

An Immersive Visualization Framework for Monitoring, Simulating and Controlling Smart Street Lighting Networks

Gianluca Paravati
Politecnico di Torino
Dipartimento di Automatica e
Informatica
Corso Duca degli Abruzzi, 24 -
10129 Torino, ITALY
gianluca.paravati@polito.it

Fabrizio Lamberti
Politecnico di Torino
Dipartimento di Automatica e
Informatica
Corso Duca degli Abruzzi, 24 -
10129 Torino, ITALY
fabrizio.lamberti@polito.it

Andrea Sanna
Politecnico di Torino
Dipartimento di Automatica e
Informatica
Corso Duca degli Abruzzi, 24 -
10129 Torino, ITALY
andrea.sanna@polito.it

Eduardo Henao Ramirez
Politecnico di Torino
Dipartimento di Automatica e
Informatica
Corso Duca degli Abruzzi, 24 -
10129 Torino, ITALY
eduardo.henaoramirez@
polito.it

Claudio Demartini
Politecnico di Torino
Dipartimento di Automatica e
Informatica
Corso Duca degli Abruzzi, 24 -
10129 Torino, ITALY
claudio.demartini@
polito.it

ABSTRACT

In recent years there has been an increasing effort in investigating integrated solutions to reduce energy consumption levels in different domains. In this perspective, tools thought as a support for the simulation and experimentation of energy saving strategies could constitute a valid help both in the design and validation phases.

Within the above scenario, this paper presents a framework for monitoring, simulating and controlling energy-aware street lighting systems capable of optimizing their energy consumption levels based on environmental conditions and established lighting policies.

The devised system allows operators to visually inspect (and control) the power saving impact of specific dimming rules in both monitored and simulated environments through the use of advanced 3D vision and interaction capabilities. In fact, a virtual interface is able to recreate the lighting effect within digital city environments, which can be explored by using immersive stereoscopic techniques. Moreover, natural user interfaces based on the recognition of user gestures and poses and voice commands are exploited to improve virtual navigation of the lighting infrastructures.

By addressing all the presented aspects, the proposed system aims at constituting an effective tool to develop and analyze energy saving solutions in the domain of (but not necessarily limited to) future public street lighting.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Simutools 2012, March 19-23, Desenzano del Garda, Italy
Copyright © 2012 ICST 978-1-936968-47-3
DOI 10.4108/icst.simutools.2012.247739

Categories and Subject Descriptors

H.5 [Information Interfaces and Presentation]: Miscellaneous

Keywords

Smart Street Lighting, 3D Visualization, Natural User Interfaces

1. INTRODUCTION

Today, advanced controlled networking architectures are becoming one of the key enabling technologies to design effective solutions for a wide range of applications, ranging from environmental sensing for natural disaster prevention to traffic control and sustainable resource management.

With a particular reference to the latter scenario, it is considered to paramount importance to develop new urban energy saving policies, e.g., considering that most of the European energy consumption takes place in cities. In this respect, the integration of sensor networks and the design of smart cities could constitute a valid help in achieving this goal; in fact, the creation of distributed networks of intelligent sensor nodes, which can measure many interesting parameters, could lead to more efficient management strategies tailored to different urban aspects. Indeed, the rationalization of public street lighting systems would play a significant role in the above scenario.

The design of system-wide intelligence and control of networked street lighting systems could allow to make informed choices about the control of the instantaneous power wasted by diverse sets of luminaries. Moreover, public street lighting should provide users with much more than a feeling of “comfort”. Specific photometric performances have to be ensured in order to allow different users accomplish multiple visual tasks.

In this work, the outcome of the activities carried out within the LITES “Led-based intelligent street lighting for

energy saving” project are presented. In this project, a technology for designing lighting networks composed by smart LED (Light Emitting Diode) luminaries has been developed. Besides being characterized by higher durability and improved quality of light, LEDs are fully gradable in light, with linear consumption and instant time response; moreover, the smart luminaries are rigged out with on-board sensors. The above aspects allow to create a responsive system with increased efficiency with dimming at the individual luminary. Moreover, other key benefits are provided by remote monitoring of status and consumption data in real time, which is expected to ensure lower maintenance costs.

The aim of the proposed system is to simulate, monitor, display, and supporting the analysis of the data generated during the functioning of an intelligent street lighting system. The operator is provided with a 3D representation of the city environment of interest, which shows both real time and recorded lighting results in the context of virtual reconstructions of simulated and on-site networks of luminaries. By virtually exploring the reconstructed 3D virtual environment, the operator can issue commands to configure the overall behavior of the system. In particular, specific rules allow the operator to control the reaction of a given set of luminaries to particular stimuli (motion, changes in ambient light conditions and so forth). Furthermore, during simulation, additional rules can be designed to govern the behavior of the stimuli themselves, thus letting the operator create synthetic environments suitable for experimenting different intelligent lighting strategies. The LITES project considers the validation of these dimming policies, from the user/perceptive point of view, as a task to be accomplished by the pilot sites, which are in charge to install, maintain and test the overall system performance. Onsite tests and user feedbacks will be used to validate dimming rules and to gather statistical information to refine models that generate input events for the simulator.

The benefits of 3D visualization within this scenario concern primarily the possibility to create synthetic environments for the simulation of light control rules and consequently explore a virtual scenario of the expected lighting effects. Moreover, this technology can be used in prospect to automatically analyze the recreated 3D environments to assess quality metrics of the perceived light in order to find a tradeoff between energy consumption levels and, e.g., safety. In this respect, the framework proposed in this paper is only a first step of a more ambitious program that aims to develop a simulation tool able to manage both dimming rules and realistic rendering. In order to allow users to experience realistic simulations, 3D representations and natural user interfaces are necessary technologies.

However, at this stage of the LITES project, the aim of the proposed simulation tool is not focused on the realism of the scenario. The goal of the solution proposed in this paper is to provide a real-time tool, running also on machines with limited hardware capabilities, to allow system administrators to set the dimming policies and check corresponding power consumption. For this reason, indirect illumination, shadows, reflections and refractions are not currently managed.

Figure 1 shows an overview of the main functionalities of the integrated visualization and man-machine interaction subsystems. The management software is in charge to receive and translate user commands into meaningful actions

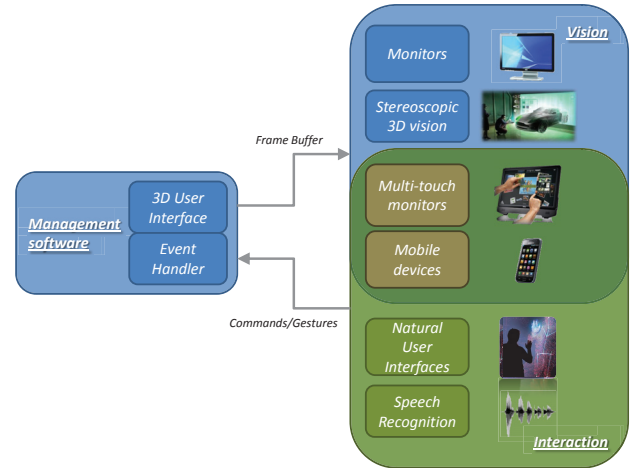


Figure 1: Overview of the main functionalities of the man-machine interaction subsystem providing multi-modal access with heterogeneous devices and control through man-machine interfaces alternative to the traditional mouse and keyboard based ones.

to be applied to the 3D scenario. An event handler module receives commands and gestures from the interaction subsystem, which supports different man-machine interaction paradigms, from vocal commands to Natural User Interfaces (NUIs). Thus, user commands can be interpreted to navigate within the digital city during the simulation, monitoring and control phases. Moreover, the system can also be configured to be used with handheld terminals (e.g., tablet PCs, smartphones and so forth), thus giving the user an ubiquitous access to the 3D visualization system. In particular, remote rendering techniques are exploited in order to provide mobile terminals with the capability to display complex 3D environments. In the context of 3D visualization, a 3D stereoscopic module gives the user the possibility to experience an immersive navigation through the digital city environments. In the following, a generic description of the framework will be given, with a particular emphasis on the simulation tool as it constitutes a peculiar component of the devised system. Indeed, the simulation and the analysis of the data generated by using different intelligent lighting solutions could represent an effective means for the design of energy-saving strategies, leading both to economic and ecological benefits. Moreover, the devised man-machine interaction subsystem will be described in details.

The rest of this paper is organized as follows. Section 2 overviews the related work reported in the literature. Section 3 presents the overall system architecture of the intelligent street lighting solution, whereas Section 4 provides an in-depth description of the simulation tool. The man-machine interaction subsystem is described in Section 5. Simulation tests and results are discussed in 6. Finally, conclusion and future works are addressed in Section 7.

2. RELATED WORK

Controlled sensor networks architectures have been proposed both in the framework of research projects and commercial solution within different areas. For example, the

OpenEnergySim project [14] proposes a platform aimed at studying the effects of intelligent transport system strategies on driving behavior; in particular, the main object of interest is pollution reduction in terms of CO₂ emissions. The system includes a networked 3D virtual world where the results of traffic simulations including pedestrians and environmental conditions can be visualized [11]. The transportation related data are managed and monitored through computer graphics also in [12].

Virtual reality has been exploited as a simulation tool for a wide number of applications. As a matter of example, in [10] pipeline network layouts are planned, optimized and simulated in a 3D digital city scenario. Situational awareness virtual reality environments are exploited in [9] to coordinate the work of emergency response teams.

In the specific context of energy-aware systems, besides the LITES research project [4] introduced in Section 1, Echelon Corporation [2] provides commercial implementations of distributed control frameworks in the China market for smart networked street lighting systems. The focus is on energy saving and remote management, and pilot projects have been installed on various chinese cities and provinces. In particular, the lighting solution exploits high pressure sodium luminaries and includes electronic ballasts networked with a signaling technology for distributed control together with ad-hoc servers developed for the enterprise connectivity. However, the system does not include an immersive visualization system capable of helping the operators to control monitoring and simulation data more effectively.

3. SYSTEM ARCHITECTURE

The overall architecture of the LITES system consists of different software and hardware modules, as illustrated in Figure 2.

The street lighting system is constituted by a network of luminaries with embedded local intelligence. Luminaries are furnished with different on board sensors, capable of measuring ambient light intensity and temperature and of identifying motion events. Based on sensor information, an embedded logic is configured to “take a picture” of the overall environmental conditions and to react by setting the proper level of light for the specific luminary.

An electric cabinet/collector gathers and transfers lighting system data from the luminaries to the management software. In particular, data encompassing sensor measurements and luminary status information are sent over a Power Line Communication (PLC) or a radio frequency link to the cabinet; here, a web-based gateway module is responsible for creating a structured description of the overall lighting system status (including notifications and alarms caused by fault-detection modules, e.g., related to broken or not working luminaries) that is then delivered to the communication manager.

The gateway simulator generates all the events needed to simulate the real world lighting scenario. Basically, through the gateway simulator a pilot site administrator is able to set virtual illumination conditions by acting on different luminary parameters, such as position, height, orientation, color temperature and so forth. Moreover, the administrator can specify local lighting rules as well as configure the behavior of synthetic pedestrians and vehicles to simulate energy consumption levels and verify the impact of the designed configuration on comfort and safety.

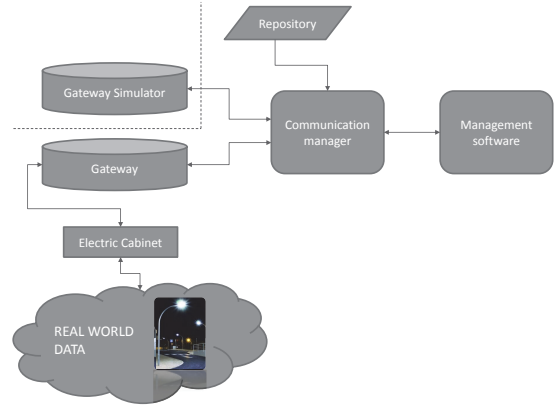


Figure 2: Overall LITES architecture. A communication manager is in charge to both manage real sensor and simulation data.

The communication manager acts as a translator, resolving the correspondences between the sensor data and status coming from the real world (or from the simulator) and a persistent repository. Specifically, the communication manager receives information from the gateways (represented in its own codification and syntax) and translates it to the specific syntax used in the repository. Moreover, it provides intelligence to the luminaries by allowing the administrator to upload the set of rules defined for lighting system.

The management software encompasses all the modules that allows to remotely manage the lighting system, by allowing the administrator to create and apply the rules controlling the intensity of light, to explore a 3D virtual scenario recreating in real time a preview of the expected lighting effect, and to analyze metering information and performance statistics.


Simulation data and real (monitoring) data are stored in the repository. All the events generated during the simulation phase (e.g., power level change events, activation of motion detection sensors and so forth) are recorded and can be used to build statistics. Information concerning the virtual environments representing the pilot sites (including terrains, buildings, streets, pavements and so forth) are also stored in the repository as well.

3.1 Operating modalities


The lighting system management software is designed to operate in three different modalities:

- simulation mode;
- monitoring mode;
- control mode.

As shown in Figure 2, the system is designed to collect data from the actually deployed street lights. In particular, each lamp provides information about its current status (working or not), the ambient light level, and movement detection. The tool has been designed to be used both for lighting design phase and operation phase. A communication component manages incoming information; its behavior is independent of the actual data source (e.g. the on-site lighting infrastructure or the simulator). Indeed, in the



(a) Sample rendering



(b) Building models on the city mesh

Figure 3: Modeling phases of a city district.

current state of the project, the collected parameters are considered in the control mode to allow the application of the control rules not only in the simulation world but also to actually control the deployed street lights. However, the operation phase will be validated when the pilot sites will be active. On the other hand, in the current state of the implementation the collected parameters are not currently used to update the rendering phase.

In the following subsections, the basic features of the operating modalities will be explained, with a particular focus on the simulation functionalities.

3.1.1 Simulation

The simulation modality allows the lighting site administrator to set up virtual illumination conditions by acting on different luminary-related parameters, including position, height, orientation, color temperature and so forth. Moreover, it is possible to study the energy-consumption levels by properly configuring local lighting rules. In order to evaluate the impact of the designed lighting configuration under realistic conditions (e.g., in terms of comfort, safety and energy saving) the site administrator can configure different behaviors for virtual pedestrians and vehicles. Further details regarding the above features will be provided in Section 4.

3.1.2 Monitoring

During monitoring, data gathered by the environmental sensors will be used to set the lighting conditions of the virtual environment in order to mimic the status of the managed scenario. In particular, sensors data will be translated into suitable computer graphics lighting parameters, thus providing operators with a reproduction of the monitored site as much realistic as possible. In order to improve information retrieval capabilities and to enhance the effectiveness of the interface, the 3D scene has been augmented by adding visual clues to depict instantaneous and average power consumption, sensor events, failure notifications, etc. The site administrator is allowed to move around the whole installation by navigating through a 3D stereoscopic system, thus having a clear view of the status of the entire lighting infrastructure.

3.1.3 Control

The 3D reconstruction of the lighting system allows the operator to select specific set of luminaries and set up local lighting rules by defining: working time interval, minimum lighting power, lighting power supplied in response to twilight switch and motion sensor events, etc. The dimming

rules can be easily modified to follow specific public lighting policies or to implement particular cost optimization strategies. Moreover, in the control modality the administrator can instantaneously change the current lighting parameters in terms of light intensity, color temperature, etc. to respond to specific local events (e.g., related to safety issues, weather conditions, etc.).

4. SIMULATION SYSTEM

The simulation tool provides the administrator with the possibility to simulate, on a given lighting site topology, the impact of designed strategies aimed at reducing the energy consumption levels. In order to let him or her visually inspect the effects of the simulation, each site is represented through a 3D graphic visualization module. To this purpose, site topology is first converted into a 3D data format (permanently stored into the repository) starting from cartographic information and digital photographs. Figure 3 shows some particulars of a reconstructed 3D model used in the simulation tool. The basic idea behind the simulation tool is to let the administrator identify characteristic patterns defined by the topology of the site model and by the knowledge of the area, to support the extraction of most probable paths followed by pedestrians and vehicles and to experiment the application of different lighting strategies with the identified patterns. In this context, lighting strategies are coded into dimming rules that will be explained later in this section. Finally, the administrator can verify the effectiveness of the tested rules by analyzing statistics and virtual metering information generated during the simulation.

The main parameters of the simulation system are its duration (in terms of starting and ending dates and time) and its speed (controlled by a slider). During the simulation, each parameter and each rule can be accessed either by selecting the corresponding identifier within a list of luminaries or by directly picking the luminary of interest in the 3D environment. By selecting one or more luminaries, different 3D icons showing the current lighting status together with some key information (e.g., the percentage of output power) are displayed.

4.1 Intelligent logic and simulation items

The intelligent logic represents the core of the simulation environment. The architecture of the intelligent logic is composed by objects (for instance luminaries, pedestrians, vehicles, waypoints and so forth), rules and actions. Objects can be quite complex, since they can be composed by or

equipped with other entities. For instance, luminaries are equipped with different kind of sensors able to react to external stimuli. The information provided by all the on-board sensors allows the system to reason and choose the appropriate action to be taken. The reasoning system is built on the top of a stack of dimming rules that are evaluated for each luminary to decide the best course of actions given certain input conditions.

In the simulation, luminaries obtain stimuli mainly from vehicles and pedestrians. Real pedestrians and vehicles generally move in well-defined and distinct spaces. For the purpose of simulating their behavior, the simulation tool implements a number of waypoints, i.e., geographically distributed entities defining the paths along which pedestrians and vehicles will move. The connections established among these points create two distinct pedestrian and vehicle networks (which, nonetheless, during the simulation could interact to let the system handle pedestrian-vehicle relations).

Each vehicle path has been modeled to be unidirectional, starting from an origin waypoint and ceasing at a destination waypoint. On the other hand, pedestrian paths have been modeled to be bidirectional. Indeed, pedestrians do not have restrictions in their movement, whereas cars can only move forward during the simulation.

Waypoints are classified into three categories: generators, crossings and destroyers. A generator waypoint is in charge to create new instances of pedestrians or vehicles. A crossing waypoint simply serves as a node to define the network of connections for the moving items. Finally, a destroyer waypoint is responsible for destroying the incoming element.

If a waypoint has been configured as a generator, a rule can be added to it in order to control the generation process. In this way, the administrator is allowed, for instance, to simulate different types of traffic during the day. Nevertheless, it is possible to devise ad-hoc generators and attractors to concentrate the generation of pedestrians (or vehicles) along particular paths to simulate highly crowded environments such as public spaces.

Waypoints are symbolized with yellow spheres inside the 3D virtual scene. In order to avoid the overpopulation of the 3D view, their visibility is not activated unless requested by the administrator.

4.2 Rules

In the management software there are specific rules which can be set for controlling both luminary and waypoint behaviors. A rule is a unique entity that has got a unique identifier and specific parameters. In the following, a description of the dimming rules (for the control of luminaries) and the waypoints rules (for the control of the generation processes during the simulation phase) will be provided.

4.2.1 Dimming rules

A dimming rule models the behavior of the single elements of the lighting system in terms of supplied power based on the elaboration of the input events. In general, dimming rules consist of homogeneous blocks taking as input two types of parameters and giving in output a target result. Parameters describe time and sensor constraints. The target result represents the output level of light intensity to be supplied by the interested set of luminaries. Moreover, rule instances can be seen as a unique box that could affect a single luminary or a specific group of them. In the first

Figure 4: Interface of the rule builder.

case, the rule is referred to as a “local rule”, whilst in the latter case it is named “global rule”. Each luminary can be associated with multiple rules.

Local and global rules are put into a stack in order of priority. In general, global rules have higher priority than local ones. In this way, local rules are evaluated only if none of the rules present on the global stack can be activated based on given inputs at specific circumstances.

As already said, the first parameter of a rule concerns the time constraints (day, date, hour, etc.). Indeed, it is possible to create rules that are evaluated at particular time intervals on specific dates or at recurrent days of the week. The day and date options can be set in an exclusive way.

On the other hand, the second parameter concerns the sensor constraints. This parameter takes into account the information coming from real or simulated sensors in order to trigger the decision events of the intelligent logic accordingly. A list of available sensors per luminary allows the administrator to select the desired triggering event for the considered rule. At the present time, two inputs are considered, which are originated by a motion sensor (e.g.,

infrared-based) and by a light sensor (e.g., a photodetector). The motion sensor generates a binary information; as a matter of fact, a rule can be activated based on the presence or absence of movement in the field of view of the sensor itself. Conversely, the light sensor captures the current level of ambient light; as a matter of fact, a rule can be scheduled to react based on the percentage of input light that reaches the sensor. The light sensor has been modeled by using a percentage value, since the output values are sensor-dependent. By setting the desired relation (e.g., equal to, lesser, lesser or equal, greater, greater or equal) between the sensor input and the light percentage value, the administrator can associate a given action to the monitored ambient light.

Finally, as said, a rule returns a target result represented by a percentage power output of the interested set of luminaries and a timer. In this way, the light flux of selected luminaries can be forced to a fraction of the nominal power in particular situations (e.g., during half of the night) and later raised up to a higher level (e.g., when motion is detected). Whenever both the time constraints and the sensor constraints of a rule are satisfied, the power output of a luminary is set to the level specified as a target by the rule itself. This power output level is hold for a time interval pointed by the specified timer value. Figure 4 shows the interface used for creating and managing the rules.

4.2.2 Waypoint rules

A waypoint is endowed with a stack with its own rules, similarly to a luminary. The waypoint rules determine the schedule for the creation of the pedestrian or vehicle instances. Therefore, the main parameter of this kind of rules concerns the time constraints. It is possible to select dates and time intervals to generate the assets. Differently from dimming rules, the time range in this case is not an option, as it is mandatory for the creation process. Moreover, it is possible to specify the type of distribution for the generation process and the number of entities to be generated during the specified time interval. For instance, an “equally distributed” option will allocate the new instances at regular time intervals, whereas a “normal distribution” option will randomly generate new instances by following a normal or gaussian distribution.

4.2.3 Rules evaluation

Indeed, rules can overlap, as they may be defined on the same period of time, on the same input/event, etc. As a matter of example, for two different rules to be activated on Sunday, the time constraints could be set from 5:00AM to 8:00AM for the first rule and from 7:00AM to 10:00AM for the second rule; assuming that the other parameters and incoming data are configured to lead to an activation impulse, the effects of the two rules effectively overlap in the time interval 7:00AM-8:00AM. If the two rules belong to the same group (i.e., local or global rules), a “collision event” occurs.

The rule manager is designed to check the consistency of the time constraints inside the rule stack. Should rules overlap, a warning showing the interested rule identifiers is generated by the system to warn the administrator about possible collision events. However, the administrator can ignore the generated warnings. In this case, at simulation time the rules evaluator will select the satisfied rule with the highest priority within the stack (i.e., the first evaluated

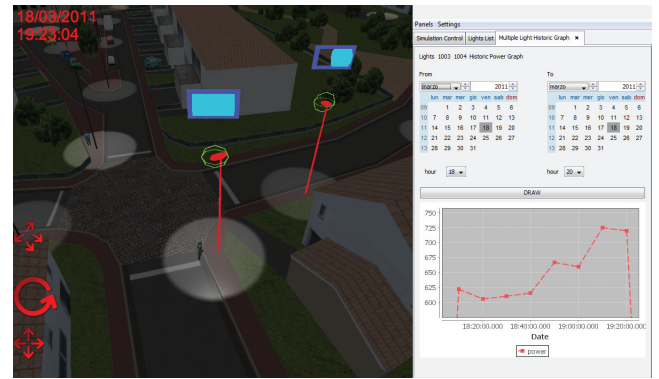


Figure 5: User interface of the management software. The left panel shows a 3D representation of the simulation environment, whereas the right side panels enclose the simulation functionalities.

rule).

On the other hand, another possibility concerns the management of time overlaps among rules belonging to different groups. In this case, a local rule could collide with a global rule. This situation generates an “overriding event”, aimed at favouring the global rule results with respect to local rule ones during the critical time interval. For instance, considering the previous example, if the first rule (5:00AM to 8:00AM) is set to be global and the second one (7:00AM to 10:00AM) is set to be local, the global rule will override the local rule between 7:00AM and 8:00AM.

It is worth observing that, although collisions among rules are undesired events (because they indicate the presence of incongruences), the overriding events are specifically designed so as to homogenize the behavior of a given group of rules at specific working hours.

4.3 Analysis tools

Both during simulation and real time site monitoring, a large amount of data is generated for the overall street lighting system to work. The long-term data are transferred into the persistent repository. Simulation and monitoring results can be analyzed by granting the access to the repository data to the site administrator. Data aggregation strategies supply the statistics and metering tools, whose goal is to provide a graphic representation of the statistical properties of the results obtained by applying a set of rules on a particular street lighting topology. In this way, the administrator could analyze economic and energy impact of specific light configurations or try to identify improved dimming rules by studying recurrent lighting patterns (e.g., associated to particular pedestrians or vehicles traffic generated in particular intervals of time).

For instance, through the plots on filtered data it is possible to analyze the power consumption per unit of time of a given set of luminaries. The administrator can apply different rule configurations on the same dataset (referred to either monitored or simulated events) to perform a series of optimization studies aimed at finding a tradeoff between energy consumption and performance.

It is worth recalling that the frequency of the stimuli on a real luminary heavily depends on its installation area (e.g., industrial or residential) and on the time. For instance, a

set of luminaries installed on the streets of a residential area will provide the highest frequency of stimuli during the early morning and the late afternoon, when people is on the work-home route, thus resulting on a larger energy consumption during these intervals of time.

5. MAN-MACHINE INTERACTION

5.1 User interface

The visualization framework is divided into various modules depending on the device used. In fact, besides allowing for the use of standard displays, the 3D visualization services offered by the designed system permit to explore the street lighting scenario using 3D stereoscopic vision. Moreover, the devised framework supports ad-hoc visualization services for mobile devices (such as smart phones and tablet PCs).

The user interface of the management software is divided into two main components (shown in Figure 5), namely a 3D visualization window and several side panels. Side panel allow the administrator, for instance, to inspect and change the configuration of simulation parameters, to query the database and to obtain operating statistics, etc.

The 3D visualization window reconstructs the managed site topology as a virtual environment within which lighting effects are recreated by using computer graphics techniques. As a matter of example, the virtual navigation allows the administrator to check the impact of luminary failures, control the behavior of the network and study the effects of the fine-tuning of designed local and global lighting rules both in terms of illumination comfort and power consumption. The virtual environment can be explored by using a set of three possible actions, i.e., tilt, rotate and pan. The rotation mode rotates the viewpoint around a fixed centre. The tilt operation leaves the camera position fixed and points it towards the location of interest. The pan mode moves the camera parallel to the screen to obtain an overall translation of the view. Moreover, it is possible to set navigation directly on the street in order to simulate pedestrian viewpoints.

The digital site environments have been modeled starting from Geographic Information System (GIS) data, CAD drawings and photographs taken in place to create the 3D geometries. Figure 3 shows some particulars of the pilot sites reconstruction. The virtual models have been constructed by using different techniques, ranging from polygonal to parametric modeling. The 3D geometries are loaded and displayed by an OpenGL-based application.

Standard OpenGL applications allow a maximum of eight light sources in a scene at once. However, the designed 3D application heavily relies on the concurrent visualization of the effect of multiple luminaries. Hence, ad-hoc shaders have been written to enable the visualization of up to 50 real time concurrent light sources, thus enabling interactive visualization of the particular scenario under control. The designed shaders provide the capability to control the color, position, orientation, inner and outer cone of each light source. In particular, the capability of OpenGL to use custom vertex and pixel shaders to harness the power of the video card has been exploited. More specifically, the shaders have been written by using the ARB Vertex Program and ARB Fragment Program, in order to support the use of the majority of the video cards commercially available.



Figure 6: Technological set-up of the man-machine interaction subsystem. Body poses are captured by a Microsoft Kinect device and stereoscopic support is provided by an NVIDIA 3D Vision kit.

5.2 Natural interaction

The aim of the interaction subsystem is to allow the lighting administrator to explore a 3D virtual scenario of the street lighting system both during simulation and monitoring to check and modify the expected lighting effects. The devised management software makes use of advanced man-machine interfaces that allow for a natural interaction with the system. In fact, the site administrator can freely explore the 3D scenario without wearing any controller. To accomplish this task, the Microsoft Kinect sensors [5] are used and a set of body poses are defined that can be translated into meaningful actions inside the 3D street lighting scenario. For example, navigation tasks can be performed by recognizing poses like lean left, lean right, walking, etc.

User body is tracked by the on-board sensors of the Microsoft Kinect's optical subsystem, composed by an infrared emitter, a depth sensor and a color camera. The infrared emitter projects a pattern of infrared light into the real environment; as the beam hits a surface, the pattern becomes distorted. The depth sensor analyzes the distorted pattern and builds a 3D map of the environment. The Microsoft Kinect is connected to a graphics workstation via a USB link, transmitting data via an uncripted feed. The RGB camera is used to get information about objects and people standing in front of the sensors. Hence, the body movements are tracked in real time and translated to meaningful commands to be passed on to the 3D visualization engine.

The body gesture recognition subsystem is built on the top of the following middleware: FFAST (Flexible Action and Articulated Skeleton Toolkit [3]), OpenNI [6], PrimeSense Nite [7], and the Kinect drivers. FFAST is a framework aimed at making easier the integration of full-body control with virtual reality applications using OpenNI-compliant depth sensors. The framework includes a custom VRPN (Virtual-Reality Peripheral Network [8]) server to stream the user skeleton over a network, allowing virtual reality applications to read the skeletal joints as trackers using any VRPN client. OpenNI interfaces the physical devices to the framework components in order to provide sensor data.

OpenNI covers communication with both low level devices (e.g., Microsoft Kinect), as well as high-level solutions (e.g., FFAST). Body poses are detected by FFAST and used to trigger relative interaction commands on the standard user interface, thus implementing an effective and robust command chain.

5.3 Stereoscopic 3D vision

One of the features that makes the simulation tool more appealing is the stereoscopic 3D support. Figure 6 shows the main components of the stereoscopic vision system. The workstation hosts the management software described in Section 3 and its graphic card supports the Quad Buffered OpenGL Stereo mode. The Quad Buffering technique is used to provide each eye its own image from a slightly different perspective. Basically, the standard double buffering technique (where the part of graphics memory used to display the current frame and the part of memory that is filled with the data for the next frame are alternatively swapped) is used for separating the left and right eye views for a total of four buffers. From an OpenGL point of view, each pair of buffers is swapped at once through hardware page-flipping. The result is a sequence of left perspective image and right perspective image frames.

As previously discussed, an OpenGL-based application loads and renders digital representations that can be used by site administrators to monitor and simulate new dimming rule strategies, thus obtaining an immediate comparison of their impact on the recreated environment. In particular, two cameras are used inside the virtual environment to generate the different views associated with each eye, thus filling the corresponding part of memory in the Quad Buffered mode.

In order to give users a smooth visualization experience, besides reducing flickering issues, a 120 Hz projector is used; furthermore, a pair of LCD shuttering glasses worn by the user blinds alternatively the left and right eyes at the same refresh rate of the projector. Figure 7 shows user interaction with the lighting system during a stereoscopic visualization session. The synchronization between LCD shuttering glasses and the images shown on the projection surface is granted by an infrared emitter connected to the graphics adapter, which is aimed at intercepting buffer swapping operations. The whole process happens so fast that the mind cannot detect the two images melding to produce a full-stereo view.

One of the main problems affecting stereoscopic 3D systems is represented by the negative symptoms causing an uncomfortable feeling during the visualization that could be experienced by the users [1]. This is mainly due to the differences between the functioning of the real human vision system and the way a stereo image pair is viewed. Indeed, two different processes are involved during viewing: focus and convergence. During normal viewing human eyes concurrently converge and focus at the same depth, depending on the location of the object of interest; on the other hand, whenever a stereoscopic image is viewed the focal point is the same for all the objects within the scene and corresponds to the depth of the projection surface, but the convergence is different for each represented object. This discrepancy can cause eyestrain.

To address this issue, some parameters affecting in real time the position and the view direction of the virtual cam-



Figure 7: Navigation of a digital site through body poses. The stereoscopic 3D vision is provided by shutter glasses worn by the user.

eras have been introduced. In particular, the separation distance and the convergence angle can be initially reduced in order to give users a more comfortable setting. Thereafter, these parameters can be modified to enhance perception of depth in the graphics simulation.

5.4 Support for mobile devices

A site administrator could need to access the system from everywhere. To address this need, an important feature has been included in the visualization subsystem, i.e., the possibility to perform the simulation, monitoring and control tasks by following a multi-modal approach, where various classes of devices can be used. Among these, mobile devices such as tablet PCs and smart phones are of main importance to give the user the possibility of an ubiquitous access. One of the main limitations of extending the devised system on mobile devices is the lack of enough resources to run all the tasks. For example, 3D exploration of the simulation or monitoring results cannot be locally performed on some of these devices. To this purpose, the designed system comprises a remote rendering module in charge to offer 3D rendering services tailored to mobile clients.

The remote rendering module makes use of interactive live video streaming techniques to deliver the 3D scene content to mobile terminals. Indeed, in this context, the generation of the virtual administrator perspective depends on the navigation commands given in real time. Since standard video streaming techniques would have as a main drawback high interactions latencies due to buffering timings, ad-hoc solutions have been investigated to reduce at minimum the round-trip delay, i.e., the time lapse between the user inputs (in terms of roto-translational commands) and their visible effect on the terminal screen (the new rendered image of the digital site scenario). In particular, the key concept of the adopted low delay streaming system is to send an image frame immediately after it was rendered on the server side and to display an image frame immediately after receiving it at the client side. Further details about the support for mobile devices can be found in [13].

6. TESTS AND RESULTS

The purpose of this section is to show simple practical

Table 1: Waypoint rules for pedestrian generation

Rule ID	Time interval (hours)	Number of instances	Distribution type
Rule 1	18:00 - 21:00	100	Gaussian
Rule 2	21:00 - 00:00	60	Gaussian
Rule 3	00:00 - 04:00	15	Homogeneous
Rule 4	04:00 - 06:00	75	Gaussian

Table 2: Dimming rules

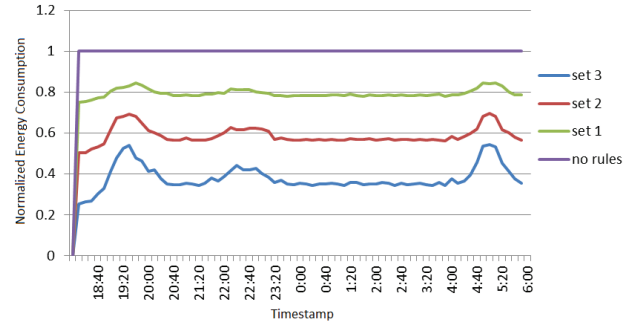

Set ID.	Power output (movement)	Power output (no movement)
Set 1	100 %	75%
Set 2	100 %	50%
Set 3	100 %	25%

cases, which will give a general idea of the simulation process. For the tests, a site corresponding to a residential area was used (Figure 3). Tests were set to start at 18:00PM and finish the next day at 6:00AM. Only pedestrian instances were used, and the generation rules were the same for all the generator waypoints. A total of four waypoint rules were generated, shown in Table 1.

The idea was to simulate a typical night, with people coming back home from work after six o'clock, then exiting again – in less proportion – to have dinner or have fun and then returning home. After midnight, the traffic is reduced significantly and has a homogeneous distribution. Finally, on early morning people start going to work again. From the luminaries perspective, the goal was to create a simple simulation, with a pair of global dimming rules controlling the entire scenario. For the tests, only the movement information was used (with two rules detecting the presence and absence of movement), as the idea was to let the light flux at full power in the presence of a pedestrian, and then lower this value in its absence. Three pairs of dimming rules were considered to show the differences between the resulting energy levels. Dimming rules are shown in Table 2.

The results were also compared to the worst case represented by luminaries without on-board intelligence, to give a better idea of the energy saving that could be achieved by means of the devised system. It is also important to mention that, for the simulations, the waypoints had all the same priority, meaning that the pedestrians would follow absolute random paths. Figure 8 shows the normalized energy consumption of all the luminaries for all the dimming rule pairs. The impact of the pedestrian generation rules can be clearly seen, with the homogenous region that goes from midnight to 4 o'clock in the morning and the three gaussian distributions with a height proportional to the amount of instances generated per waypoint. The most interesting aspect is related to the difference between each pair and the top line showing the energy consumption for a luminary with no intelligence. The plot in Figure 9 shows the average energy saving for the three pairs.

Another useful information that can be obtained from the simulations is the amount of energy consumed per luminary. An example is shown in the plot in Figure 10, where the consumed energy is depicted for a subset of luminaries. Here, it can be immediately seen that luminaries 1004, 1007 and 1009 are consuming more energy than the others. This be-

**Figure 8: Normalized energy consumption of all the luminaries.****Figure 9: Average energy saving (percentage) for all the dimming rule pairs.**

havior was actually expected for the given rules. However, results are also due to the fact that, based on its location, a luminary could be more or less prone to switch to full power. With this information, the administrator could identify which luminaries are located in crowded spots of the network, and can devise new particular rules for these subsets of luminaries that need more efficient dimming strategies.

The administrator could check the validity of the simulations, by changing the present set of rules or adding new specific rules to further compare results, defining an iterative workflow that attempts to solve the energy consumption requirements by moving from general strategies based on global rules to very specific – and more efficient – approaches exploiting local rules.

Though this case of study and the simulation model are relatively simple, it shows the potentiality behind the concept of intelligent street lighting. The quantity of information and details that can be generated by using robust conditions for the stimuli generation (pedestrians and vehicles) and by using more detailed sets of lighting rules could be limitless. As a simulation tool, the designed application could provide with significant information that can give a hint of the behavior of a set of rules applied to virtual luminaries, before they are applied to the real devices. The validation of the system will be carried out by the pilot sites, and the results of the tests will be used to refine the models that generate the input events for the simulator.

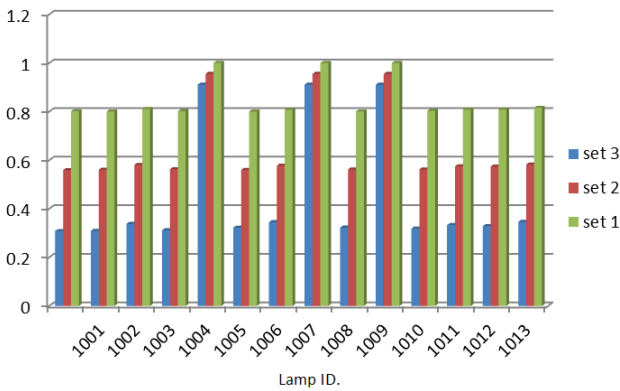


Figure 10: Normalized energy consumption per luminary for a subset of luminaries.

7. CONCLUSIONS AND FUTURE WORKS

In this paper, a street lighting management system based on intelligent LED luminaries has been presented. Luminaries are equipped with embedded sensors (monitoring environmental and lighting parameters) and on-board processing capabilities to implement local lighting rules with the aim of reducing energy consumption. The overall system allows lighting site administrators to address the tasks of monitoring, controlling, simulating and evaluating dimming strategies using virtual reality visualization and by following a multi-modality approach ensuring access to the server-side functionalities of the system via different client devices. Moreover, the simulation environment provides administrators with an up-to-date man-machine interface supporting stereoscopic 3D vision and natural interaction by interpreting body poses to explore the network of smart luminaries within a 3D digital scenario.

The paper also presented a set of simple practical simulations to show the capabilities of the designed tool; however, the simulation scenarios (e.g. the behavior of pedestrians) will be tuned when real data from pilot sites will be available. In this way, more complex simulation rules can be set and tested.

At the time of writing, the effort of the project was focused on providing a real-time tool, also supported by machines with limited hardware, to allow system administrators to tune the dimming policies and check corresponding power consumption. For this reason, the rendering techniques used for testing the overall framework are quite simple and indirect illumination, shadows, reflections and refractions are not currently managed. In this direction, future works will be aimed at completing the preliminary solution proposed in this paper by improving the realism of the scenario; in particular, a more complete and sophisticated shading model (encompassing indirect lighting and shadows) combined with the integration of the software with real lighting models of the lamps will lead to improve a correct representation of the overall environment. Moreover, the improvements in rendering realism can be exploited to automatically determine optimal illumination policies in the street lighting management system. In particular, renderings of the reconstructed site can be automatically analyzed to find out a quality measure of the perceived light, with the aim of finding a tradeoff between energy saving and ap-

propriate illumination, thus reducing costs and maintaining safety of the area.

8. ACKNOWLEDGMENTS

This article is developed within the frame of the project LITES “Led-based intelligent street lighting for energy saving” (G.A. no. 238916), co-funded by the EC under the Competitiveness and Innovation Framework Programme - ICT Policy Support Programme.

The authors would like to thank the anonymous reviewers for their valuable comments and suggestions to improve the quality of the paper.

9. REFERENCES

- [1] *3D Stereo - Consumer Stereoscopic 3D Solution*. NVIDIA Technical Brief.
- [2] Echelon Street Lighting Solutions. <http://www.echelon.com/solutions/streetlight/>, November 2011.
- [3] FAAST web site. <http://projects.ict.usc.edu/mxr/faast/>, November 2011.
- [4] LITES Project. <http://www.lites-project.eu>, November 2011.
- [5] Microsoft Kinect web site. <http://www.xbox.com/en-US/kinect/>, November 2011.
- [6] OpenNI web site. <http://www.openni.org/>, November 2011.
- [7] Primesense web site. <http://www.primesense.com/>, November 2011.
- [8] The Virtual-Reality Peripheral Network. <http://www.cs.unc.edu/Research/vrpn/>, November 2011.
- [9] W. Hai, S. Zabuawala, N. Zeferjahn, and J. Yadegar. Towards a real-time 3d situational awareness visualization for emergency response in urban environment. In *Proceedings of the 10th IEEE International Symposium on Multimedia*, 2008.
- [10] H. Jianchun, H. Jinxing, T. Qingyuan, and G. Shanshan. Layout optimization of urban underground pipeline based on 3d digital city. In *Proceedings of the Joint International Conference on Theory, Data Handling and Modelling in GeoSpatial Information Science*, 2010.
- [11] M. Miska, E. Santos, E. Chung, and H. Prendinger. Opentraffic - an open source platform for traffic simulation. In *Australasian Transport Research Forum 2011 Proceedings*.
- [12] M. Pack. Visualization in transportation: Challenges and opportunities for everyone. *IEEE Computer Graphics and Application*, 30(4):90–96, 2010.
- [13] G. Paravati, F. Lamberti, A. Sanna, and C. Celozzi. High-performance solutions for adaptive and customizable streaming of interactive content to mobile devices. In *Multimedia Services and Streaming for Mobile Devices: Challenges and Innovations*, USA, 2011. IGI Global.
- [14] H. Prendinger, A. Nakasone, M. Miska, and M. Kuwarhara. Openenergysim: Conducting behavioral studies in virtual worlds for sustainable transportation. In *IEEE Forum on Integrated and Sustainable Transportation Systems*.