

A Service Infrastructure for Human-Centered IoT-Based Smart Built Environments

Denis Gračanin¹(✉), Mohamed Handosa¹, and Hicham G. Elmongui²

¹ Virginia Tech, Blacksburg, VA 24060, USA
gracanin@vt.edu

² Alexandria University, Alexandria, Egypt

Abstract. Smart built environments enhanced with technology can improve the lives of individuals, groups, and the broader community. Internet of Things (IoT), a collection of networked and interacting embedded devices, could provide the necessary infrastructure and enabling technologies to design, develop and deploy smart built-environments. We describe an approach to modeling IoT-based smart built environments that uses a large-scale virtual environment where a building model is aligned with the physical space. This approach takes advantages of affordances and embodied cognition in a large physical space to model user interaction with built spaces. The built space contains ‘smart objects’ with embedded sensors/actuators/controllers (e.g., kitchen appliances). A ‘smart object’ has the corresponding virtual object in the virtual environment. We build on our work on the conceptual design of interaction middleware and context sensitive interaction interoperability frameworks to develop support for a living ecosystem of services, a service framework, to support interactions with a smart built-environment. We illustrate the proposed framework on a case study, a living lab for smart built-environments that includes a kitchen, a living room, a bathroom, an office, and a bedroom.

Keywords: 3D user interactions · Internet of Things · Built environments

1 Introduction

Internet of Things (IoT), a collection of networked and interacting embedded devices, in connection with pervasive computing, could provide the necessary infrastructure and enabling technologies to design, develop and deploy smart built-environments. Designing and deploying IoT into the built environments provides capacities that can change how systems behave and how users interact with them. Such smart built environments enhanced with technology can improve the lives of individuals, groups, and the broader community. IoT could support sustainable and comfortable living through design, simulation, planning, monitoring, optimization, and visualization tools. Similarly, ambient intelligence would allow us to learn about people activities, their responses and interests,

how to make intuitive and data-based decision about what people want, and which action to perform.

The physical and social structures within a built environment are subject to continuous change. Thus, adding information technology and related services into our homes has to be done considering the multitude of activities architectural space support. House is a social place where family and friends meet and interact. As a consequence, service support for multiple users must accommodate for resource and preference conflicts. The conflict-resolution approach, while based on priorities and access control, also depends on the context and ambient intelligence. At the intersection of interaction design and built environments is thinking about how to “architect interaction.”

The design factors range from cultural, environmental, to more specific interaction goals such as effective navigation to the exit of a building. However, the common design goal for the services is to enhance the user experience and achieve the desired comfort level by providing easy and flexible ways to interact with the surrounding environment as well as to provide automatic actions based on decisions inferred from collected data and user preferences.

In the context of energy and living comfort, an energy-efficient house is effective only if the inhabitants are comfortable. Therefore, quantitative measures such as energy consumption or temperature must be augmented with more qualitative or subjective measure based on the individual and group preferences and desired comfort level. There are many efforts along those lines (e.g., Nest thermostat) but we are still lacking a comprehensive, holistic approach that should overcome the disconnect between quantitative and qualitative measures.

We build on our work on the conceptual design of interaction independence middleware and context sensitive interaction interoperability frameworks to develop support for a living ecosystem of services, a service framework, to support interactions with a smart built-environment. These services are created and destroyed over time based on the evolution of users’ needs and habits [14].

User tasks in built environments usually involve interaction with physical objects that respond to user actions with some kind of feedback. Services should allow users to both issue commands and receive feedback. They exist on multiple scales ranging from the design of door knobs to the placement of buildings.

The services are organized into layers. The IoT layer services provide connectivity and communication to embedded devices, sensors and actuators. The data layer services provide data collection and data fusion. The resource layer services support resource management (e.g., energy, water). The comfort layer services allow users to manage comfort settings. Finally, the UI layer services support multi-modal interactions across a variety of interface and interaction devices in a smart built-environment. In addition, the service management subsystem manages the service life-cycle and deals with user collaboration and conflict resolution, context awareness, behavior detection and safety/security/privacy issues.

Technology does not only dominate our work lives but has become an integral part of our domestic lives as well. Instead of replacing existing media, new technology is adapted for the existing patterns of use [9]. The ‘smart house’, or

‘home of the future’, such as the Aware House built by Georgia Tech, augments the traditional home with a rich computational and communicational infrastructure, such as “smart” devices and sensors that can detect and interact with the inhabitants of the house in novel ways.

We illustrate the proposed framework on a case study, a living lab for smart built-environments [17]. We describe the implementation of the framework and provide several examples of service creation and supporting multi-modal interactions in a single-user and multiuser setup. These examples demonstrate service for kitchen activities and proactive light control in a living room. We also present a corresponding simulation testbed that uses an enclosed space $50 \times 40 \times 40$ feet with real-time tracking and spatial audio capabilities to provide a multi-user mixed-reality/virtual reality environment. Several users can simultaneously move in the physical space and interact to test the service ecosystem before its deployment in the living lab.

2 Related Work

The new sensor, mobile, and control technologies have a great promise in connecting users to their environment. Smart built environments (e.g., a smart house) enhanced with technology can support better living to improve the lives of individuals, groups, and the broader community, including increased awareness of information in the user’s surroundings, integrated control over factors in one’s surrounding and home environments, and increased ability to support a sustainable living for both individuals and groups. A smart house can be used in many contexts, such as residence, office, store, classroom, etc., or can be a part of a larger place that includes the surrounding area and other buildings. Each smart house or building and their inhabitants contribute to the overall data collection, including participatory sensing and crowd-sourced data, thus supporting a ‘smart city’ [20].

The notion of a smart home, aware home, or green home has been explored toward incorporating environmentally friendly initiatives and techniques into the built-home environment. Technology plays a key role in monitoring information generated by these homes and in controlling certain aspects of the home. Awareness can be raised through the use of mobile notification systems, which attempt to deliver current, important information to the user in an efficient and effective manner. Some key Human-Computer Interaction (HCI) related questions regarding the presentation of information via mobile and other devices: enabling individuals to understand how their decisions may impact the environment, and enabling groups to reach consensus on decisions.

In addition, supporting interactions with built environments raises some unique usability challenges. Given the multitude of devices that all the inhabitants of the space can each interact with, it becomes apparent that these interactions can become quite complex. For example, they can depend on a given augmented device, the identity of the user that wants to interact with it, the presence of other users in the same or other location, the time, and any other

contextual information. As users can take on different roles in the household, there is a potential for conflicting rules within policies, in which case a conflict resolution mechanism needs to be put in place.

Harrison and Dourish emphasized the difference between space and place, defining a place as space with added socio-cultural understandings [11, 18]. They argue that “place, not space, frames appropriate behavior.” The physical and social structures within the home are subject to continuous change. Thus, adding information technology into our homes has to be done considering the multitude of activities architectural space support.

Aipperspach et al. discussed the dangers of introducing information technology into the home without considering potential detrimental effects on its inhabitants [2]. They suggest the idea of a semi-smart home, where areas are intentionally kept free of technology, except for so-called tourist devices (devices that may visit certain areas, yet have a specific ‘home’).

Nissenbaum introduced the concept of contextual integrity to address concerns about the effects of information technology on privacy from the perspective of the law [24]. Contextual integrity is based on norms of appropriateness and norms of information flow and can be used to determine whether privacy expectations have been violated.

Crabtree and Rodden studied domestic routines as related to communication and collaboration [9]. They introduce three major concepts: ecological habitats as places where communication media live, activity centers as places where media are produced and consumed, and coordinate displays as places where media are made available to coordinate activities.

The concept of Internet of Things (IoT) [6] describes the pervasive presence of things or objects which use a unique addressing scheme to interact with each other and cooperate with their neighbors to reach common goals. These physical objects have a social existence that could be supported through the IoT (an Internet of social things) [23].

The applications of IoT range from automation and manufacturing to assisted living and e-health. Designing and deploying IoT into the built environments provides capacities that can change how systems behave and how users interact with them. IoT could support sustainability through design, simulation, planning, monitoring, optimization, and visualization tools [15].

The IoT architecture and the corresponding implementation [10, 22] differ from the traditional network architecture. A large number of devices have to be connected, most of them with limited computing and networking capabilities. The IoT devices will be deployed in various contexts [3], including wearable devices, house appliances/sensors, embedded devices/smartphones, manufacturing plants and environmental sensors. The can span large urban areas to support ‘smart cities’ [20].

The evaluation of IoT and distributed sensor systems includes real-world testing, miniature prototypes and software simulations [5]. Sensor device emulators could be used to fully replicate the behavior of the deployed sensors. While software simulations are more convenient compared to the real-world testing, most of

them focus on the low-level networking aspects, with well-defined topologies and arrangements of objects. However, recent efforts are focusing on efficient simulation methodologies for large-scale IoT systems in urban environments from an application-layer perspective [8].

Distributed test system frameworks for open-source IoT software have been proposed [26] to support continuous integration techniques and a permanent distributed plugtest for network interoperability testing. Since one of the key function of IoT systems is data collection and processing, it is important to understand how data analytics solution will work for IoT systems. Benchmark toolkits, such as IoTABench [4], support testing of IoT use cases.

Three-dimensional (3D) Virtual Environments (VEs) allow users to explore virtual worlds without actually ‘being there.’ VEs have been used in a variety of applications including education, training, architectural walkthroughs, scientific visualization, art, and entertainment [7]. Although VE applications can be used for many different purposes, there are some fundamental interaction tasks used in VE applications.

3D Distributed Virtual Environment (DVE) is a software system in which multiple users interact with each other in real-time, even though those users may be located around the world [27]. DVEs are challenging applications due to several reasons. First, DVEs are multimedia-rich environments that include 3D graphical objects as well as audio and video streams. Second, users interact in real-time, such as in multi-player online games. Third, the number of users could reach hundreds or thousands in a single application who can compete or collaborate individually or in groups. That, in turn, implies the high variations in users’ requirements.

Experience with 3D user interfaces and interactions with the virtual worlds can be a starting point for interaction with the real world and architectural artifacts [16,19]. By providing a virtual world that models the architectural space and is connected to controls, sensors and actuators in the architectural space, we can use human-computer interaction based approaches and interaction techniques to support human-architecture interactions.

VE applications, such as Second Life [28], can be used to map or represent real-world sensors and embedded devices to virtual devices and objects. 3D user interface and visualization enable users to evaluate IoT computing applications using a VE to interactively test various IoT configurations [12,21,25].

3 Approach

We build on the preliminary work [13] that uses Virginia Tech’s Cube facility (http://www.icat.vt.edu/facilities/living_labs) that provides an enclosed space (dimensions $50 \times 40 \times 40$ feet) with real-time tracking and spatial audio capabilities. Several users can simultaneously move in the physical space and view the simulation results from different points of view. The physical location and orientation of a user is used in real-time to determine the corresponding view in the VE. The user can also observe other users and collaborate to change the simulation (Fig. 1).

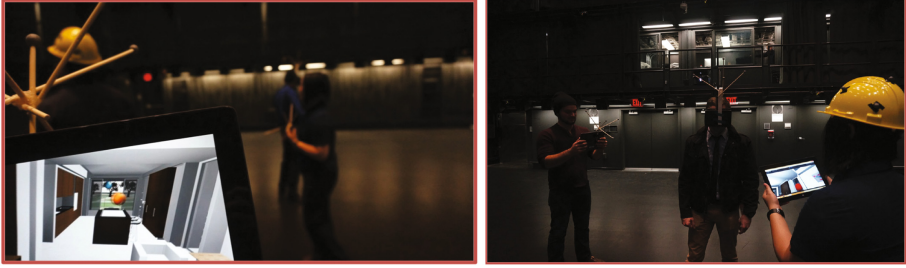


Fig. 1. The users navigating the virtual environment and the corresponding physical space (Cube).

The modeling framework (Fig. 2) provides a connection between the real world IoT-based collection of devices, sensors, and actuators and the corresponding virtual world representations [1]. Each sensor is modeled based on the set of the corresponding environmental and control parameters. The physical space and virtual space can synchronize their state by sharing the same IoT server. The IoT server records both sensory readings and user commands and keeps track of the most recent ones.

Besides mirroring a physical space state and providing a 3D virtual user interface to interact with it, the virtual space can serve two other purposes. First, it can simulate a non-existing physical space, where the virtual sensors will produce virtual streams of data and user commands will change the state of the virtual model. This allows for exploring different user interaction designs and techniques before implementation in the physical space. Second, it can playback a history of events stored at the IoT server, including both sensory readings and user commands. This can help exploring and studying different usage patterns by having the ability to reproduce them in VE settings.




Fig. 2. Mapping between the real and virtual worlds.

Constructing a VE to clone a physical IoT-based built space involves a set of steps. First, identifying the characteristics of the built space structure to be modeled in the VE. Second, identifying the list of devices to be modeled together with their physical characteristics. Third, identifying the set of sensors to be

modeled together with the corresponding equations for producing virtual data streams, as needed. Fourth, determining the set of virtual actuator representing the set of supported functionalities for each device (e.g. turn on/off). Fifth, finding how a user will trigger different actions upon virtual actuators. Finally, building the VE model and integrating it with the IoT server as well as providing the user interface for accepting user commands.

Creating the list of devices starts with the initial requirements gathering and preliminary design. The built environment under design is first analyzed from the energy perspective to identify components (appliances, consumer electronics, lights, etc.) that consume energy and should be monitored. Next, user interactions are studied to identify components (things) that a user interact with. In addition to the obvious choices (appliances, light switches) that already provide control capabilities, the designer can also select purely mechanical devices such as doors or drawers that then need to be equipped with embedded controllers, sensors and actuators.

Identifying, for each device, its physical characteristics for modeling purposes means that the control parameters (e.g., the warmer temperature or the extent of drawer opening) are used in equations modeling the real-world actions (e.g., energy consumption or the time needed for the warmer drawer to open). The individual sensors and actuators are registered with the IoT server so that the sensory readings and user commands are automatically recorded. Identifying the set of virtual sensors that will produce virtual data streams is based on the developed physical models for the individual devices to mimic the data generated by the real world sensors. Identifying the basic functionalities of the device and related actuators determines all possible user interactions. For example, a microwave can have a very elaborated user interface with many buttons/keys. Determining the virtual actions that the device can perform through the virtual actuator determines what are the visual manifestations of user actions (e.g., changes in position/orientation/size). Initially, only a basic subset of interactions is supported in the VE and then gradually expanded as the development progresses.

Representing the devices in the VE means that a sufficiently detailed model (geometry, parts, animation) has to be developed to accurately represent the user's actions and their consequences. Pressing a button in the VE to open a warmer drawer must result in the drawer moving with the appropriate speed that results from the user action in the VE.

Synchronizing the physical built-space and the VE using the same IoT server allows for creating a very flexible testbed. First, since the real world data is stored on the IoT server, the recorded data can be replayed and studied in the VE. Second, the developed physical models can be compared to the real world measurements and refined as needed. Third, once the VE is sufficiently accurate, the VE can be used to observe user interactions and behavior in a physical space (Cube) that matches the real world scale. Finally, the VE can be used, through the IoT server, to control the real world system.

In short, the system is developed in an iterative, bottom-up fashion where the IoT server is connected to both the real world devices and the virtual devices modeled in the VE. Modeling individual devices and composing them into a larger built environments produces reusable artifacts.

4 Case Study

Figure 5 shows the FutureHAUS prototype (<http://www.futurehaus.tech/>) consisting of several modules (a kitchen, a living room, a bathroom, an office, and a bedroom) that leverages the use of prefabricated architectural components for a home. The components can be accessed and controlled through a whole-house interface which manages and monitors appliance performance and energy use. The modules are instrumented with IoT-based electronic sensors and actuators that facilitates the use of the living space makes it more energy efficient. The goal is to integrate smart technologies into a prefabricated system while elevating the human experience of all household activities (e.g. cooking, and socializing).

The initial idea and the architectural conceptual (Fig. 3 top left) were used to create the corresponding VE (Fig. 3 center and Fig. 4). After exploring and refining the VE (Figs. 1 and 3 bottom left) the modules were constructed (Fig. 3 bottom right and Fig. 5).



Fig. 3. Kitchen and living room modules.

The kitchen has a set of devices (appliances), where each device has an IoT-based embedded controller controlling a set of sensors and/or actuators. The living room has controllable color lights and a sofa with pressure sensors. A local network connects the controllers to the IoT server. The controllers communicate with the IoT server by sending sensory readings periodically and receiving user



Fig. 4. VE for the built environment.



Fig. 5. Kitchen and living room modules.

commands, if any, before triggering the actuators to execute them. Applications and users can interact with the system by querying the readings or posting commands. The list of devices includes an inductive cooktop, a faucet, a dishwasher, a coffee maker, an oven, a warmer, a refrigerator, a microwave, LED lights, and a sofa.

4.1 Device Modeling

Modeling each device should consider different sets of parameters representing the physical characteristics of the device, the environmental parameters that can affect those characteristics, and the control parameters. For example, the warmer model considers its power consumption, temperature, and the length of its drawer. The environmental parameters include room temperature and humidity. The control parameters include the desired warmer temperature and the open percentage of the warmer's drawer. Different warmer characteristics were represented using basic physical models. For example, the energy consumption was initially approximated using Eq. 1.

$$E = \frac{T_{select} - T_0}{T_{max} - T_0} P_{max} t \quad (1)$$

where T_{select} is the desired warmer temperature, T_{max} is the maximum warmer temperature, T_0 is the initial warmer temperature, P_{max} is the maximum power and t is the elapsed time. This physical model (equation) can be further modified based on the real-world measurements.

Different appliances were individually designed and tested in terms of modeling and user interactions. Figure 6 shows a VE setup for a virtual warmer and a virtual refrigerator.




Fig. 6. FutureHAUS warmer and refrigerator appliances.

4.2 Light Control

Exploring different types of interaction techniques with IoT devices in a built environment, we have developed a smartphone application as well as a 3D user interface to control the LED-lights in the living room module as shown in Fig. 7. The smartphone application maps the layout of the living room lights to a 2D GUI allowing the user to control the color and intensity of individual lights through the smartphone. However, there are some limitations associated with the use of a smartphone for controlling devices. First, mapping devices that exist in a 3D space to a 2D GUI can be tricky for both users and developers. Second, relying on a smartphone to control devices places an explicit barrier between users and devices. Third, switching between different applications/screens to control different devices can be frustrating to the user. The 3D version of the user interface tries to address those limitations associated with the use of 2D GUIs by allowing for hands-free control of lights using a gestures-based interaction technique. The user can control different parameters of individual lights by simply pointing at the light source of interest using one hand and changing its color/intensity using the other hand. This allows for a seamless interaction while avoiding explicit intermediary devices.



Fig. 7. Left: The living room with LED-lights. **Right:** Gesture-based lights control.

5 Conclusion

The current implementation demonstrates the feasibility of the proposed approach. The users can explore the VE while navigating and interacting in the physical space that matches the built environment under study. One of the main challenges is how to quickly develop sufficiently accurate physical models of individual appliances and other devices. Our experience with the case study points out that for a well-defined application domain such built environments, it's feasible to build a library of models that can be reused.

We are continuing our work to model all of the house modules. The ongoing work proceeds in several directions, including preparing user studies and exploring additional application domains, such as manufacturing systems, and training.

Acknowledgments. This work was supported, in part, by a grant from Virginia Tech Institute for Creativity, Arts and Technology (ICAT).

References

1. Ahmed, H., Gračanin, D.: Context sensitive interaction interoperability for serious virtual worlds. In: Debattista, K., Dickey, M., Proença, A., Santos, L.P. (eds.) Proceedings of the 2nd International Conference on Games and Virtual Worlds for Serious Applications (VS GAMES 2010), pp. 159–166, 25–26 March 2010
2. Aipperspach, R., Hooker, B., Woodruff, A.: FEATURE: The heterogeneous home. *Interactions* **16**(1), 35–38 (2009)
3. Anderson, J., Rainie, L., Duggan, M.: The internet of things will thrive by 2025. Technical report, Pew Research Center, Washington, D.C. 20036, May 2014
4. Arlitt, M., Marwah, M., Bellala, G., Shah, A., Healey, J., Vandiver, B.: ioTAbench: an internet of things analytics benchmark. In: Proceedings of the 6th ACM/SPEC International Conference on Performance Engineering, pp. 133–144. ACM, New York (2015)

5. Armac, I., Retkowitz, D.: Simulation of smart environments. In: The Proceedings of the IEEE International Conference on Pervasive Services, pp. 322–331, 15–20 July 2007
6. Atzori, L., Iera, A., Morabito, G.: The internet of things: a survey. *Comput. Netw.* **54**(15), 2787–2805 (2010)
7. Bowman, D.A., Kruijff, E., LaViola, J.J., Poupyrev, I.: 3D User Interfaces: Theory and Practice. Addison-Wesley, Boston (2005)
8. Brambilla, G., Picone, M., Cirani, S., Amoretti, M., Zanichelli, F.: A simulation platform for large-scale internet of things scenarios in urban environments. In: Proceedings of the First International Conference on IoT in Urban Space, pp. 50–55. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), Brussels (2014)
9. Crabtree, A., Rodden, T.: Domestic routines and design for the home. *Comput. Support. Coop. Work* **13**(2), 191–220 (2004)
10. daCosta, F.: Rethinking the Internet of Things: A Scalable Approach to Connecting Everything. Apress L. P. (2013)
11. Dourish, P.: Re-space-ing place: “place” and “space” ten years on. In: Proceedings of the 2006 20th Anniversary Conference on Computer Supported Cooperative Work (CSCW 2006), pp. 299–308, 4–8 November 2006
12. Eastman, C., Teicholz, P., Sacks, R., Liston, K.: BIM Handbook: A Guide to Building Information Modeling for Owners, Managers, Architects, Engineers, Contractors, and Fabricators. Wiley, New Jersey (2008)
13. Gračanin, D., Eck II, T., Silverman, R., Heivilin, A., Meacham, S.: An approach to embodied interactive visual steering: bridging simulated and real worlds. In: Proceedings of the 2014 Winter Simulation Conference (WSC), pp. 4073–4074, December 2014
14. Gračanin, D., Matković, K., Wheeler, J.: An approach to modeling internet of things based smart built environments. In: Proceedings of the 2015 Winter Simulation Conference (WSC), pp. 3208–3209, 6–9 December 2015
15. Gračanin, D., McCrickard, D.S., Billingsley, A., Cooper, R., Gatling, T., Irvin-Williams, E.J., Osborne, F., Doswell, F.: Mobile interfaces for better living: supporting awareness in a smart home environment. In: Stephanidis, C. (ed.) UAHCI 2011. LNCS, vol. 6767, pp. 163–172. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-21666-4_19](https://doi.org/10.1007/978-3-642-21666-4_19)
16. Gračanin, D., Zhang, X.: Caffeneve: A podcasting-capable framework for creating flexible and extensible 3D applications. *ACM Computers in Entertainment* **11**(1), 1–30 (2013)
17. Handosa, M., Gračanin, D., Elmongui, H.G., Ciambone, A.: Painting with light: gesture based light control in architectural settings. In: Proceedings of the 2017 IEEE Symposium on 3D User Interfaces, 18–19 March 2017
18. Harrison, S., Dourish, P.: Re-place-ing space: the roles of place and space in collaborative systems. In: Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work (CSCW 1996), pp. 67–76. ACM, New York (1996)
19. Lazem, S., Gračanin, D.: Social traps in second life. In: Debattista, K., Dickey, M., Proença, A., Santos, L.P. (eds.) Proceedings of the 2nd International Conference on Games and Virtual Worlds for Serious Applications (VS GAMES 2010), pp. 133–140, 25–26 March 2010
20. Lea, R., Blackstock, M.: Smart cities: an IoT-centric approach. In: Proceedings of the 2014 International Workshop on Web Intelligence and Smart Sensing, p. 12:1–12:2. ACM, New York (2014)

21. McGlinn, K., Hederman, L., Lewis, D.: SimCon: a context simulator for supporting evaluation of smart building applications when faced with uncertainty. *Pervasive Mob. Comput.* **12**, 139–159 (2014)
22. Mukhopadhyay, S.C. (ed.): *Internet of Things: Challenges and Opportunities. Smart Sensors, Measurement and Instrumentation*, vol. 9. Springer, Cham (2014)
23. Nansen, B., van Ryn, L., Vetere, F., Robertson, T., Brereton, M., Douish, P.: An internet of social things. In: *Proceedings of the 26th Australian Computer-Human Interaction Conference on Designing Futures: The Future of Design*, pp. 87–96. ACM, New York (2014)
24. Nissenbaum, H.F.: Privacy as contextual integrity. *Wash. Law Rev.* **79**(1), 119–158 (2004)
25. Prendinger, H., Brandherm, B., Ullrich, S.: A simulation framework for sensor-based systems in second life. *Presence Teleoperators Virtual Environ.* **18**(6), 468–477 (2009)
26. Rosenkranz, P., Wählich, M., Baccelli, E., Ortmann, L.: A distributed test system architecture for open-source IoT software. In: *Proceedings of the 2015 Workshop on IoT Challenges in Mobile and Industrial Systems*, pp. 43–48. ACM, New York (2015)
27. Singhal, S., Zyda, M.: *Networked Virtual Environments: Design and Implementation*. SIGGRAPH Series. ACM Press, Addison Wesley, Reading (1999)
28. Varvello, M., Picconi, F., Diot, C., Biersack, E.: Is there life in second life? In: *Proceedings of the 2008 ACM CoNEXT Conference (CoNEXT 2008)*, pp. 1–12. ACM, New York (2008)