VLC-Enabled Human-Aware Building Management System

Yehuda E. Kalay[1,2](✉), Haripriya Sathyanarayanan[1], Davide Schaumann[3], Albert Wang[4], Gang Chen[4], and Ramdas G. Pai[5]

[1] University of California, Berkeley, CA, USA
{kalay,haripriya_snarayanan}@berkeley.edu
[2] Technion, Israel Institute of Technology, Haifa, Israel
[3] Jacobs Technion-Cornell Institute at Cornell Tech, New York, USA
davide.schaumann@cornell.edu
[4] University of California, Riverside, CA, USA
{aw,gachen}@ece.ucr.edu
[5] School of Medicine, University of California, Riverside, CA, USA
ramdas.pai@medsch.ucr.edu

Abstract. "Smart" buildings that can sense and detect people's presence have been in use for the past few decades, mostly using technologies that trigger reactive responses such as turning on/off heating/ventilating, lighting, security, etc. We argue that to be considered truly smart, buildings must become "aware" about the locations and activities of their inhabitants so they can proactively engage with the occupants and inform their decision making with respect to which actions to execute, by whom and where.

To help assess the potential impact of "aware" buildings on their occupants, we are developing a multi-agent simulation-powered building management system that can sense human and building assets, extrapolate patterns of utilization, simulate what-if scenarios and suggest changes to user activities and resource allocation to maximize specific Key Performance Indicators (KPIs). The system is able to evaluate the implications of potential conflict resolution strategies and account for individual and collaborative activities of different types of users in semantically rich environments.

Sensing in our case is based on Visible Light Communication (VLC) technology, embedded in a building's LED lighting system. It can detect the actors, where they are located and what they do. To understand what happens in each space at any given time the information derived from the VLC system is combined with models of users' activity schedules, profiles, and space affordances.

We demonstrate our approach by hypothetically applying it to a Cardiac Catheterization Laboratory (CCL). The CCL is high-intensity hospital unit, second only to the Emergency Department in terms of the urgency of the cases it must handle. An aware building will help both patients and staff to allocate their (always scarce) resources more efficiently, saving time and alleviating stress.

Keywords: Smart environments · Human behavior simulation · Space utilization · Hospital environments · Visible light communication

N. Streitz and S. Konomi (Eds.): HCII 2020, LNCS 12203, pp. 207–222, 2020.
https://doi.org/10.1007/978-3-030-50344-4_16

1 Introduction

Sensing technologies that enable buildings to detect people's presence have been around for the past few decades. Their use was limited mostly to triggering reactive responses to people's presence (heating/ventilating, lighting, security, etc.). We argue that truly "smart" environments can leverage sensed information about the locations and activities of their inhabitants to proactively engage with the occupants and inform their decision-making processes with respect to which activities to execute, by whom and where. We claim that such buildings will be "aware" of their own status, as well as the status of their occupants and the activities that are performed within (and around) them.

To help assess the potential impact of such "aware" buildings on their occupants, we are developing a simulation-powered building management system that can sense the location and activities of human and building assets, simulate what-if future scenarios and suggest user activities and resource allocation that will maximize specific Key Performance Indicators (KPIs). Our system is able to evaluate the implications of potential conflict resolution strategies using a multi-agent simulation system that accounts for individual and collaborative activities of different types of users in semantically rich environments.

Sensing in our case is based on Visible Light Communication (VLC) technology, embedded in a building's LED lighting system (Pan et al. 2019). It can detect who are the actors, where they are located and what they do, as well as how the spaces themselves are used. Information derived from the VLC system is combined with models of actors' activity schedules, profiles, and space affordances, to understand what happens in each space at any given time. This data forms the current state of the building occupancy and utilization, and is used to simulate alternative possible future states for each actor and to resolve possible conflicts that may occur. The simulation and decision-making process are driven by a previously developed narrative-based modeling system that can simulate human behavior in buildings (Schaumann et al. 2017a, b, 2019). It produces alternative future states, revealing the consequences of enacting different resource allocation strategies. A priority function is used to evaluate and compare the alternative futures and choose the one that maximizes some previously agreed-upon utility function. Once the decision is made, the system uses VLC to communicate the information to the relevant actors who enact them.

While the approach being developed is, in principle, agnostic of the sensing and communication technology used, VLC has been chosen because unlike radio frequency it is highly localized and it does not interfere with the building's other sensitive instruments, which is critical in the case of hospitals (our chosen case study). In addition, because it is embedded in the building's LED lighting system, it requires little additional infrastructure compared to other technologies.

We demonstrate our approach by hypothetically applying it to the Cardiac Catheterization Lab (CCL) in a major hospital: an intervention cardiology unit in a hospital with imaging equipment used to diagnose the arteries and chambers of the heart and treat any stenosis or abnormality found. In addition to treating about 20 scheduled patients every day, the CCL also treats 1–2 emergency cases, known as ST-Elevation Myocardial Infarction (STEMI), every day. These acute heart attack cases require immediate attention, and can, therefore, disrupt the scheduled activities of the CCL. A smart building will help allocate resources more efficiently, saving time and alleviating stress in the Cath Lab, the Emergency Department, the in-patient wards, and more.

2 Building Automation

The effects of a given built environment on the people who inhabit it become apparent only after the environment has been built and occupied, making it difficult to assess the impact of technologically rich environments on their inhabitants. Buildings that have embraced truly advanced technological innovations are few, and these mostly focus on improving energy, lighting, and security performances (Jalia et al. 2019). To help understand and assess the potential impact of buildings that can adapt their performance to the dynamic needs of their inhabitants we have identified three levels of building automation (Kalay 2004) feedback regulated adaptability; model-based adaptability; and total environmental adaptability.

Feedback regulated adaptability is based on the concept of automation, where the results of some action are compared against some desired performance measures. Departure from the desired condition triggers another action to bring the results closer to the desired ones. In building automation, the ubiquitous thermostat demonstrates this principle: it instructs the HVAC system to heat (or cool) the air inside a building, until it reaches the desired pre-set temperature. Enabling buildings to sense and respond to such needs is a relatively simple, reactive kind of automation. It has been implemented in areas of control, regulation and supervision of electrical, mechanical and climatic control equipment.

Adding a functional model to networked building systems and appliances allows for a proactive adaptability approach to building automation: it helps to manage the building in expectation of events, rather than in response to them. A functional model of a building is one where the occupants' behavior patterns are programmed in advance, based on learning their typical preferences, so the building can anticipate and position itself to support recurring events, not only to respond to them (Mozer 1988).

We posit that true adaptability will be reached when the building not only responds—reactively or proactively—to its inhabitants' behavior, but it actively engages, even manages them. Such active management depends on much more information than the locations of the inhabitants and prevailing environmental conditions. Broadly, it must include information about *spatial* conditions, *activities*, and the *inhabitants* themselves:

– *Spatial* information includes the configuration of the building (rooms and the connections between them), the intended purpose of each room (e.g., a patient room in hospital, an Emergency Department, a nurse station, etc.), the environmental conditions prevailing in each space (light, temperature, noise, etc.), and current location of each inhabitant within these rooms/corridors.
– *Activities* information includes each inhabitant current, past and future activities: what is s/he doing now, what did s/he do earlier, and what is his/her schedule for the foreseeable future. It also includes information about customary scheduled activity sequences, and what to do in case of unplanned activities (e.g., 'Code Blue' in a hospital).

– *Inhabitants'* (which we call 'actors') information includes the identity of each actor, his/her profile (role in the organization—doctor, nurse, patient, visitor, etc.), abilities, degree of fatigue, and more.

Once the building management system has access to all this information it can form an image of the current state of the whole building and its inhabitants, and predict alternative future states, by means of simulation. It can then evaluate the cost/benefit of alternative future states and recommend choosing the one that seems most profitable in terms of some pre-determined Key Performance Indicators (KPI).

3 The Power of Seeing the Whole Picture

We call the ability to sense the overall state of the building and its occupants, and be able to extrapolate and evaluate future states, "the power of seeing the whole picture."

VLC affords such abilities in buildings, in terms of locating people and equipment in indoor spaces when individuals and equipment are tagged with VLC transponders. Figure 1 shows a hypothetical experiment of locating staff and patients in a medical ward in a hospital. This ability not only allows locating individual staff members, patients and visitors, but also to make assertions about their status: are they busy, free, or in need of medical attention. Likewise, it allows making assertions about the state of occupancy of spaces. When coupled with predefined space profiles, thermal, light, smoke and other detectors, it can inform the building management system about the specific conditions of every space.

The ability to see the whole provides an overview of some situation, not visible from the individual actor's point of view. Furthermore, as evident from Fig. 1, this ability extends from the present to the past: it is possible to trace previous locations of individuals and equipment at prior points in time.

It is our contention that this ability can also be extended into the future, by way of simulation, which will allow the building management system to predict the future locations and activities of the inhabitants. It could, therefore, consider alternative "futures" and help choose the one most desired (according to some predefined criteria). It is this ability which comprises total environmental adaptability.

We are developing a simulation-powered Building Management System (BMS) aimed at Total Environmental Adaptability that leverages the power of "seeing the whole." It will sense the presence and location of humans and building assets, extrapolate patterns of behavior and utilization, simulate what-if scenarios and suggest modifications to user activities and building operations to maximize specific Key Performance Indicators (KPI).

4 Method

The system is composed of four main components: (1) sensing, (2) simulating, (3) evaluating, and (4) acting.

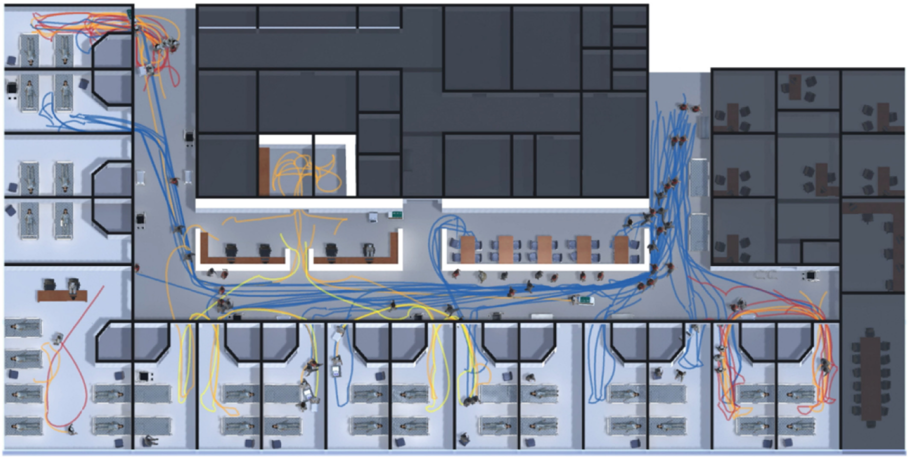

Fig. 1. Locating and tracking people in a hospital ward.

4.1 Sensing

For our case study we have chosen sensing by Visible Light Communication (VLC) system, which is a preferred communication method in hospitals—the domain of our research. Existing RF wireless technologies have many major limitations, including in particular, interference with medical devices that may potentially put patient safety at risk, such as ventilators, pumps, telemetry, defibrillators, brain stimulators, ophthalmic equipment, and many medical implants (Camulli 2014; Berger and Gibson 2013). Another major challenge is that medical-grade wireless must comply with extremely high security and regulatory standards to ensure patient privacy. Such high security cannot be ensured by existing RF wireless networking because RF signals are publicly open and can be intercepted. Any passcode protection may be cracked with enough effort.

Wireless VLC has many advantages over traditional RF technology: it can be embedded in energy-efficient solid-state LEDs (Kim and Schubert 2008). Other than greener lighting, LEDs can be switched ON/OFF at a speed of tens of MHz without flickering visible to the eye, enabling VLC at high data rate by modulating LED light (Conti 2008). The optical spectrum is unlicensed, unrestricted and orders of magnitude wider (300 THz) than the crowded RF spectrum, making wireless streaming of big data possible for large number of users. Visible light allows more emission power for higher data rates and better quality of service (QoS) without risking human health. VLC is unable to penetrate walls, hence ensures high security and privacy. Being interference-free, VLC can co-exist with RF technologies. LED also allows visible light positioning (VLP), due to its beaming nature, which offers indoor locating and navigation in hospitals, parking structures and shopping malls. VLC devices are cheaper than RF components (e.g., expensive 60 GHz mmWave devices). In a sense, VLC wireless is "free" because it is built on existing LED lighting infrastructure, providing VLC wireless streaming at a beyond-Gbps speed.

Ubiquitous VLC wireless systems will consist of modulated LEDs (lamps) for broadcasting and user terminals (smartphones with embedded photodetector as transceivers) to realize full-duplex optical wireless streaming anywhere, anytime for anyone. Figure 2 shows a typical VLC scenario in a hospital.

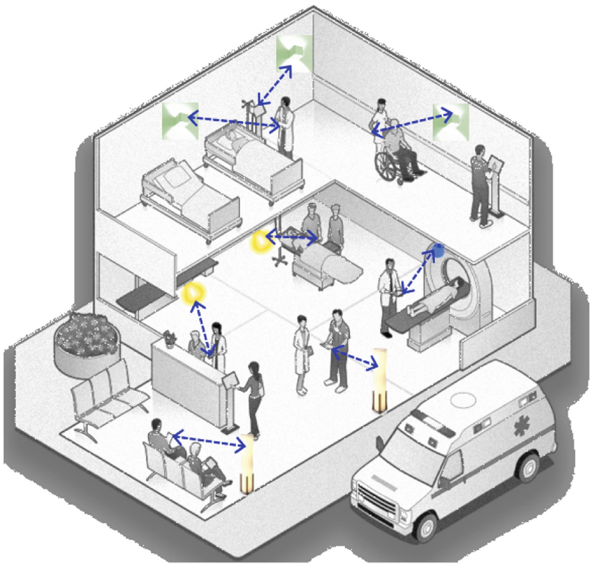

Fig. 2. VLC enables real-time communication positioning in a hospital.

4.2 Simulating

The purpose of the system is to use sensed data, combined with other data, to help make decisions about future actions. Since these decisions involve human activities, which are dynamic and depend on many factors such as spatial, occupational, and personal conditions, the method chosen to help predict future situations is simulation.

Simulation methods have been used to analyze the dynamic relationship between human activities and the surrounding environments in both existing and not-yet-built environments. They include particle-based methods, which describe pedestrians as homogenous particles subject to physical and social forces of attraction and repulsion (Helbing and Molnar 1995). Fluid-based methods describe people flow in fluid-like terms (Henderson 1971; Hoogendoorn and Bovy 2004; Hughes 2003). Cellular automata models provide an inherently spatial representation of occupancy, whereby each cell indicates its occupancy state and transition rules govern the evolution of a cell state (Blue and Adler 1999). Process-driven models consider structured sequences of activities that require a set of resources (e.g., people, equipment) and take a certain (usually stochastic) amount of time (Marmor et al. 2012).

In these models, space is often abstracted in the form of a graph where nodes represent rooms and link represent stochastic traversal times. In hospitals, however, several different processes may take place in the same space, and one process may affect the others. Unplanned social interactions between staff members and patients, for example, have been proven to affect the performing of other medical tasks (Seo et al. 2011). The aforementioned approaches cannot consider interactions among multiple parallel processes occurring in the same space.

We use a Multi-Agent System (MAS) where autonomous agents inhabit virtual environments and sense, plan and act individually or in groups to achieve a specific goal (Yan and Kalay 2004; Helbing et al. 2002; Zheng et al. 2009; Belhaj et al. 2014; Kapadia et al. 2015). While these approaches provide efficient solutions to simulate collision-free movement and social interactions (Thalmann and Musse 2013), they mostly focus on the abstract movement of occupants while ignoring the setting where a behavior is enacted (the specific building type) and the context-dependent activities that people engage in (task-based behaviors in healthcare facilities).

Recent work on narrative-based modeling (Schaumann et al. 2017; Simeone et al. 2013) demonstrated a different approach to simulating day-to-day occupancy scenarios in complex facilities, like hospitals. The approach is centered on the concept of *narratives*, which are rule-based scripts that coordinate the collaborative behaviors of heterogeneous actors (doctors, nurses, patients) who perform a structured sequence of activities (checking a patient, distributing medicine) that unfold in semantically rich spaces.

A key aspect of this approach involves distributing ‘intelligence’ among the different components of the model. To that effect, the narratives are responsible for the high-level organization of low-level activities into task-based procedures. To relieve the narratives from the need to handle low-level calculation processes, both actors and spaces entities are equipped with autonomous calculation abilities and can dynamically update their status based on contextual social and spatial conditions. That status can be retrieved by the activities and the narratives during their execution so that they can make the most informed decision at any given time.

A narrative manager determines which narrative to trigger at a given time, depending on the current state of the world (the current simulated time or the proximity of actors in a space). In contrast to other approaches that simulate scheduled activities in workplaces, the execution of narratives can adapt to dynamic conditions. For instance, unplanned narratives (such as staff-patient interactions) can cause delays to planned narratives (a patient check).

4.3 Evaluating

Measurable performance indicators are obtained from the simulation results, that can be compared to predefined absolute threshold measures or relatively to one another. They may include patients’ length of stay, average waiting times, overall throughput, congestion, staff work schedules and shift loads, staff or space utilization, and other space occupancy indicators.

Performance indicators consists of hard and soft criteria. Hard criteria are quantitative, measurable performances, such as walking distance, length of stay, and throughput of a clinic or ward. Soft criteria are typically qualitative, based on subjective perceptions, such as social, psychological, and organizational policies. In many cases, the same performance results may be valued differently by different stakeholders. To create a building-wide management system, it is therefore necessary to create a shared world-view that incorporates the relative merits of each action from different points of view, and reconciles the differences among them in light of shared, higher-level objectives (Kalay 2004). Since evaluation criteria differ from one other, the evaluation process requires a tradeoff mechanism that balances competing needs: it needs to choose optimization of one performance criterion over others, or strike a balance in the degree to which any performance criterion is achieved, assuring that overall performance is maximized (Kalay 2004).

The process of evaluating the simulation results includes three main phases: (1) identifying the relevant KPI for the simulation; (2) simulation results for the KPIs; (3) comparison of simulation results to benchmark values and normalizing them so they can be compared to one another for evaluation. Evaluating the results and drawing conclusions on the performance of the simulated behavior of the system through the assignment of weights and priorities to arrive at overall scores for a ranking (see Table 2).

The first determines which KPI are needed in relation to time, space, activity, and actors. The second phase is a direct result of the simulations with the metrics for the KPIs. The third phase normalizes the results of the simulation with different numeric measurements to one scale of scores, for example: from 0 (worst) to 5 (best). The normalization is achieved by comparing the simulation (absolute) results to benchmarks based on organizational goals, policies, culture, professional guidelines, norms and regulations, evidence-based design, or expert's opinion.

The evaluation is a study of the implications of the simulation results, based on evidence from research, experience, or precedents. For example, the implications of high level of noise that was predicted in the simulation can be evaluated by their impact on increased patient's stress, pain and depression, decreased patient's privacy and communication with family and staff, decreased staff effectiveness and increased potential of medical errors (Ulrich et al. 2008).

4.4 Acting

Once the comparative evaluations are completed, it is then possible to recommend enacting the most desired—or least disruptive—action. This action is communicated to the relevant stakeholders via the building's two-way communication system, which as mentioned earlier in our case is by means of the VLC system. If the preferred action involves the building's mechanical systems, such as HVAC, lighting, etc., the preferred action may be communicated directly to the assets involved. Like other "recommender" systems, such as GPS-based driving instructions, the actors may accept or ignore the recommended action. Either way, their action will be sensed by the building, and become input for the next round of simulation/evacuation/action.

5 Case Study

We have implemented our system on a hypothetical Cardiac Catheterization Laboratory (CCL). A CCL is a hospital unit equipped with imaging equipment used to visualize the arteries and the chambers of the heart and treat any stenosis or abnormality found. It performs diagnostic, interventional, and electro physiology procedures, serving outpatients, inpatients, and emergency cases.

The CCL while being one of the departments in the hospital, impacts and is impacted by other departments. The challenge is to evaluate those impacts to avoid conflicts and maximize the overall efficacy of the hospital. The simulation process necessarily requires abstraction of a complex system into a simplified model, and experimenting iteratively on it to test the relationship among many variables interacting in complex and often unpredictable ways (Shannon 1998).

The hypothetical case study CCL has five Cath Labs: three Cardiac Catheterization (CC) labs, one Electro Physiology (EP), lab, and one Hybrid Cath Lab. Typically, 20–25 procedures are planned for each workday. In addition, the CCL handles 1–2 unplanned emergency cases every day. The CCL is staffed by 15 medical staff members and operate from 7.30 am to 5 pm every day and may run overtime depending on the procedures and other emergencies. Typically, a diagnostic procedure involves a team of three staff members (a cardiologist and two nurses) and lasts 20–30 min, while an interventional procedure involves a six-member team (cardiologist, anesthesiologist, three nurses and a technician), along with a nurse in the observation area, and lasts 45–90 min. The labs interact with a 15-bed Cardiac Acute Care Unit (CACU) where patients are prepared for the procedure and recover from it (see Fig. 3).

Fig. 3. A typical floorplan of a CCL

5.1 Planned Events

Figure 4 depicts the typical planned activities workflow of the CCL operations. It comprises of four activity blocks that begin after pre-procedure preparation of the patients, either at the CACU for outpatients, the hospital nursing wards for inpatients, or the Emergency Department, depending on the type of patient. These include the 'patient transfer to the procedure room,' 'patient preparation,' 'procedure' and 'transfer back for recovery.' The times depend on the type of patient and type of procedure. There can be additional waiting time dependent on the availability of the cardiologist for the procedure, including other causes for delays such as transfer times and availability of staff for the procedures. The procedure workflow for a patient in the CCL is shown (see Fig. 4) with all the actors and activities involved at various stages in the activity blocks.

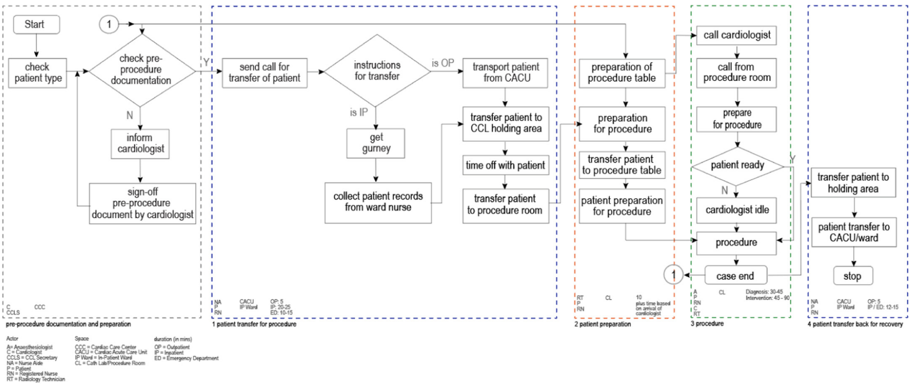

Fig. 4. Typical Cath Lab workflow for the planned procedures.

5.2 Unplanned Events

To demonstrate the proposed system, we look at the impact of an *unplanned* or unscheduled event on the CCL and the overall hospital. The emergency case known as STEMI (ST-Elevation Myocardial Infarction) was taken for the unplanned event. The STEMI is a very serious type of heart attack during which one of the heart's major arteries (one of the arteries that supplies oxygen and nutrient-rich blood to the heart muscle) is blocked. For STEMI patients, access to a facility with percutaneous coronary intervention (PCI) capabilities is time-critical: Door-to-Balloon (D2B) time must be less than 90 min. The D2B is typically reduced by sending STEMI patients directly to the CCL by pre-activating the CCL, instead of sending them to the ED (Martel et al. 2017).

Although the frequency of STEMI cases is typically 1–2 a day, they can be considered as 'unplanned' events that disrupt planned events at the CCL: a STEMI protocol requires an immediate activation of a suitable Cath lab and medical team to prevent delays in care.

Unless the CCL includes a standby un-used CL and medical team to handle such cases (as some hospitals do), this presents a suite of challenges within the CCL, as it may disrupt planned treatments for scheduled patients and medical teams: not all CLs are suitable to treat STEMI patients, and those that are might be occupied with ongoing procedures at different stages of completion. Cardiologists may be occupied or have a scheduling conflict if assigned the STEMI patient. Furthermore, these challenges go beyond the CCL, as changes to the planned schedules can have a negative effect on other outpatients and inpatients: patients that were scheduled to undergo treatment may be bumped, requiring rescheduling (of outpatients) and longer stays (for inpatients). Hence, an action that may seem optimal for the CCL may adversely affect other units of the hospital, and thus be less optimal overall. Such overall implications are often not visible to the CCL staff: resolving them requires "seeing the whole picture," as discussed earlier.

It is the goal of the system described here to critically evaluate all the options facing the CCL in case of a STEMI, and to recommend the overall most suitable plan of action for the hospital as a whole (subject, of course, to the constraints of the STEMI protocol and others).

5.3 Sensing

The overall best course of action is determined by detecting the state of the CCL in term of spaces, actors, and activities, when a STEMI protocol is declared. Table 1 shows the hypothetical state of the five Cath Labs depicted in Fig. 3 in terms of the types of patients and medical teams involved, and the type of ongoing procedures, their expected duration and possibility for interruption of the ongoing procedures by the unplanned STEMI event.

Table 1. Narratives of the planned procedures in the CCL. IP = In-Patient; OP = Out-Patient; CC = Cardiac Catheterization; HCL = Hybrid Cath Lab; EP = Electro Physiology

Space	Type	Procedure	Patient	Duration	Time left
CL1	CC	Intervention	IP1	80 min	40 min
CL2	CC	Intervention	IP2	45 min	45 min
CL3	CC	Diagnosis	OP1	30 min	30 min
CL4	HCL	Diagnosis	OP2	30 min	10 min
CL5	EP	EP	OP3	35 min	30 min

From Table 1 the current state can be determined as follows: the procedure in CL1 is mid-way, the procedures in CL2 and CL3 are yet to commence, the procedure in CL4 is nearing completion, while CL5 is an EP lab, making it unsuitable for the treatment of the STEMI.

5.4 Simulating

The Cath labs need to be evaluated to determine which one of the available labs can be assigned for the STEMI (CL1, CL2, CL3, or CL4), with simulation of the consequences of choosing any one of them and evaluating their relative merits. The Event-Based simulation described earlier is used to determine:

- The procedure in CL1 cannot be interrupted, therefore that lab is not available to treat the STEMI.
- If CL2 is chosen to treat the STEMI, Patient IP2 (an inpatient) who was scheduled to be treated will have to be delayed or rescheduled. The patient will be taken back to the inpatient ward, where (s)he will stay at least another day to be treated (we assume their condition allows such postponement of the treatment). Patient IP2 will not be discharged from the hospital as previously planned and will instead continue to occupy a bed in the cardiac in-patient ward.
- Furthermore, the continued hospitalization of IP2 will prevent admission of an incoming patient from the Emergency Department, who was scheduled to be hospitalized in the cardiac in-patient ward. Instead, she will have to stay in the ED for another 24 h, at a great inconvenience to her and the ED staff.
- If CL3 was chosen to treat the STEMI, Patient OP1 (an outpatient) who was scheduled to be treated will be delayed. He will be taken back to the CACU, delaying treatment of other outpatients scheduled for the day. Since the policy of the CCL is to treat all out-patients that were scheduled for the day rather than sending them home to be treated another day, the CCL clinical staff will have to stay for a longer shift.
- If CL4 is chosen to treat the incoming STEMI, it will take 10 min to complete the ongoing procedure, and another 15 min for the turnaround time (TAT) to make it ready for the incoming STEMI. This will result in an overall 25 min delay in treating the STEMI.

5.5 Evaluating

The results of these simulations are evaluated comparatively to a list of Key Performance Indicators partially drawn from the literature (Brodeschi et al. 2015; Mozer 1988) and discussed with an expert/lead-cardiologist at the hospital's CCL, resulting in sixteen relevant and feasible KPIs that were chosen for this demonstration. The KPIs were grouped into three categories: operational, user related and space related, with eleven operational KPIs, three user-related KPIs, and two space-related KPIs. The KPIs were structured and ranked based on relative importance for the CCL, inpatient and emergency departments, along with executive KPIs for the organization.

For user specific KPIs that address 'patient satisfaction,' proxies such as patient wait times and staff load schedules were used. Inter-departmental relations were accounted for based on the goals or executive KPIs for the hospital to ensure there were no undesired trade-offs where processes within the CCL interact with processes outside the CCL.

Th e simulation results for the KPIs were evaluated based on the method described above to arrive at overall scores for the Cath labs that included the assignment of weights and priorities after normalization.

Figure 5 shows the final evaluation and scores based on category and the Cath lab. While different Cath labs score high in the different categories, if the user-related KPIs are given more weightage by the organization, it is CL3 that is the preferred choice to minimize undesirable consequences such as patient and staff dissatisfaction.

In case of conflicts, the organization's policies and preferences are obtained for the recommended action. Similarly, the scores help understand conflicting needs and consequences of actions on the KPIs.

Table 2. Key performance indicators

Key Performance Indicators As defined by Stakeholders		Simulation Results Key Performance Indicators		Benchmark Values	
Operations	Average LOS	Duration of inpatient hospitalization	___days	4.8 days	↓
	Bed turnover rate	Number of discharges / Number of beds	___%	68%	↑
	Cancelled Procedures	Percentage of procedures cancelled/rescheduled	___%	2%	↓
	Average LOS ED	Time between arrival and departure of patient from ED	___hours	9 hours	↓
	First case on-time starts	Percentage of first cases of the day that start on time	___%	100%	↑
	Average procedural time	Procedural time for different procedures	___mins	Diagnosis:30 mins Intervention:50 mins	↓
	Turnaround time (TAT)	Turnaround time between cases (TAT)	___mins	17 mins	↓
	Average time in pre/post procedure holding	Time in pre & post procedure holding area by procedure	___mins	20 mins	↓
	STEMI D2B ≥ 90 minutes	STEMI patients with D2B ≥ 90 minutes	___#	90 mins	↓
	Average LOS post procedure	Duration of time post procedure recovery	___hours	4 hours	↓
	Overall Patient throughput	Time taken by patient from admission to discharge	___hours	6 hours	↓
User	Average Patient Wait Times	Patient wait times	___hours	2 hours	↓
	ED Waiting Time	Wait time in ED	___mins	50 mins	↓
	Staff Load Schedule	Engagement time for treatment and volume of patients	___hours	8 hours	↓
Space	Bed Occupancy	Number of beds occupied/Number of beds	___%	80%	↑
	Room / Asset utilization	Number of rooms utilized/number of rooms	___%	80%	↑

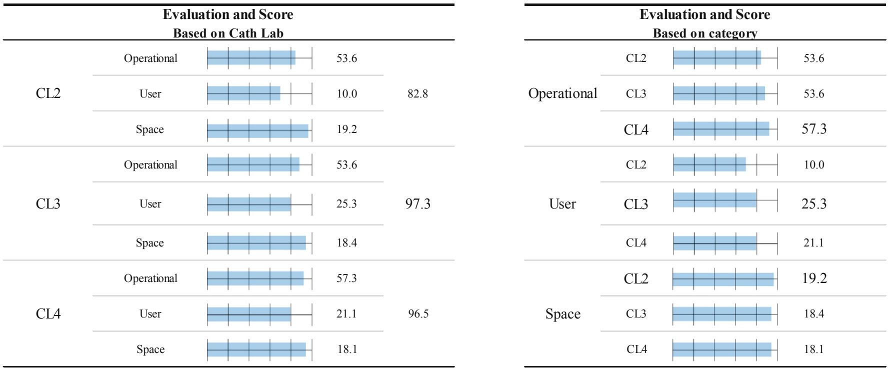

Fig. 5. Scores based on category and Cath Lab for ranking

5.6 Acting

Based on the comparative evaluations, the most desired—or least disruptive—action is communicated to the relevant stakeholders. This is done via the building's two-way communication system, which as mentioned earlier in our case is by means of the VLC.

6 Implementation Details

We use Unity 3D as simulation engine. Unity® 3D is a popular video game engine that features advanced physics and artificial intelligence libraries to model collision avoidance and path-finding.

Spaces have been modeled using Autodesk's AutoCAD® and then imported into Unity 3D using the FBX format. The spaces, actors, activities, narratives and narrative manager have been modeled directly in Unity 3D using Microsoft C#. Activities and narrative have been modeled as co-routines. Narrative co-routines are composed of a structured set of activity coroutines that are nested within the narrative and executed one after the other while yielding at each time step a status to the parent narrative (running or completed). Based on such a status, each narrative updates its own status, which is reported to the narrative manager. To run multiple narratives involving different actors concurrently, we have leveraged Unity 3D's ability to run co-routines in parallel, emulating multi-threading processing.

7 Conclusions

The notion of "smart" buildings has gained popularity with the advent of digital technologies that allow ever-growing sensing and control of building assets and inhabitants. Such systems have mostly been used to improve the energy behavior of buildings and their security by relieving the building's occupants from having to manages such chores. In this paper we argue that true "smartness" will be achieved when the building's omniscience can be harnessed in the service of its inhabitants. Such omniscience implies that the building, unlike its inhabitants, "knows" all that is happening within and around it at any moment. When coupled with operational procedures and occupant profiles, such knowledge can be used to predict and evaluate future events and recommend choosing the most beneficial one for each and every inhabitant. We have demonstrated such abilities in the case of a hospital environment.

To help assess the potential impact of such "aware" buildings on their occupants, we are developing a simulation-powered building management system that can sense human and building assets, extrapolate patterns of utilization, simulate what-if scenarios and suggest changes to user activities and resource allocation to maximize specific Key Performance Indicators (KPIs). Different from existing approaches, our system is capable of evaluating the implications of potential conflict resolution strategies using a multi-agent simulation system that accounts for individual and collaborative activities of different types of users in semantically rich environments.

Sensing in our case is based on Visible Light Communication (VLC) technology, embedded in a building's LED lighting system. It can detect who are the actors, where they located are and what they do. It can also detect how spaces are used. Information derived from the VLC system is combined with models of actors' activity schedules, profiles, and space affordances, to understand what happens in each space at any given time. This data on the current state of the building occupancy and utilization is used to simulate alternative possible future actions for each actor and to resolve possible conflicts that may occur. The simulation and decision-making process is based on a narrative-based modeling system previously developed to simulate human behavior in buildings. It produces alternative future states, revealing the consequences of enacting different resource allocation strategies. A priority function is used to evaluate and compare the alternative futures and choose the one that maximizes some previously agreed-upon utility function. Once the decision is made, the system uses VLC to communicate the information to the relevant actors (i.e., the occupants) who can enact them.

While the approach being developed is agnostic of the sensing and communication technology used, VLC has been chosen because unlike radio frequency it is highly localized, and it does not interfere with the building's other sensitive instruments, which is critical in the case of hospitals (our chosen case study). In addition, because it is embedded in the building's LED lighting system, it requires little additional infrastructure compared to other technologies.

We demonstrate our approach by applying it to a hypothetical Catheterization Lab in a hospital, which treats both scheduled (planned) and emergency (un-planned) patients. Un-planned emergency cases, known as STEMIs, may create scheduling conflicts, as they require immediate attention that can disrupt on-going operations. An aware building can help the CCL staff make decisions that are optimal from the overall hospital point of view, saving time and alleviating stress of both staff and patients.

To build the system we have interviewed medical staff in a major hospital and conducted a literature search for the most pressing issues that need to be handled by a CCL. Models of patient flow, staff work schedules, and space/equipment utilizations were developed. These provide the system with knowledge of past, present and future events at any given time, where each event is comprised of an actor-space-activity triad.

Results indicate that smart environments of the kind described here hold promise to enhance decision-making capabilities of building inhabitants, thus enabling building management strategies that support human needs and efficiency requirements, especially in mission-critical facilities, such as hospitals.

Acknowledgment. The research reported in this paper was made possible through Award #1838702 of Division of Information & Intelligent Systems (IIS), Directorate for Computer & Information Science & Engineering (CISE), U.S. National Science Foundation, and Grant #340753 of the European Research Council (ERC).

References

Belhaj, M., Kebair, F., Ben Said, L.: Agent-based modeling and simulation of the emotional and behavioral dynamics of human civilians during emergency situations. In: Müller, J.P., Weyrich, M., Bazzan, A.L.C. (eds.) MATES 2014. LNCS (LNAI), vol. 8732, pp. 266–281. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-11584-9_18

Berger, H.S., Gibson, H.M.: Managing your hospital RF spectrum. Biomed. Instrum. Technol. **47**, 193–197 (2013). The wireless challenge

Blue, V., Adler, J.: Cellular automata microsimulation of bidirectional pedestrian flows. Transp. Res. Rec. J. Transp. Res. Board **1678**(1999), 135–141 (1999)

Brodeschi, M., Putievsky Pilosof, N., Kalay, Y.E.: The definition of semantic of spaces in virtual built environments oriented to BIM implementation. In: Computer Aided Architectural Design Futures, pp. 331–346 (2015)

Camulli, E.: 5 Wi-Fi performance challenges in hospitals. 7 Signal, March 2014. http://7signal.com

Conti, J.P.: What you see is what you send. Eng. Technol. **3**, 66–68 (2008)

Kalay, Y.E.: Architecture's New Media: Principles, Theories and Methods of Computer-Aided Design. MIT Press, Cambridge (2004)

Mozer, M.C.: The neural network house: an environment that adapts to its inhabitants. In: Coen, M. (ed.) AAAI Spring Symposium on Intelligent Environments, pp. 110–114. AAAI Press, Menlo Park (1988)

Kim, J.K., Schubert, E.F.: Transcending the replacement paradigm of solid-state lighting. Opt. Express **16**(26), 21835–21842 (2008)

Helbing, D., Molnar, P.: Social force model for pedestrian dynamics. Phys. Rev. E **51**(5), 4282 (1995)

Helbing, D., Farkas, I.J., Molnar, P., Vicsek, T.: Simulation of pedestrian crowds in normal and evacuation situations. Pedestr. Evacuation Dyn. **21**(2), 21–58 (2002)

Henderson, L.F.: The statistics of crowd fluids. Nature **229**, 381–383 (1971)

Hoogendoorn, S.P., Bovy, P.H.: Pedestrian route-choice and activity scheduling theory and models. Transp. Res. B Methodol. **38**(2), 169–190 (2004)

Hughes, R.L.: The flow of human crowds. Annu. Rev. Fluid Mech. **35**(1), 169–182 (2003)

Jalia, A., Bakker, R., Ramage, M.: The edge, Amsterdam: showcasing an exemplary IoT building. Technical report, Centre for Digital Built Britain, University of Cambridge, UK (2019)

Kapadia, M., Pelechano, N., Allbeck, J., Badler, N.: Virtual Crowds: Steps Toward Behavioral Realism. Synthesis Lectures on Visual Computing: Computer Graphics, Animation, Computational Photography, and Imaging, vol. 7, no. 4, pp. 1–270 (2015)

Marmor, Y.N., Golany, B., Israelit, S., Mandelbaum, A.: Designing patient flow in emergency departments. IIE Trans. Healthc. Syst. Eng. **2**(4), 233–247 (2012)

Martel, T.J., Nambudiri, V., Kirkman, D., Martel, J., Pappas, L., Jehle, D.: Activation of the cardiac catheterization lab for STEMI patients. Emerg. Med. **49**(6), 259–262 (2017)

Pan, Z., et al.: Visible light communication cyber-physical systems-on- chip for smart cities. J. Commun. **14**(12), 1141–1146 (2019)

Schaumann, D., Pilosof, N.P., Sopher, H., Yahav, J., Kalay, Y.E.: Simulating multi-agent narratives for pre-occupancy evaluation of architectural designs. Autom. Constr. **106**, 102896 (2019)

Schaumann, D., Date, K., Kalay, Y.E.: An event modeling language (EML) to simulate use patterns in built environments. In: Proceedings of the Symposium on Simulation for Architecture and Urban Design, p. 21. Society for Computer Simulation International (2017)

Schaumann, D., Breslav, S., Goldstein, R., Khan, A., Kalay, Y.E.: Simulating use scenarios in hospitals using multi-agent narratives. Build. Perform. Simul. **10**(5–6), 636–652 (2017)

Seo, H.-B., Choi, Y.-S., Zimring, C.: Impact of hospital unit design for patient-centered care on nurses' behavior. Environ. Behav. **43**(4), 443–468 (2011)

Shannon, R.E.: Introduction to the art and science of simulation. In: 1998 Winter Simulation Conference, Proceedings (Cat. No. 98CH36274), vol. 1, pp. 7–14. IEEE (1998)

Simeone, D., Kalay, Y.E., Schaumann, D.: Using game-like narratives to simulate human behavior in built environments. In: Proceedings of the 18th International Conference on Computer-Aided Architectural Design Research in Asia, Singapore, pp. 199–208 (2013)

Thalmann, D., Musse, S.R.: Crowd Simulation. Springer, London (2013). https://doi.org/10.1007/978-1-4471-4450-2

Ulrich, R.S., Zimring, C., Zhu, X., DuBose, J., Seo, H.B., Choi, Y.S., Quan, X., Joseph, A.: A review of the research literature on evidence-based healthcare design. HERD: Health Environ. Res. Des. J. **1**(3), 61–125 (2008)

Yan, W., Kalay, Y.E.: Simulating the behavior of users in built environments. J. Archit. Plann. Res. **21**(4), 371–384 (2004)

Zheng, X., Zhong, T., Liu, M.: Modeling crowd evacuation of a building based on seven methodological approaches. Build. Environ. **44**(3), 437–445 (2009)