio of M_1 , CR_0 , and is given by the formula:

$$K_1 = K_0 * ((K_2 * CR_0 + K_3) * CR_0), \text{ where } K_0 = 0.88426$$

K_2 and K_3 were fixed at the values of 0.019958, and 0.049247 respectively. K_0 is effectively the ratio between illuminance at the back of the wall at glare height (4 feet) versus the illuminance at the monitor height (6 feet).

This formula was developed with the original small room data. When we extended our analysis to other room shapes, and a larger set of blind conditions, we kept the form, but allowed for variation in K_0 , K_2 and K_3 , or all three. Generally, with the glare data, the fits with K_0 and K_2 varied appear to provide the best compromise fit. They do not give as good a fit as when all three variables are free, but the resulting values of K_0 appear to be the most reasonable.

At the time it was developed we noted that there was a side-to-side variation in the degree of fit, with those points closest to the side walls having significantly more poorly than all the remaining points. In the fits to the large rooms there were considerably more side-to-side points, but we found surprisingly that we got almost the same pattern: those points closest to the wall were different from all the remaining points. Based on the data that we have, it appears that illuminance is significantly lower

within about 3 feet of a sidewall (the most extreme point in all the rooms examined), but the magnitude of the effect depends upon how wide the room is. A simple modification of K_1 to account for this edge effect significantly improves the fits, and is therefore included in our current analysis:

$K_1(\text{new}) = K_0 * X_{\text{edge}} * ((K_2 * CR_0 + K_3) * CR_0)$, where X_{edge} has a different value depending upon whether it is within 3 feet of the side wall, or farther.

B.2.1 Large deep room

In Table B-1, below we compare the fit coefficients for the original large deep room study (labeled Orig in columns 2 and 4), and the coefficients in the current study where blind height was varied (labeled New in columns 3 and 5).

Table B-1: Fit Coefficients of the Original Deep Room Study and Current Study

	Original	New	Original	New
K_0	0.86025	0.88946	0.94401	0.97622
K_2	0.01996	0.01996	0.01996	0.01996
K_3	0.04925	0.04925	0.68278	0.68571
$X_{\text{edge}} > 3$	1.057	1.054	1.057	1.054
$X_{\text{edge}} < 3$	0.774	0.784	0.774	0.784

As can be seen from the table, the edge effect is almost identical in both sets of data. There is about a 3% difference in K_0 between the two sets of simulations, but the drop-off of illuminance with distance from the window, which is fit by K_3 , seems almost identical. The increase in K_3 for the deep room has been noted before. It appears likely that it is large part due to the workstations blocking light from the windows.

Table B-2: Degree of Fit Achieved with Each of the Fits Described Above

	Original	New	Original	New
R^2	0.974	0.966	0.990	0.982
Maximum	0.28	0.34	0.28	0.36
Minimum	-0.82	-1.07	-0.47	-0.78
RMS	0.167	0.192	0.103	0.140

The fits with variable blind height are not as good as the fits with blinds either fully or half deployed, but they are still quite good. For the large deep room, these fits suggest that varying blind height does not have a major impact upon the predictability of glare from M_1 , as long as direct sun is excluded from the measurement locations.

B.2.2 Small room

Table B-3, below, shows the fits for the 2nd set of small room fits (column two in the table: Orig data), plus the fits to latest set of data that explicitly looked blind deployment (columns 3 and 4). The original

small room data consisted of two sets of simulations. The coefficients listed in column 2 of Table B-1, minus the edge effect coefficients, were those derived for the first set of simulations. The best fitting coefficients for the second set of simulations, including an edge effect, are shown in column 2 of Table B-3. There was no significant improvement in the fit when K_3 was changed, but the edge effects were significant. The value of K_0 is significantly higher than what was found in the original data set, but it is change of only 5%. The first set of simulations was limited to a blind angle of 30°, while the second set contained blind angles from 0° to 75°, and it appears that the net effect of this change in blind tilt angles was to slightly increase the ratio of the glare illuminance at the back of the room at four feet high relative to that at six feet high.

The fits for the new simulations with variable blind deployment echo those found in the large room, in that varying K_3 does lead to a better fit (although not by much). Furthermore, letting K_3 float, results in a fit with about a 3% larger value of K_0 , which again echoes the effect found in the large room. The small room fits also partly echoes the large room fits in that the value of K_0 is larger for blind deployment data, than for the data with the blind deployed. The echo however is only qualitative in this case, as the effect is significantly larger in the small room than in the large room. The last item of interest in this table is that there is a significant edge effect, although it is substantially smaller than in the large room, which is much wider, and this effect is again essentially unaffected by blind deployment.

Table B-3: Fits for the Second Set of Small Room Fits

	Original Data	New K_0	New K_0 & K_3
K_0	0.91354	1.01837	1.04947
K_2	0.01996	0.01996	0.01996
K_3	0.04925	0.04925	0.18892
$X_{edge} > 3$	1.043	1.044	1.044
$X_{edge} < 3$	0.979	0.979	0.978

Table B-4: Degree of Fit Achieved with Each of the Fits Described Above

	Original Data	New K_0	New K_0 & K_3
R^2	0.991	0.976	0.978
Maximum	0.356	0.333	0.323
Minimum	-0.364	-0.544	-0.505
RMS	0.117	0.152	0.149

For the small room, as with the larger room, the fits with changes in blind deployment are not as good as those with the blinds fully down, but again the difference is fairly small. Both sets of fits indicate that blind deployment is not a major problem for glare prediction as long as deployment is sufficient to exclude direct sun.

B.3 M_1 vs blind deployment and shade height

In the above section we showed that a reasonably good evaluation of room glare can be made via the readings of M_1 , regardless of the state of the blinds. In the second stage of the analysis we examined how M_1 varied with the control variables; shade deployment and blind tilt angle. We previously had examined the effect of blind tilt angle, but only for full or half deployment of the blinds, and only for a clear sky.

The perpendicular distance between the slats of a flat blind is proportional to the cosine of the blind tilt angle. The flux that passes through the blind should therefore be related to the cosine of the tilt angle, but it should not be expected to be directly proportional, both because of inter-reflections between the slats, and because real slats are curved, not flat. A semi-empirical formula that appears to work surprisingly well is to set M_1 to be proportional to a constant times the exponential of the cosine of the blind tilt angle: $\exp[B_1 * \cos(BTA)]$.

For blind deployment the maximum transmittance is when deployment (BD) is zero, and the minimum is when deployment is 1. The minimum will depend upon the blind tilt angle. We expect transmissivity through the window to exhibit a simple linear trend versus BD from the minimum to maximum transmissivities. However, the room cavity transmittance, and hence the illuminance on M_1 , depends not just on the overall window transmittance, but also the angular distribution of the light through the window. The inter-reflected light from the blinds, particularly when the blind tilt angle is high, is primarily directed towards the ceiling, which has a much higher reflectance than the floors, or walls, and is definitely more visible from M_1 than the area behind a workstation. We again examined semi-empirical formulas based on the variable of interest, which in this case is BD. Good fits were obtained using the exponential of BD times a constant, B_2 , plus a constant, B_3 , times the product of BD and $\cos(BTA)$. For these fits, at low blind tilt angles, the effective constant for blind deployment is small, and the resultant function is almost linear. Conversely, at high blind tilt angles the effective constant becomes moderately large, and the exponential deviates significantly from a simple linear fit.

The simulations covered both overcast and clear skies, and covered a number of different sun to window azimuth angles. These factors will, at a minimum, affect the illuminance on the window, and thus need to be included as scale factors. Since the sky factors are not control factors they were simply fit as categorical variables so that the data could be adjusted to a common level. A second concern is that sky type and azimuth may also affect the constants B_1 through B_3 . The full fit as described here is a linear fit for the logarithm of M_1 :

$$\ln(M_1) = B_{\text{intercept}} + B_{\text{sky}} + B_{\text{azimuth}} + B_1 * \cos(BTA) + B_2 * BD + B_3 * BD * \cos(BTA),$$

where $B_1 - B_3$ may be functions of sky type and azimuth.

B.3.1 Large deep room

For the large deep room, the JMP statistical program rejected interaction terms between $B_1 - B_3$ and sky type or azimuth. The best fit is described in Table B-5, below:

Table B-5: Fit of $\ln(M_1)$ vs Linear Fit Equation

Item	Value
B_1 (tilt)	-0.6883
B_2 (deployment)	-4.1459
B_3 (cross term)	3.8601
R^2	0.977
Minimum Error	-0.403
Maximum Error	0.420
RMS Error	0.139

The values for the parameters B_{sky} and B_{azimuth} have not been listed in the table because the simulations only covered one sun altitude angle. In addition, these two parameters do not affect the control algorithm, as it can only vary shade height and blind tilt angle. Note more importantly that the data for shade deployment with only one exception did not go below 0.25. The fit gives a non-physical trend versus blind tilt for shade deployments below 0.18. A fit without B_1 gives more physically reasonable fits for small shade deployment values, but is less accurate for the remaining range ($R^2 = 0.972$, RMS error = 0.152), which is the more useful range.

From the values in Table B-5, it is easy to show that the fitted response of M_1 to blind deployment is nearly linear at a blind tilt angle of 15°, begins to noticeably deviate at 30°, and is significantly different from linear at 45°. A separate check using a linear fit of blind deployment versus M_1 (a logarithm fit against $\ln(M_1)$), shows that the linear fit is slightly better at 15°, much better at 30°, worse at 45°, and much worse for larger angles. This is consistent with the actual data starting out with an almost linear response to deployment when the blinds are nearly open, changing to a non-linear response as the blinds are tilted and various inter-reflection effects become important.

A quick check was made to see if the results for the deep room are consistent with the more abbreviated results (full and half height only) obtained earlier for the standard large room. For the large room there were 30 data points that were fit separately in groups of 5 for the two blind deployments and with different azimuths, to the form:

$$\ln(M_1) = B_{\text{intercept}} + B_1 * \cos(\text{BTA})$$

After correcting for the depth of the room and the placement of M_1 , the intercepts for the large room averaged 0.14 larger than for the deep room, while the slopes averaged -0.44 smaller. The results look similar, but are not identical.

B.3.2 Small room results and discussion

We have a slightly abbreviated data set for the small room, as five of the simulation runs are significantly out of line in comparison to the remaining data. Inclusion of these points produces fit values that are non-physical. These points were from a separate earlier set of simulations, and I have therefore assumed that there was an isolated input error, and I have therefore dropped them from the

analysis. The remaining data set consists of 85 runs, and is sufficient to give a reasonably good picture of the blind effects.

The small room results do not perfectly match the large room results in that the JMP program found a statistically significant interactions between sky type the coefficients B_1 to B_3 , and a moderately strong interaction between azimuth and B_1 . Azimuth in the fits without interaction is simply entered as a categorical variable, since when it doesn't interact with the blind or shade coefficients there is no need to fit a pattern to it. In the fit with interactions there is no clear pattern in the fit coefficients as the azimuth increases from zero to 180 degrees, so the fit is useless for predictive purposes. It is possible that we are seeing the effects of reflections off the furniture, in which case the fit wouldn't generalize anyway.

The fits without the azimuth interaction term are still fairly good. The probability value for the sky interactions is 0.7%, but even the basic fit without this term still returns fairly good predictions. Table B-6 shows a comparison of the degree of fit for the 3 fits:

Table B-6: Comparison of Small Room Fits

	Basic Fit	Sky Interactions	Sky and Azimuth Interactions
R²	0.964	0.969	0.978
Minimum Error	-0.657	-0.543	-0.490
Maximum Error	0.402	0.368	0.422
RMS Error	0.183	0.172	0.148

Table B-7 shows the fit coefficients for the large room as compared to those for the small room. For the small room, the coefficients are shown for the basic fit with no sky interactions, and then with the fit with the sky interactions for the clear sky and overcast sky separately:

Table B-7: Comparison of Deployment and Blind Tilt Values for Small vs Large Room

Parameter	Small Room Values			
	Large Room	Basic Fit	Clear Sky	Overcast sky
B₁(tilt)	-0.688	-0.933	-1.740	-0.471
B₂(deployment)	-4.146	-4.059	-4.894	-3.494
B₃(cross term)	3.860	3.935	4.867	3.277

The parameters in Table B-7 are for a semi-empirical fit, so it is important to note that with one exception the data did not go below a blind deployment value of 0.25. The fits above are only valid so long as $B_1 + B_3 \cdot \text{shade height}$ is greater than zero. As noted earlier, this limits the fit for the large room to shade deployments above 0.18. For the small room, values of shade height below 0.14 (overcast sky) to 0.36 (clear sky) the combined tilt coefficient is less than zero, and formula predicts that M_1 increases as the tilt increases. This is clearly non-physical, as tilt must have increasingly less effect as blind

deployment goes to zero. Simple fits without B1 are not quite as good over the data range that we actually expect to be useful ($R^2 = 0.956$, and an rms error of 0.2). A better form is probably possible, but it is not worth the effort at this time.

Table B-7 shows that although we can get good fits for a particular room geometry, there are substantial differences for different rooms, and in the case of the small room, for different sky types. Our interest is primarily in relative differences with respect to changes in deployment or tilt. Because there is an interaction term in the above table, the differences in the fit aren't as clear as they might be. Table B-8 shows the shade deployment coefficient as a function of the blind tilt angle. From Table B-8 it can be seen that the biggest difference for shade deployment is at high blind tilt angles.

Table B-8: Shade Deployment Coefficients

Small Room Values				
Blind Tilt	Large Room	Clear Sky	Overcast Sky	Basic Fit
0	-0.29	-0.03	-0.22	-0.12
15	-0.42	-0.19	-0.33	-0.26
30	-0.80	-0.68	-0.66	-0.65
45	-1.42	-1.45	-1.18	-1.28
60	-2.22	-2.46	-1.86	-2.09
75	-3.15	-3.63	-2.65	-3.04
90	-4.15	-4.89	-3.49	-4.06

Because glare sometimes cannot be blocked by blind tilt until the blinds are below eye height, our proposed blind algorithm was to lower the blinds first, and then increase tilt. With the shade system that we have, a change in blind deployment is a two-step process. The shade motor only lowers or raises the blinds when they are at 90° tilt. Figure B-1 shows the relative value of M_1 versus deployment at 15° tilt, which is the tilt that we proposed, and at 90° tilt, which is the tilt that the blind will actually have while moving:




Figure B-1: Relative glare monitor (M_1) value as a function of the shade deployment. There are eight curves. Clear, OC and avg sky indicates that the curve is for the small room with the clear sky or overcast sky interaction terms, or no interaction term respectively. "large room" indicates that the curve is for the large room. The degree values in the legend give the blind tilt.

The important thing to notice about the fits for a blind tilt of 15° is not the differences between them, but rather that regardless of the room or sky, blind deployment makes little difference to the illuminance in the room at this tilt angle. A 30% effect is fairly negligible, especially in comparison with the factor of 20 difference obtained when the blind tilt is close to 90° . The disparity between the two sets of curves raises an important human factors issues. For example, if the shades are up, then deploying them to a 50% level means first cutting the daylight component to ten or 20 percent of its initial value, and then raising it back up to 80 to 90 percent of its initial value. Dropping the blinds again to get to full deployment repeats this process of first dropping the light level dramatically, and then raising it back up to near to its initial level. The process of lowering or raising the blind and then adjusting the blind tilt appears as if it is likely to be quite disruptive, and therefore should be minimized. Given that fully deployed open blinds only appear to reduce the glare illuminances by at most 25 to 30 percent it appears that it might be reasonable to simply limit operation of the blinds to fully up or fully down. Note that depending upon the sill height and window orientation, fully down could be defined as whatever height blocks direct sun. The main lesson from Figure B-1 is to limit the number of height changes, and not whether the blind actually goes to the window sill. However, if the blind does not go down to the sill, this will restrict the range of glare control possible, and will also affect the rate of glare control possible by tilting the blind.

If the blinds are operated in an "up" or "down" mode, then blind tilt is only a concern when the shades are fully deployed. For this condition the terms in B_1 and B_3 collapse to a single term in B_1' , where we no longer have to consider shade deployment. Figure B-2, below, shows of M_1 versus blind tilt angle for the fully deployed condition:




Figure B-2: The fits of the relative value of M_1 versus blind tilt in both the large and small rooms. Clear, OC and avg sky indicates that the curve is for the small room with the clear sky or overcast sky interaction terms, or no interaction term respectively. "large room" indicates that the curve is for the large room.

The curves show that the fits of relative value of M_1 versus tilt angle are, in an absolute sense, fairly insensitive to sky conditions, and even to room shape. The absolute plot hides the fact that at large tilt angles the ratios of the largest and lowest M_1 values for the fits are moderately large. Our earlier memo on this topic indicated that there can be almost a factor of two difference due to sun position, and possibly another factor of two between low reflectance and high reflectance blinds. These are the ratios for the relative values at a blind tilt of 90°. At lower angles, the differences are smaller. In short, Figure B-2 shows that in a gross sense the trend of M_1 versus blind tilt is similar across sky condition and room type, but a look at the logarithms or ratios shows that there are significant differences in glare reduction at the large tilt angles.

There appear to be two fairly different possible control strategies. The data from the earlier blind memo suggests that M_1 is the least responsive to blind tilt for the condition which potentially causes the most glare (sun directly normal azimuth to the window on a clear day). This suggests using a low value for B_1' of 2.8. Conditions which potentially cause less glare will have higher values of B_1' , but since this when there is less potential glare, the blind tilt needed to block it will be less, and the error between using B_1' and the actual value will be reduced. Alternatively, an adaptive strategy can be used if we have the computational power for it. In the adaptive mode, the controller initially uses a default B_1' to tilt the blinds, and then checks the predicted drop or increase in M_1 versus the predicted value. The value of B_1' is then adjusted using the simple logarithm fit. The new B_1' becomes the default B_1' value for the next adjustment. If the value of M_1 is outside a tolerance range, which we need to determine, the blind tilt is adjusted using the new default, and the process is repeated.

B.4 Summary

The data suggest that M_1 can be used to predict glare illuminances within a worst case factor of two error regardless of blind deployment, sky condition, blind tilt and room geometry. The root-mean-square error is much smaller, being on the order of 10 to 20 percent. The biggest errors are when there

is some sun in the space. This suggests that M_1 can be used to control the blinds, and thus glare, in a fairly reasonable fashion.

The fits of the effect of blind tilt and deployment suggest a control strategy of fully deploying the blind and then adjusting the tilt to zero degrees when glare illuminances first reach a target value. Further increases in glare are then controlled by increasing the tilt. The effects of deployment and tilt are monotonic, or nearly so, but there are differences in the magnitude of the effect that depend upon room geometry or sky conditions. This indicates that we will need a commissioning protocol, as well as an adaptive algorithm to get the proper tilt angle to control glare.

Appendix C General commissioning procedure for glare control

C.1 Summary of needed values, capabilities, and actions

- 1) Location and orientation of window: Latitude, Longitude and window azimuth. Determined at commissioning
- 2) Time of day - determined at commissioning
- 3) Algorithm to compute sun location - built in
- 4) Optional: Target time to lower blinds - set at commissioning - allow user to modify
- 5) Optional: Target time to raise blinds - set at commissioning - allow user to modify
- 6) If exterior sensor deployed - set target vertical illuminances for deployment of blinds (default 10,000 lux) and retraction of blinds (default 5,000 lux) - set at commissioning - allow user to modify
- 7) If more than one blind deployment level is allowed - set deployment levels at commissioning and allow user to modify. Also set exterior target levels for the intermediate blind deployment locations. For a two steps (up, partial, and full), set defaults at 5,000, 10,000, and possibly 20,000 lux as the target deployment and retraction levels.
- 8) Measure and input dimensions of room. Check that rear monitor is located on the back wall, is not obstructed, and is reasonably centrally located.
- 9) Locate any critical glare locations, and determine ratio of vertical illuminance at that point to vertical illuminance on the rear glare monitor. The blinds should be set to approximately 75° for this measurement.
- 10) The target glare level can be built in (2,000 lux), but needs to be user modifiable.
- 11) Set averaging time for increasing blind tilt - default 10 -15 seconds - user modifiable
- 12) Set averaging time for decreasing blind tilt - default 10 minutes - user modifiable
- 13) Set averaging time for increasing electric light level - default short - user modifiable.
- 14) Set averaging time for decreasing electric light level - match decreasing blind tilt.
- 15) Measure the illuminance on M1 as a function of power or control level. The number of values needed will depend on whether the system is dimmable, and whether there are banks of lights that can be switched separately. Blind tilt is not critical for this measurement, but medium tilts are preferred. If the measurements are made during the day, the system will need to be able to recognize that there is a non-zero illuminance when there is no electric light, so that it can subtract it out to get the electric light curve.
- 16) System should have the capability to self-calibrate for lumen maintenance by checking the ratio of M1 to the control value or power during night-time operation.
- 17) Determine the location of the workstations where the light level should be controlled. Measure the power to the lights (or light control level) and ratio of the light level at the relevant workstations to the illuminance on M1 with the lights on, and off. The lights should be allowed to warm up for a minute to reach a reasonably stable brightness before measuring the illuminance ratio. The measurements are likely to be most accurate if taken at night, but can be taken during the day if several readings are made to guarantee that a stable value was found. Blind tilt is not critical for the electric light measurements.
- 18) Locations distant from the blinds may need supplemental electric light during daytime hours. Locate the important points and measure the ratio of the horizontal illuminance to the illuminance

on the glare monitor M1. The electric lights should be off, and the blinds should be tilted to 60 to 75 degrees. The controller will need sufficient memory for each point entered.

- 19) Set target horizontal illuminance level. This is the level to be obtained by the electric lights alone. Set at commissioning - allow facility manager to modify.
- 20) Set daylighting float ratio. A value of 1 means that estimated daylight offsets electric light on the workstations. A value of 0.5 means that an estimated daylight level of A, offsets only A/2 electric light. This ratio compensates for inaccuracies in the estimation of daylighting level. Set at commissioning, and allow user to modify. Limit minimum value of the float ratio to 0.25.

C.2 Glare control System

The basic glare control system consists of a motor operable blind (deployment & tilt), illuminance monitoring sensors, and a control system that uses the signal from the monitor sensors to control blind deployment and tilt, and supplemental electric lighting. It is assumed that any significant direct sun penetration will be glaring, and that the blinds will be deployed to block direct sun. Without direct sun, the windows become large area glare sources, and the glare is correlated to the illuminance at the eye (vertical illuminance). The glare control algorithm uses the illuminance on the monitoring sensor as a measure of vertical illuminance, and controls the blind angle to keep the illuminance below a target value(s).

Blind deployment is fairly slow, and disruptive, while blind tilt is much less noticeable. Blind tilt is only effective at reducing direct sun and glare illuminances when the blinds are fully, or almost fully deployed. This suggests that an appropriate algorithm for glare control is to first lower the blinds when direct sun begins to penetrate the window facade, or when sky glare exceeds the target glare value. Once lowered, the blinds can then be tilted to the appropriate tilt angle to block the sun, or control glare.

C.2.1 Blocking direct sun

The glare monitoring sensor is located at or near the back of the room, with the default location being on the back wall facing the window, at a height of six or seven feet. The monitor is capable of giving reasonably accurate estimates of the vertical illuminances in the room as long as there is no direct sun penetration into the room. It can be off by a factor of 10 if there is direct sun. It therefore may not be capable of providing the proper signal to block direct sun.

The IES handbook (5th edition of course) lists vertical illuminances on the window plane, when the sun is not in the plane of the window as ranging from 4,000 to 7,000 lux for a clear sky, and up to about 20,000 lux for a partly cloudy sky. Overcast skies range up to about 8,000 lux. The higher values here are capable of causing glare, but since there is no direct sun for these conditions, the rear monitor should be capable of determining when the blinds need to be lowered (rms error from target illuminance about 20%, worst case error about 40%).

One procedure to determine when direct sun is potentially on the facade, or capable of penetrating the window to a detrimental extent is to compute the sun location, and compare it to the window orientation. This procedure would require a clock, plus a processor to run an algorithm based on the clock to compute the sun location. Commissioning would involve determining the latitude and

longitude of the site, and the angular orientation of the facade. A disadvantage of this procedure is that it will trigger blind deployment even if the sun is blocked by architectural or landscape features, unless the algorithm also contains a database of angles that are obscured. It will also trigger blind deployment when heavy clouds obscure the sun.

It is possible to only deploy the blinds when an exterior vertical sensor exceeds a target illuminance. Unfortunately, if the target value is low enough to detect direct sun, it may also trigger on a bright overcast sky. When the blinds are fully retracted, the maximum glare illuminance in the room is likely to be 25 to 50% of the illuminance on the exterior surface of the window. If we assume a glare target value of 2,000 lux, this limits the exterior value to the range of 4,000 to 10,000 lux. The lower value is below the amount from a bright overcast sky, and is exceeded by the both the partly cloudy and clear skies by the time the solar altitude reaches 5°, except when the sun-window azimuth exceeds about 60° for a partly cloudy sky, in which case it can almost reach 10°.

It is not clear that early deployment of the blind is that serious a problem. If the blind is deployed, but open, there is still a fairly good view, and only a minor reduction in light. In fact, because blind deployment is noisy and noticeable, it may be advantageous to simplify the deployment algorithm. For example, if the window orientation is 5 -10° south of west on up to 45° north of west (in the northern hemisphere) it might be preferable to let the glare monitor determine whether to lower the blinds if there is excessive sky glare in the morning, while the sun is out of the plane of the window, and, then, if they are not already lowered, lower them near noon (solar time), when the occupant is ready to go to lunch, and is therefore less likely to be bothered by their operation. Once deployed, the blinds would then be left deployed until the end of the work day. For these orientations, when the blinds are up in the morning the sun is either out of the plane to the windows, or at an extreme angle so sun penetration is low. Under these conditions the rear glare monitor can be counted on to properly lower the blinds and adjust the blind tilt if sky glare is excessive. In the afternoon, the sun is in the plane of the windows, and the issue will be that of determining the proper tilt of the blinds to block direct sun.

For the south orientation, to 45° west of south, and 80° east of south, the simplified algorithm would be to deploy the blinds in the morning at the start of the business day, and retract them in the evening at the end of the business day. Glare control during the day is managed by blind tilt. If the windows are more than 45° to the west of south, it might be worthwhile to leave the blinds up until about 11AM during the summer solstice period, but they would still need to be lowered by 9AM for the rest of the year. These are solar times, so if daylight savings time is in force, this still comes in at noon. If the window is oriented east of south, daylight savings time pushes the time that the blind can be raised during the summer to 2PM, which may not be considered an auspicious time. For an eastern orientation, the blinds can be lowered at the beginning of the workday, and possibly raised at noon or 1 PM. Alternatively, it may be better to simply leave them deployed at 0° tilt during the work day, as you may not wish to raise the blind on an overcast afternoon, only to have to deploy it again when the clouds finally clear away. It may be better for an east orientation to operate the blinds in the same manner that they are operated in the south orientation. For the northern orientations, direct sun is only a problem during the summer months in the early morning or late afternoon (way late with daylight savings time). Sky glare is also less of a problem, and blinds may only be needed if there is reflected

glare. This would most likely be a problem in the morning hours, so even here it may be possible to use a simplified deployment algorithm.

The simplified deployment algorithm limits the number of times that there is major blind motion, in order to reduce the distraction and annoyance caused by the motion and noise of the deployment. In the simplified case the control hardware has to include a clock. Commissioning would consist of entering one or more deployment and retraction times. This should be no more complicated than entering a set-back time on a thermostat. Note, however, that the clock might need an option to disable daylight savings time. If it is a west facing facade, switching to daylight savings time can make a difference of almost 70° in the sun azimuth in the southern-most portion of the United States. An occupant might wish to delay the deployment towards the end of the lunch period in this case.

The disadvantage of the fully simplified algorithm is that it may save less energy or result in less view. For an east facing facade, an alternative to simply leaving the blinds deployed in the afternoon because the sun might come out, is to set a lower limit of 5,000 lux on the window, below which the blinds are raised. This is a value for a fully overcast sky. It is unlikely that the glare monitor would need to lower or tilt the blinds when the illuminance on the window is this low, or lower. Partly cloudy conditions which might cause the illuminance on the window plane to fluctuate rapidly and significantly when a cloud obscures the sun should not be a problem when the sun is not in the plane of the window.

A second alternative is to allow for a user defined intermediate blind deployment position. Our algorithm for computing glare assumes an eye height for a seated person. If a person has a sit/stand workstation it might be quite useful to have a deployment stop at standing eye height. This would not prevent direct sun penetration, so the glare monitor would not be able to accurately predict the glare potential. In this case control has to be taken by the external sensor. Instead of using 10,000 lux as a signal for full deployment, or retraction, this level now becomes the signal for partial deployment. In the evening, the 5,000 lux level becomes the signal for full retraction. In the morning, a second higher user defined level (default 20,000 lux) becomes the signal for full deployment and tilt control. Tilt during the hours where there is direct sun would be controlled to block sun penetration through the blind.

Once the blinds are deployed, there is still the issue of making sure that blind tilt is controlled to block direct sun. The solar angle that is blocked by the blinds as a function of tilt can be estimated by assuming that the slats are flat. If the slat width is equal to the spacing between the slats, the profile angle is equal to 45° minus $\theta/2$, where θ is the tilt angle. If the ratio of the slat width to the slat spacing is r , then the profile angle, ϕ , is given by the expression:

$$\phi = \frac{1 - r \cdot \sin \theta}{r \cdot \cos \theta}$$

The solar altitude, β , is related to the profile angle through the expression:

$$\tan \beta = \frac{\tan \phi}{\cos \alpha}$$

where α is the azimuthal angle between the sun and the plane of the window.

At a zero degree tilt, the maximum solar altitude that would penetrate the space is 45° at normal incidence, but only 15° at an azimuth of 75° ($r = 1$). If the blind tilt is raised to 15°, which was the default minimum in many of our simulation runs, the blocking angles become 37.5° and 11° respectively. Real blinds will tend to have r slightly greater than one so that the blinds can shut all the way when fully tilted. This will give somewhat lower values than listed above.

If the control system has the capability of determining the sun location these blind tilts can be calculated. This would ensure blocking when needed, but as mentioned earlier, it would also lead to tilting under overcast conditions or when there are obstructions. One potential solution to maximize the open-ness of the blinds is to use an exterior illuminance sensor to determine if the blinds need to be tilted to block the sun.

The simplest algorithm would be to set a target value for the vertical illuminance on the window. If the value is exceeded, and the solar location algorithm indicates that there could be direct sun on the window, the blinds are tilted to block it. If we assume that people can avoid looking directly at the sun itself, then our glare criteria is still on the order of 2,000 lux at the eye. If we also assume that the observer is at least 2.5 feet from the window, then we can look to see what the maximum ratio was between the glare illuminance and the exterior illuminance. For blind tilt angles of 0° and above, this ratio appears to be about 0.2. This means that an exterior vertical illuminance must be 10,000 lux or higher before it is likely to cause glare.

This target value is above the vertical illuminance from a CIE overcast sky, so it will not change the tilt for an overcast sky. It is below the vertical illuminance level for a partly cloudy sky when the solar altitude is at or above 10° when the sun is facing the window, rising to 20° when the relative azimuth is 85°. At these low angles the sun contributes a relatively small fraction of the vertical illuminance, and the glare monitor would probably work fine, however when the sun-window azimuth is below 80 - 85° the direct sun contribution rises very rapidly with solar altitude, so the angular band where the monitor by itself would be fine is small. At the higher azimuths the solar contribution under partly cloudy conditions is very low at all solar altitudes, so this would be a region where the blinds may block the potential sun penetration more than is needed. For clear sky conditions the only time target illuminance would not be exceeded is when the solar altitude is under 5°, or the sun to window azimuth exceeds 80°.

C.2.1.1 Hardware sensor

If finer control than that offered by the algorithm above is desired, it should be possible to put a small radiation/illuminance sensor on the upper edge of a blind slat. If there is direct sun, the sensor signal should drop abruptly when the blind is tilted to block the sun. This technique unfortunately requires that the blinds be tilted to check if there is direct sun. The technique would therefore require some "hunting" of the blind tilt in the case that there is no direct sun. I have not modeled the magnitude of the change that would indicate direct sun, but one can presumably set a target of 2X or more, as otherwise the direct sun is too dilute to really matter.

C.2.2 Controlling sky glare:

The results of our many simulations show that when the direct sun is blocked, a rear wall mounted window facing sensor can provide good estimates of the glare illuminances as a function of the distance from the window. We have a basic functional form, and a limited number of free parameters. The rear sensor can be calibrated either directly to the best fitting theoretical form, or it can be commissioned by a set of direct measurements.

The advantage of commissioning from measurements is that it removes the variability associated with differences in room shape, size, finish, and furnishings. But this is an advantage only if the measurements provide a reasonably accurate procedure.

If the control system is installed in a room where there is a known worst glare location that needs to be controlled then all that is needed is the ratio between that location and the glare monitor. The variability of the ratio at any given point is lower than the variability between even the best fit and the points in general throughout the space. Reasonably accurate commissioning requires that there be no direct sun, which is most easily met by fully deploying the blinds, and setting them to 75° tilt. The highest potential glare is for points closest to the windows, and centered between the side walls. A check of the two closest centered locations gave a maximum to minimum ratio of 1.1 to 1.2 for the variability in the ratios of the glare to monitor illuminances for both the small and large rooms. This suggests that the maximum error for a direct commissioning measurement would be 20%. If we assume that the target glare level is a user modifiable parameter, then this error becomes irrelevant.

In the case above, only the vertical illuminance ratio need be known. If we want to take advantage of the functional form for illuminance as a function of distance we would need to also have several dimensional and location measurements as well. The functional form for the illuminance is based on cavity ratios, so it requires that control system knows the width, depth, and height of the room, as well as the distances from the window of the glare location in question. If the glare location is within three feet of the side wall that information can be input as well, although this just changes the scale factor, and therefore can be accounted for by adjusting the target glare level, as mentioned above.

The functional form has three constants, K_1 , K_2 , and K_3 . In our simulations we found that we could get good fits over many different conditions if we fit K_1 , and fixed K_2 and K_3 . This was particularly true for all the small room fits. If only K_1 is free then there is only one unknown, and the value of K_1 can be determined directly from the measured ratio of the glare illuminance, E to the monitor value, $M1$:

$$K_1 = \left(\frac{E}{M1} \right) \cdot (1 + (K_2 \cdot CR + K_3) \cdot CR)$$
$$K_2 = 0.01996$$
$$K_3 = 0.04925$$

Since K_1 is directly proportional to the illuminance ratio, the uncertainty of its determination is identical to uncertainty in the ratio itself. As mentioned earlier, the maximum uncertainty appears to be about 20%.

For the large room, fixing K_2 and K_3 gave good fits, but there was a distinct improvement in the fits if K_3 was allowed to vary, or if the room was partitioned into zones with fixed K_2 and K_3 , but different values of K_1 . Partitioning is less effective than letting K_3 float, but still made a very significant difference in the large rooms. Relative to the partitioned fit, floating K_3 results in illuminances that are higher in the front of the room, and then lower until the cavity ratio hits 11 - 16, where they once again becomes larger. The room is too shallow to show the increase in the K_3 floated estimates relative to the partitioned estimates. The K_3 floated estimates gives significantly better results in this room. In the large deep room the K_3 estimates do begin to rise relative to the partitioned fit estimates, and the difference in the accuracy of the two fits is slight.

From a theoretical point of view there is justification for both types of fits. The small room, which does not show a need for either of these changes to the default form, has its workstations along the side walls, while the large rooms have them across the floor so that there is significant restriction of light flow from the window. The first glare location is in front of these workstations, while the remaining locations are even with or behind at least one row of workstations. This provides a good justification for the partitioned fit, but there is also reason to believe the workstations will affect the effective cavity ratio, so there is support for a modified K_3 model, or even a combination model utilizing both factors.

Both fits are quite good, so from a practical point of view the question is which approach is easiest and most accurate to implement. The answer here comes down unequivocally in favor of the partitioned fits. If the control algorithm is set up to recognize a partitioned fit function, then commissioning simply consists of measuring the glare to monitor ratio at two distances. The closest distance covers all values closer to the window. The farthest distance covers all values farther from the window. For points in between, the function interpolates the value of K_1 . The uncertainty in the two values of K_1 is again solely determined by the uncertainties in the glare ratio measurements, so this procedure is very stable.

One could, in theory, use two measured glare ratios to fit a K_1 and a K_3 . This would not be a good idea, as the value for K_3 turns out to be very sensitive to the uncertainties in the measurements. For the small room, where there are over 30 sky condition results with a blind angle of 75°, the value of K_3 varied from 0.000015 to 0.15, a thousand to one variation. The partitioned K_1 values of 7% and 18% for the closer and more distant K_1 values. We do not have a theoretical formula for K_3 , so we are left with no practical method of taking advantage of its sometimes better predictive power. For the large room, which is the worst case, the best fit had maximum fit errors of 18%, while the partitioned fit had maximum fit errors of 32%. These latter values are still relatively small in terms of the overall control function.

C.2.2.1 *Blind tilt algorithm*

The value of M_1 lets us predict the vertical glare facing the window, but we still need to actually control the glare by tilting the blind. Depending upon the speed of the sensor-control network, this can be done either as a direct feedback loop, or it can be done by a default prediction and then correction. If the blind angle control only has a set of discrete deployment ratios and tilt angles, then a third alternative is to simply increment by one step whenever glare exceeds the target value, and decrement by one step when it drops a sufficient amount below the target value.

Our simulations show that for a fixed sky condition $M1$ follows approximately the following relationship:

$$M1 = M10 \cdot \exp((4 \cdot z - 1) \cdot \cos \theta)$$

where $M10$ is the $M1$ illuminance level when the blind is at zero degrees tilt, z is the blind deployment ratio, with zero being not deployed and one being fully deployed, and θ being the blind angle, with zero degrees indicating that the slats are horizontal and the slats are open, and 90 degrees indicating that the slats are vertical and the blind is closed. The formula is valid for $z \geq 0.5$.

If we set a target glare level, Eg , at a particular cavity ratio, CR , in the room, then we can determine the value of $M1$ that limits glare to the value Eg :

$$M1t = Eg/k$$

, where $k = K_1/(1 + (K_2 \cdot CR + K_3) \cdot CR)$, where $M1t$ is the target value for $M1$.

When the value of $M1$ exceeds $M1t$, the default blocking blind angle is calculated as:

$$\cos \theta = \frac{\ln(M1t/M1 + (4 \cdot z_1 - 1) \cdot \cos \theta_1)}{4 \cdot z_1 - 1}$$

where θ_1 and z_1 are the current blind angle and deployment.

If the blind is not fully deployed and the calculated blind angle exceeds 90°, then the blind deployment has to be incremented to the next step, and the calculation is redone with the new blind deployment. If the blind is fully deployed and the calculation still indicates an angle greater than 90°, then the blind just goes to 90°. The above algorithm is based on the default sensitivity, S . The sensitivity value can be updated by comparing the value of $M1$ after the change in angle with the predicted value. The sensitivity is $S = A \cdot z - B$, with, with the default values for A and B being 4 and 1 respectively. If only the blind angle is changed, then there is no information on the individual values of A and B , so just the value of S is updated. In addition, there must be at least two changes if both unknowns are to be solved for separately. If the initial setting of the blind to block the sun results in $M1 > M1t$, and the default calculation indicates that the blind angle must be reduced to meet the target glare level, then the default S has to be in error. In this situation the blind angle must be incremented and a new value of $M1$ measured in order to get sufficient information to calculate a new S and the needed blind angle to limit glare. The simulations suggest that S should not vary by more than ± 30% from the default value. Values larger than this may indicate a problem with the hardware.

If the control system changes blind angle in discrete steps, then the function of the default blind algorithm is to determine whether it is necessary to move the blinds more than one step. This may happen under partly cloudy conditions. It is not as likely to happen under clear or overcast skies, as the change in illuminance should be smooth. Changes of more than one step are more likely to occur if the steps are small, or unevenly spaced. With regards to the latter, it should be noted that the relationship between $M1$ and the blind angle in our simulations was that of an exponential of the cosine of the blind angle. If the control system has blind angle steps that are equal, then the relative changes in light level

will be unequal. For example, if the blind increments in 10° steps, then the relative changes in light level are 5%, 14%, 25%, on up to 68% for the last step. To make the ratios of the steps equal the blind angle steps have to be evenly spaced in the cosine of the angle: 0° , 27° , 39° , 48° , 56° , 64° , 71° , 77° , 84° , and 90° (40% change for each step).

There is one other operational/commissioning issue associated with the operation of the blinds, and that is the speed of operation, or time delays. The maximum speed is going to be determined by the hardware. Time delays are an explicit part of the control algorithm. There is a trade-off between controlling the light level, and irritating the occupant with too frequent motion of the blind. We do not have data on this user satisfaction aspect of blind operation. We can set defaults based on intuition, and we can let the user modify the default time delays.

We expect that the user will not want to tolerate glare, so the time delay to tilt the blinds to block glare should be short. We probably don't want it to be so short that the blinds respond to a momentary reflection, but then again, if the reflections are repetitive and frequent (reflections off vehicle windshields), we might want to block it. This analysis suggests that the on trigger for increasing blind tilt be a rolling average over a fairly short time in the range of seconds to a few minutes at most. As a guess, we can set a default of 10 - 15 seconds.

Partly cloudy conditions can produce significant shifts in the glare level. The reductions in glare are not the problem. It is the reduction in the overall illumination that may be problematic. These changes will be largest at the front of the room, where there may be little need for electric light when the daylighting is strong. This suggests that the immediate response to a sudden dip in the daylighting should be an increase in the electric light, followed by a slower response where the blinds begin to open and the lights begin to dim. Any time delay that is chosen represents a trade-off between energy saved (dimming the lights) and the intrusiveness of the change in the blind angle (size and speed of change in the light level, and noisiness of change in volume and time. With ten blind angle steps, the average change in illumination from the blinds is 40%, which is noticeable. The impact of this change can be spread out over several seconds, but this is still likely to be noticed. This suggests that the trigger for decreasing blind tilt should be a rolling average of M1 over time range of minutes, and not seconds. As a guess we can set the default time at 10 minutes, and provide a user variable time from one minute to 20 minutes.

C.3 Commissioning the electric light contribution:

The contribution of the electric lighting to the illuminance on the monitor is expected to depend upon the type of lighting, the location of the fixtures and the general layout of the room. We have shown, however, that the contribution is relatively insensitive to the status of the blinds, particularly for lighting near the back of the room. Our worst case comparison was for just the row of lights in the front of the room lit, with the blinds either fully retracted, or fully deployed and with a blind tilt of 75° . The ratio of the monitor illuminances for these two conditions was 24% for our pendant lighting layout, and 39% for our panel light layout. Changing the blind tilt to 45° drops these ratios to 8 and 12 percent respectively. If the rear row of lights are on, the ratio drops to under 4 percent.

We do not expect high electric lighting levels in the front of the room while the blinds are fully deployed and blocking sun glare, as this is a condition where there should be significant amounts of light in the front of the room from daylight. It therefore seems safe to suggest that commissioning consist of measurements with the blinds at moderate tilt. The most accurate procedure would be to do the measurements at night, when the lights can be left on long enough to stabilize. This procedure could be automatically built in, with the system simply checking the increase (or decrease) of the monitor value any time the lights are switched during non-daylight hours. This capability would be useful for maintaining accuracy in glare control as the system ages, and fixture efficacy declines.

If the lights have a dimming capability, the controller will need extra memory to be able to store the dimming curves. The commissioning algorithm will also need enough intelligence to determine if and where there is hysteresis in the light versus power dimming curve.

For the initial setup it may be convenient to do the commissioning during the day by taking two or three measurements in succession, capturing both the increase in light when the lights are first turned on, and the decrease in light when they are turned off.

If there is more than one lighting condition possible (front row on, back row on, individual lights on or off) then the controller will need to have sufficient input and memory for each condition.

C.4 Commissioning the work plane illuminance

The blind condition has little effect on work plane illuminance from the electric lights. In our simulations, if only the rear row of lights was on, blind operation made a maximum of a 40% difference in the light level on the work place locations near the front of the room. However, it only did so for locations which were distant and shielded from the lights, and thus only had light levels of 10 to 30 lux. Locations where there was a significant of light showed variations of less than 10% with blind condition. The blind condition can be ignored during the commissioning of the work plane illuminance from electric lighting.

Commissioning of the work plane light levels from the electric lighting will require that work place illuminances are measured at those locations for which control is desired. The control algorithm will need sufficient memory for each of them. The dimming curve, if needed, will presumably be the same as measured for the glare monitor.

Daylight does not affect work place illuminance in the same manner that it affects vertical illuminance. Prediction of the workplace illuminance from the glare monitor illuminance is accurate to only a factor of two. If the glare monitor is the only indoor sensor available, a design decision has to be made as to what type of errors are acceptable.

With normal blinds, illumination near the windows will be higher than farther away. With workplane illuminance discomfort glare is generally not the issue, and instead the goal is to get sufficient minimum levels. When the glare is controlled by the blinds, the control system needs to know whether the workplane illuminance from daylight is sufficient, or whether supplemental electric light has to be added. The lowest daylight illumination occurs at the back of the room when the blinds are fully or nearly fully tilted. In the large room, restricting the location and blind angle improves the predictability of

workplace illuminance using just the glare sensor to a maximum error of $\pm 40\%$, and a standard deviation of 15 to 17% (75° and 60° blind tilt respectively). This is an acceptable error range. Commissioning is similar to what is done for glare. The simulations showed that the default form used for glare predicts the average fall-off with distance. Commissioning therefore involves shutting the blind to 60 to 75 degree tilt, and measuring the ratio between the relevant workplace illuminance and glare monitor reading. A single reading is sufficient to determine K1 for the workplace illuminances.

With the system commissioned in this manner, the system will be able to provide sufficient electric lighting at the rear of the room when the blinds are tilted. On the other hand, it will, on average, overestimate the illumination by about 60% when the blind tilt is only 15°. This is a less serious error than overestimating when the blind tilt is 75°, as the workplace illumination at 15° is typically 5 - 6 times higher than when the blinds are at 75°. I have identified four options for dealing with this problem of overestimation:

- 1) Ignore it. The light levels are generally higher under low tilt conditions anyway. At 15° blind tilt the workplace illumination will generally be 300 lux or more anyway, and it doesn't matter if it is overestimated, since no additional light is needed.
- 2) Float the target workplace illuminance level. For most people, higher illumination levels are actually an amenity. We do not provide them because the cost of providing them is more than the benefit. If some of the light is provided by daylight, it changes the cost, and raises the optimal level. If T is the target level to be provided by the electric lighting when there is no daylight, set the target level when there is daylight to $T - xD$, where $0 < x \leq 1$, and D is the estimated daylight level. A value of x of about 0.5 will generally assure sufficient light levels when the blind tilt is 15°, and yet will still save significant amount of energy (cost) at higher blind tilts.
- 3) Provide task lighting. This is basically the same as ignore it, except that the occupant is provided with an option to manually increase the light level if they feel they need it. Task lighting is a nice amenity even if another option is taken to try to control the light level.
- 4) Put in an explicit blind tilt correction. The problem with this approach is that the correction factor does not appear to have a simple functional form. I have not attempted to fit the data at this time.

The above discussion assumes that we only have the rear monitor sensor available for estimating the workplace illuminance levels. We have shown that the workplace illumination levels can be fit more accurately if we have data from both the rear monitor sensor and a downward facing ceiling mounted sensor located near the front of the room. The signals from the two sensors, $M1$ (the rear glare sensor) and $M2$ (the ceiling sensor), are combined to give a joint signal MJ , where $MJ = M1^a \cdot M2^{(1-a)}$. In the large room the best value of a was 0.6, while in the large deep room the best value was about 0.7. Deviations of ± 0.1 in the value of a had a fairly minor effect on the degree of fit. Deviations of ± 0.2 approximately double the variance of the fit.

The idea of using both sensors to control the workplace illumination level is appealing, but there are a number of practical problems with this approach. The first problem is simply that the degree of

improvement is mixed. In the large room, the approach leads to fits with a maximum error of 40%, but in the large deep room the error is over 65% ($\log 0.52$). A second issue is that the values of the exponent a do differ for the two rooms we simulated. Although the difference is small, we do not have data on other room geometries that proves that the exponent would remain in this same range. I also have not evaluated the sensitivity of the exponent to the location of the ceiling sensor. This is something which could be modeled, although it might take some more simulation runs. It would perhaps lead to a procedure for determining the exponent, but I do not have a procedure at this time. We could try determining the exponent by fitting it against measurements at two blind tilt angles, but unfortunately the exponent is sufficiently sensitive to the measured values to make it unreliable. I therefore suggest that this is a procedure which would take further development before we could guarantee that it would work. I suggest that it is not practical at this time.

Appendix D Return on investment calculations

Item	Unit	Retail Price	Unit	Total Unit	Total Unit Cost after Discount	Labor Hour per Unit	Total Labor Hour	Total Cost
LED Luminaire	EA (8'x10' spacing)	\$200.00	274	\$38,360.00	0.33	91.33		\$42,926.67
Blind	Area (m ²)	\$61.90	652.8	\$28,285.82	0.00	0.00		\$28,285.82
Blind Motor	EA/10 Linear Feet	\$305.00	168	\$35,868.00	0.50	84.00		\$40,068.00
Dynalite Multisensor	EA/Zone	\$53.99	60	\$2,267.58	0.17	10.00		\$2,767.58
Glare Photosensor	EA/Zone	\$20.00	60	\$840.00	0.17	10.00		\$1,340.00
Dynalite Blind Driver	EA/Side/Floor	\$222.92	12	\$2,140.03	0.50	6.00		\$2,440.03
Anemometer	EA/Building	\$117.00	1	\$117.00	1.00	1.00		\$167.00
Outdoor Sensor Box	EA/Façade	\$200.00	4	\$800.00	1.00	4.00		\$1,000.00
Temperature Sensor	EA/Zone	\$28.95	60	\$1,215.90	0.17	10.00		\$1,715.90
ECoMIC Zone Controller	EA/Zone	\$200.00	60	\$8,400.00	0.33	20.00		\$9,400.00
Infrastructure & misc.	EA/Building	\$10,000.00	1	\$10,000.00	3.00	3.00		\$10,150.00
Total				\$128,294.34		239.33		\$140,261.00

Note: The retail prices and discounts are only rough estimations. In many cases, the prices and discounts are proprietary information specific to each project.

Figure D-1 Retrofitting costs for the perimeter zone in the DOE reference medium office building.

	Perimeter Lighting Energy (GJ)	Savings (%)	Energy (GJ)	HVAC Savings (%)	Charge	Total Energy Saving	Total Demand Charge	Total Demand Saving
Before Retrofit	388.02		1003.47		\$52,842.25		\$18,375.92	
After Retrofit	163.26	57.92%	999.62	0.38%	\$49,498.35	\$3,343.90	\$17,545.49	\$830.43
After ECoMIC	65.30	83.17%	299.89	70.12%	\$36,942.10	\$15,900.15	\$13,094.73	\$5,281.19

Figure D-2 Energy and demand charges estimation.