

Chapter 16

Building Internet of Things-Enabled Digital Twins and Intelligent Applications Using a Real-time Linked Dataspace



Edward Curry, Wassim Derguech, Souleiman Hasan,
Christos Kouroupetroglou, Umair ul Hassan, and Willem Fabritius

Keywords Decision support · Internet of Things · Intelligent systems · Smart environments · Energy management · Water management · Dataspaces

16.1 Introduction

Smart environments have emerged in the form of smart cities, smart buildings, smart energy, smart water, and smart mobility. A key challenge in delivering smart environments is creating intelligent applications for end-users using the new digital infrastructures within the environment. In this chapter, we reflect on the experience of developing Internet of Things-based digital twins and intelligent applications within five different smart environments from an airport to a school. The goal has been to engage users within Internet of Things (IoT)-enabled smart environments to increase water and energy awareness, management, and conservation. The chapter covers the role of a Real-time Linked Dataspace to enable the creation of digital twins, and an evaluation of intelligent applications.

The chapter starts in Sect. 16.2 with a description of Digital Twins and the role that Boyd's OODA Loop can play in their realisation. Creating digital twins and intelligent application using a Real-time Linked Dataspace is detailed in Sect. 16.3. The results from the smart energy and water pilots are detailed in Sect. 16.4. Section 16.5 discusses experiences and lessons learnt, and the chapter concludes in Sect. 16.6.

16.2 Digital Twins and Intelligent Applications with a Real-time Linked Dataspace

Driven by the adoption of the Internet of Things (IoT), smart environments are enabling data-driven intelligent systems that are transforming our everyday world, from the digitisation of traditional infrastructure (smart energy, water, and mobility), the revolution of industrial sectors (smart autonomous cyber-physical systems, autonomous vehicles, and Industry 4.0), to changes in how our society operates (smart government and cities). To support the interconnection of intelligent systems in the data ecosystem that surrounds a smart environment, there is a need to enable the sharing of data among systems.

16.2.1 Real-time Linked Dataspaces

A data platform can provide a clear framework to support the sharing of data among a group of intelligent systems within a smart environment [1] (see Chap. 2). In this book, we advocate the use of the dataspace paradigm within the design of data platforms to enable data ecosystems for intelligent systems.

A dataspace is an emerging approach to data management which recognises that in large-scale integration scenarios, involving thousands of data sources, it is difficult and expensive to obtain an upfront unifying schema across all sources [2]. Within dataspaces, datasets *co-exist* but are not necessarily fully integrated or homogeneous in their schematics and semantics. Instead, data is integrated on an *as-needed* basis with the labour-intensive aspects of data integration postponed until they are required. Dataspaces reduce the initial effort required to set up data integration by relying on automatic matching and mapping generation techniques. This results in a loosely integrated set of data sources. When tighter semantic integration is required, it can be achieved in an incremental *pay-as-you-go* fashion by detailed mappings among the required data sources.

We have created the Real-time Linked Dataspace (RLD) (see Chap. 4) as a data platform for intelligent systems within smart environments. The RLD combines the pay-as-you-go paradigm of dataspaces with linked data, knowledge graphs, and real-time stream and event processing capabilities to support a large-scale distributed heterogeneous collection of streams, events, and data sources [4].

16.2.2 Digital Twins

Within the business community [32], the metaphor of a “Digital Twin” is gaining popularity as a way to explain the potential of IoT-based assets and smart environments. A digital twin refers to a digital replica of physical assets (car), processes

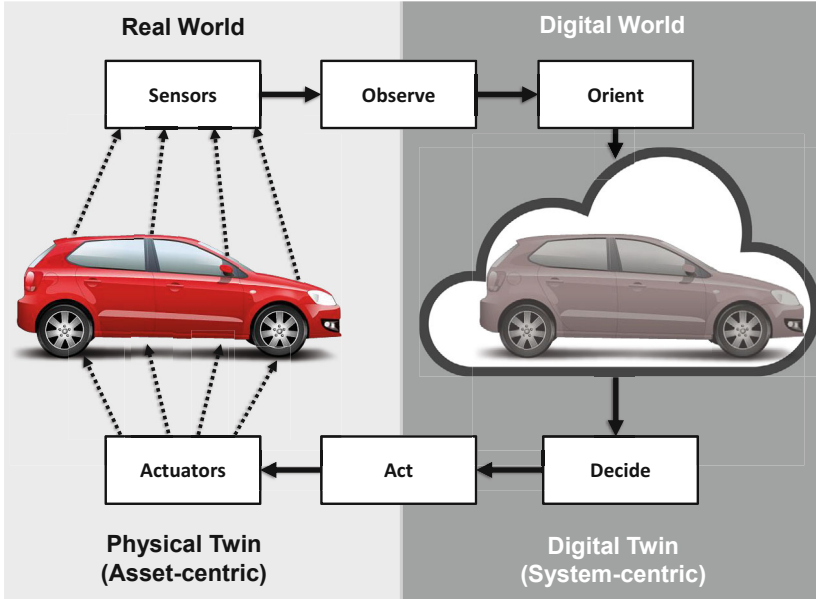


Fig. 16.1 A digital twin provides a digital representation which can be analysed to optimise the operation of the “physical twin”

(value-chain), system, or physical environment (building). The digital representation provided by the digital twin can be analysed to optimise the operation of the “physical twin”. The digital twin provides a digital representation (Fig. 16.1) (i.e. simulation model, data-driven model) that updates and changes as the physical twin changes. Digital twins can provide digital representations ranging from human organs such as the heart and lungs to aircraft engines and city-scale twins. For example, the SmartSantander smart city project has deployed tens of thousands of Internet-connected sensor devices in large cities across Europe [33]. The sensing capabilities of these devices are wide-ranging, including solar radiation, wind speed and direction, temperature, water flow, noise, traffic, public transport, rainfall, parking, and others. The devices provide a digital representation of the state of the real world, in the case of SmartSantander a digital representation of the city, enabling visibility into processes and operations of the city that can be analysed and optimised.

With the use of advanced analytics and artificial intelligence techniques, the digital twin can learn the optimal operating conditions of the physical twin and optimise the physical twins’ operations in areas such as performance, maintenance, and user experience. One of the most promising outputs from such an analysis is the possibility to find root-causes of potential anomalies which can happen (prediction) and improve the physical process (innovation).

Digital twins are a sophisticated example of a cyber-physical system which is constructed from multiple sources of data including real-time IoT sensors, historical


sensor data, traditional information systems, and human-in-the-loop input from human operators and domain and industrial experts. The core of a digital twin requires a holistic and systematic approach to data management and decision-making; at the heart of a digital twin is an *OODA Loop*.

16.2.3 The OODA Loop

John Boyd hypothesised that individuals and organisations undergo a continuous cycle of interaction with their environment. Boyd developed the “OODA Loop” [341] as a decision process by which an entity (either an individual or an organisation) reacts to an event by breaking the decision cycle down to four interrelated and overlapping processes through which one cycles continuously: *Observe*, *Orient*, *Decide*, and *Act* (OODA). Boyd initially applied the OODA Loop to military operations, and it was later applied to enterprise operations. More recently, it has been considered as an approach for processing observations within cyber-physical systems [14]. In this latter context, we apply the OODA Loop as a high-level design guide for intelligent energy and water systems within smart environments. As illustrated in Fig. 16.2, the four OODA processes applied to an intelligent application within a smart environment are:

- *Observation*: The gathering of data from the smart environment to understand its state.
- *Orientation*: The analysis and synthesis of data to form an assessment of the circumstances within the smart environment. Moving from data to information, knowledge, and insights.
- *Decision*: Consideration of the options to determine an appropriate course of action. The goal is to optimise the operation of the smart environment. The use of predictive modelling can play a significant role here.
- *Action*: The physical execution of decisions via actuation (both automated and human). Once the result of the action is observed, the loop starts over.

Fig. 16.2 Boyd’s OODA Loop [341] applied to intelligent applications within a smart environment



16.3 Enabling OODA for Digital Twins and Intelligent Applications

As shown in Fig. 16.3, and detailed in [4], we use the OODA Loop to align the different development phases of digital twins and intelligent applications with the relevant RLD support services.

16.3.1 Observation

The RLD support services facilitate the observation phase by minimising the required effort for a data source to join the RLD. Support services such as the Catalog, Access Control Service (see Chap. 6), and the Search and Query Service (Chap. 10) are the primary services that enable the collection of data sources and IoT data and the maintainability of its associated metadata. The incremental approach of the RLD made it easier to gradually improve the collection of observations from the smart environment by adding a new sensor, thing, or dataset to the RLD. The 5 star

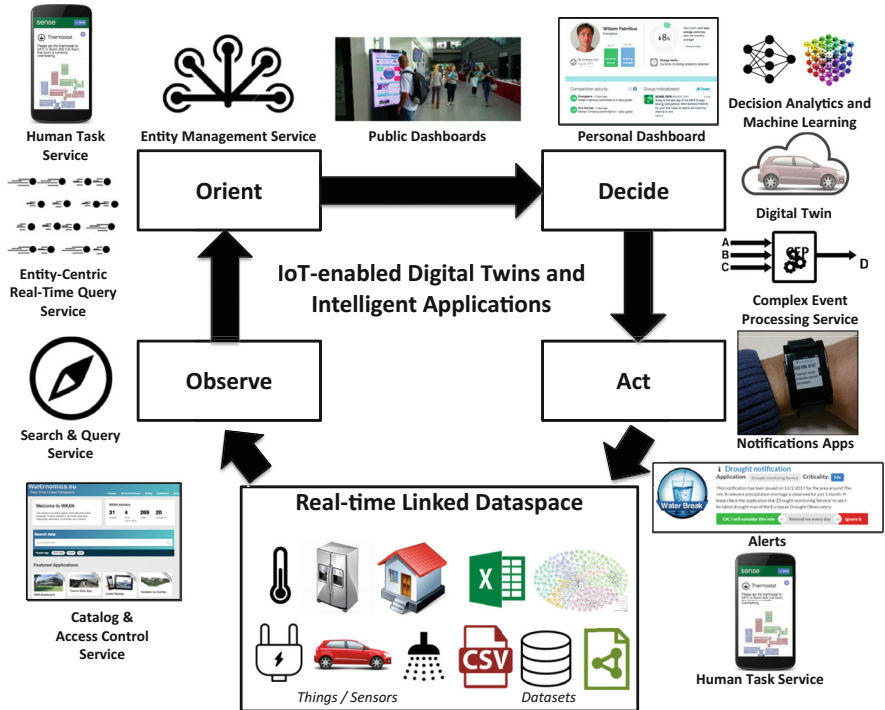


Fig. 16.3 Role of RLD and its support services across the phases of the OODA Loop [4]

pay-as-you-go model for data management (see Chap. 4 for more details) was useful for specifying and planning the level of service needed for each data source.

The human task service enables the engagement of users in maintaining a high-quality catalog of managed entities. Active participation of users in a smart environment improves their engagement and sense of ownership while supporting a higher accuracy of data maintained by the dataspace. In one of our pilot deployments we noticed a direct benefit of using the human task service for the collaborative management of the entities in the environment to provide a more accurate and rich understanding of the environment's state [256].

16.3.2 Orientation

The primary objective of the orientation phase is to support situational awareness of the smart environment. The real-time query services (see Chap. 10) enable users to understand the current and historical state of the smart environment. The Entity Management Service (EMS) builds awareness regarding the entities in the environment through entity linking and enrichment (which can be supported by the Human Task Service). Together with the real-time query services, the EMS provides entity-centric views of the smart environment and reduces the overall effort to integrate entity data from different real-time streams and contextual data sources.

Within all the pilots, a key goal is to increase the visibility, understanding, and awareness of energy and water use. Using the RLD support services, we can build dashboards to provide situational awareness for users with targeted information on energy and water consumption. Within the different pilots, this is manifested in a variety of ways and at different time frames, from informing the residents in their smart home as they live their day, supporting the detailed analysis required by building managers and operational staff, to brief encounters with “frequent-flyer” passengers as they pass through the airport. User orientation in the pilots was driven by public displays, interactive touchscreen displays and tablet applications (see Fig. 16.4). These user interfaces communicate current and historical energy and water usage within the environment, convey information about the importance of energy and water, tips on how to improve consumption, and games to calculate the users' footprint in real-time. The displays are also personalised to target different users by using appropriate metaphors to communicate relevant messages to them. The intelligent applications in the orientation phase make extensive use of real-time, historical, and contextual data sources to enhance the user experience (see Chap. 17).

16.3.3 Decision

Once users have built a certain level of awareness regarding the energy or water consumption of their environment, they can use their expertise to start taking

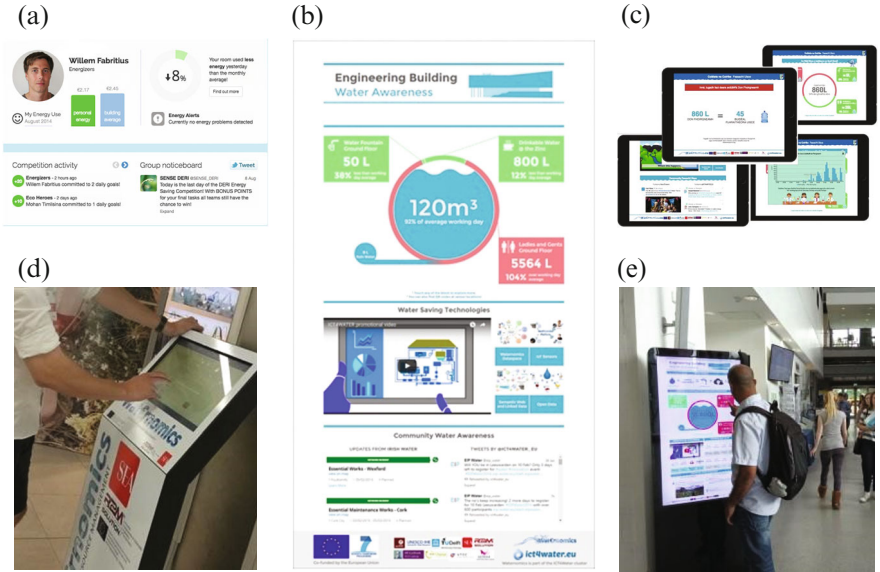


Fig. 16.4 Public interactive displays and personalised dashboards: (a) Smart office, (b) Smart building, (c) Smart school, (d) Smart airport, (e) Smart building

decisions towards more sustainable behaviours. In the decision phase, a critical aspect of the dashboards is to provide users with targeted information on usage, goal setting, targets for conservation, and tips to improve their consumption behaviour. This is where decision-making takes place. For example, managers can define consumption thresholds to serve as sources of “alerts”, notifying them of excessive usage, goals attained, or the detection of a possible fault (e.g. Complex Event Processing Service, see Chap. 11). Developing decision support applications relies on the entity-centric real-time query service to analyse data from the environment, interpret it, and decide on the appropriate course of action.

A specific example of decision support is the Water Retention Time Observer application (see Fig. 16.5) that determines the amount of time drinking water resides in water pipes and creates alerts in case of potential issues. In public spaces, drinking water quality is a significant concern for building managers: is the water safe to drink? Currently, this can be managed by selecting a popular location to place the drinking water fountains to ensure people are always using them, thus ensuring that freshwater is always flowing through the pipes. However, in some public buildings, drinking water fountains can remain unused during long holidays and weekends. Consequently, drinking water can reside for extended periods in the pipes. In this context, the water retention time observer can assist building managers by providing timely notifications regarding low water quality in drinking water pipes. This is achieved by creating a simple digital twin of the water network to detect inactivity in specific measurement points in the water network and sending a notification if

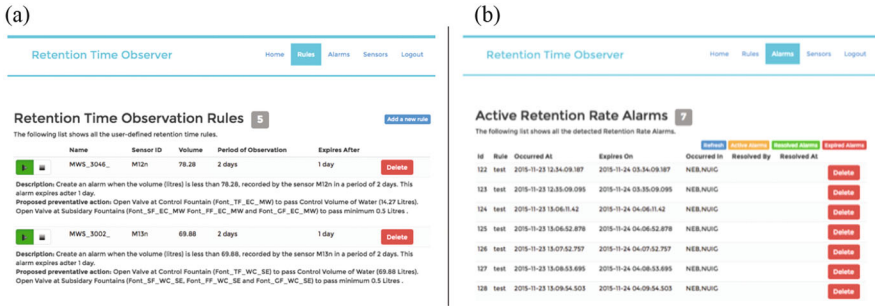


Fig. 16.5 Water retention time observer: (a) Observation rules. (b) Active alarms

stagnant water is detected. Within this digital twin, we aimed to enable notification to attract users' attention only when necessary. This was a key lesson from our work to enhance user experience, which is discussed further in Chap. 17.

16.3.4 Action

Intelligent applications in the action phase of the OODA Loop help users in smart environments meet their goals for energy and water consumption by taking appropriate actions. The complex event processing service of the RLD is used to express these goals as a set of rules which can generate alerts and suggested preventive actions. Actions are then communicated to the users in the smart environment using an appropriate mean of communication: emails, notifications on the dashboards, messages on smart devices, and human tasks.

The occupants of the environment can participate in taking energy or water saving actions. In the smart building pilot, we implemented a collective energy management system where the RLD was used for the identification of energy-saving tasks. The tasks were routed to the building occupants using the human task service (see Chap. 9) to take energy conservation actions such as turning off the light in empty rooms or closing a window when an air conditioner is on in a room. Figure 16.6 shows an example of these "Citizen" actuation tasks.

The role of a building manager is a demanding one that often has personnel working in the field. An anytime-anywhere notification mechanism was needed for managers. To minimise the search friction between actionable information and users, a well-designed notifications system is needed. The wearable info-centre application was developed to enable notification through the wearable technology for high-priority alerts. Figure 16.7 shows an example notification using the wearable info-centre.

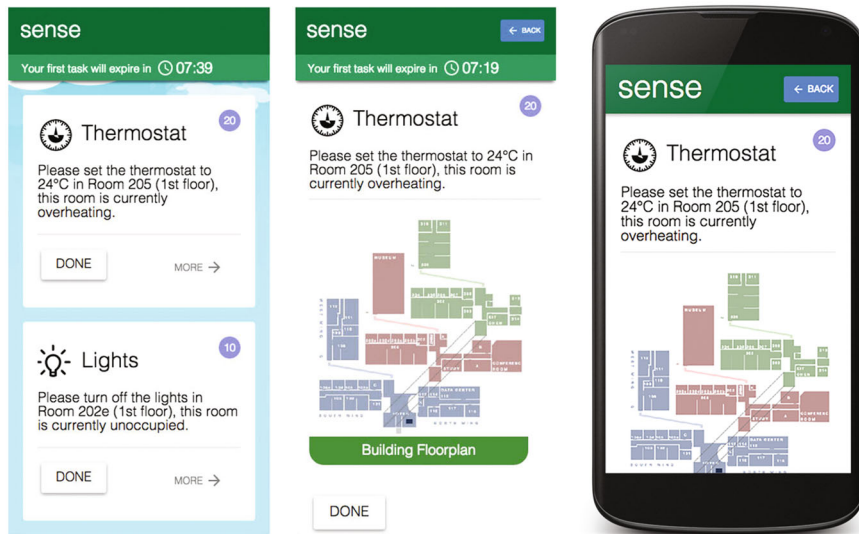


Fig. 16.6 Example of citizen actuation tasks within the smart environment



Fig. 16.7 Example notifications within the smart environment

16.4 Smart Energy and Water Pilots

This section presents the results and insights gained from deploying the RLD and intelligent applications in the smart environments described in Chap. 14. Each pilot followed a similar methodology for design, deployment, and evaluation [63]. In this section, we detail the energy and water savings achieved in the pilots, the performance of the human task services in engaging users to save energy, and a set of experiences and lessons learnt from deploying the RLD in the pilots.

16.4.1 Energy and Water Savings

During the initial period of the pilots, energy and water metering data was collected from existing monitoring systems to establish baselines for consumption across the pilots. During the control period, the users within the pilots had access to the data generated by the metering infrastructure system through traditional information systems (e.g. building management systems, and basic public dashboards within the airport, office building, and school). The data collection period for each pilot spanned between 6 and 16 months, which also included a range of user interventions such as pre-surveys, focus groups, interviews, and feedback cycles. The RLD was used to develop intelligent energy and water systems and decision support analytics across the pilot smart environments. Table 16.1 details the characteristics of the pilots during the study period, the number of events generated in the environment, the number of intelligent applications/twins deployed, and savings achieved in terms of energy and water. In terms of energy and water savings, the RLD supports these impacts in three fundamental ways:

- Connecting data across silos provided “big picture” entity-centric views of the resource consumption within the smart environments. These views made it easier for the users within the smart environments (e.g. building managers) to identify waste and efficiency opportunities as the data produced within the environment was structured and organised around real-world entities. Entity-centric views were the basis of the digital twins created.
- The pay-as-you-go approach was useful for building the business case and getting “buy-in” from users by enabling quick wins to demonstrate the benefit of the approach. These early wins that demonstrated energy and water savings encouraged non-technical business users to engage with the project and system more

Table 16.1 Summary of the impact of intelligent energy and water systems in smart environments [4]

| Pilot site | Location | Study period | Events per year | Intelligent applications/twins deployed | Actual savings measured | Estimated annual savings |
|---------------|----------------------------------|--------------|-----------------|-----------------------------------------|---------------------------------------------------|-------------------------------------------------------|
| Smart Airport | Linate Airport, Italy | 10 months | ~11.5 million | 8 | 2954 m ³ 3013 kg CO ₂ | 54,000 m ³ 55,080 kg CO ₂ |
| Smart Office | Insight, Ireland | 6 months | ~8 million | 4 | 24% energy reduction | – |
| Smart Homes | Thermi, Greece | 16 months | ~2.3 million | 11 | 30% water reduction | – |
| Mixed Use | Engineering Building, NUI Galway | 16 months | ~36 million | 8 | 174 m ³ 177 kg CO ₂ | 8089 m ³ 8251 kg CO ₂ |
| Smart School | Coláiste na Coiribe, Ireland | 12 months | ~1 million | 5 | 2179 m ³ 2223 kg CO ₂ | 9306 m ³ 9492 kg CO ₂ |

actively. The project team could build a business case around intelligent applications/digital twins and decision support tools that would reduce resource usage and its associated economic costs. The savings identified can be used to justify the necessary investment needed in data integration.

- The RLD enabled highly specialised decision analytics and digital twins that provided action and notification alerts for each of the pilot smart environments, including leak detection, fault detection, and abnormal usage patterns. These alerts and notifications were crucial for building managers and operational staff who do not have time to study and analyse the data generated in the smart environment.

16.4.2 Human Task Service Evaluation

To determine the effectiveness of the human task service of the RLD within a smart environment, we performed two types of human tasks in the smart office pilot: (1) an entity enrichment data management task, and (2) a citizen actuation task for energy savings.

16.4.2.1 Human Task for Entity Enrichment

This experiment focuses on a data management task that requires the user to enrich the description of an entity by collecting location information on sensors within the smart environment. Accurate location data is needed by the energy management system to make appropriate recommendations about temperature control and energy usage in the monitored building. We do not assume this metadata on sensor locations, and room characteristics are available at the start of the experiment. This situation simulates the case when it is difficult to gather all metadata upfront, or the metadata becomes invalid due to changes in the environment. The objective of the experiment was to use human tasks to enrich the sensor entities in the RLD with the support of building occupants.

The occupants of the building were contacted through email to participate in the experiment. If they consented, they were asked to look for sensors around them in the building and to scan a QR code on the sensors using their mobile phones. This would resolve the URL associated with the QR code in a web browser, where they would then be asked to perform a relevant task. This action connects the user to the human task service within the RLD and enables the linkage between human tasks and physical sensors. Once the participant submits the location of the sensor, additional tasks are pushed to them to collect further metadata about the surrounding environment. Three tasks collect information about lights, heaters, and windows in the room. The collected data is then used to enrich the description of the sensor and room entities in the EMS.

The evaluation is based on the comparison of occupant-contributed metadata versus gold-standard data. The gold-standard data was created manually by studying

Table 16.2 Description and results of entity enrichment tasks against the gold-standard data [256]

| Task | Description | Assignment method | Accuracy |
|-----------------|------------------------------------------------------------------------------------------------|-------------------------------------|----------|
| Sensor location | This task requires participants to specify the location of the sensor. | Task pull based on the QR code. | 85.71% |
| Room lights | This task asks participants to specify the number of fluorescent lights installed in the room. | Task push based on person location. | 100% |
| Room heaters | This task asks participants to specify the number of heaters in the room. | Task push based on person location. | 83.33% |
| Room windows | This task asks participants to specify the number of windows in the room. | Task push based on person location. | 100% |

the physical space. Table 16.2 shows the accuracy of data submitted by occupants of the building within 5 h of sending the invitation email to the building occupants. The reported accuracy is based on the data submitted by the first few participants for each sensor and room. The human task service achieved more than 80% accuracy in describing the sensors and rooms within 5 h. The accuracy could be increased if the results from multiple users are used to verify the accuracy of the contributions.

16.4.2.2 Human Actuation for Energy Savings

The second evaluation of the human task service focuses on tasks for humans to save energy by performing citizen actuation [265]. When the energy management system detects an abnormal energy usage in a room in the building (i.e. high energy use for both the time of day and room status [booked for a meeting or not]), a notification via Twitter is sent to an appropriate user to request the user to check on the issue. This is the actuation request. Often the cause of the energy consumption abnormality is due to a light or equipment (e.g. projector or air conditioning) being left on in an empty room. This interaction between the user and the human task service, together with the relevant energy sensor readings, is illustrated in Fig. 16.8.

Within the smart office pilot, we collected data over a 32-week control period. Weekend data was removed from the experiment, as the users would not be on site. Fifteen volunteers were selected for the experiment. For each request, one volunteer was chosen at random to receive the request. The results of the experiment are illustrated in Fig. 16.9 with the max, min, median, and average energy consumption for the control and actuation days. Overall, the results show that the energy usage on average declined compared to the control during the weeks (experimental weeks) users received actuation requests and completed the actions of turning off electrical components. The average saving was 0.503 kWh, when compared to the average energy used in the control weeks of 1.93 kWh. This equates to a decrease in energy usage by 26%. Each actuation week's energy usage was equal to or lower in value to the lowest control week apart from 1 week (which, compared to the other control weeks, was lower).



Fig. 16.8 Citizen actuation task using the human task service of the RLD [220]

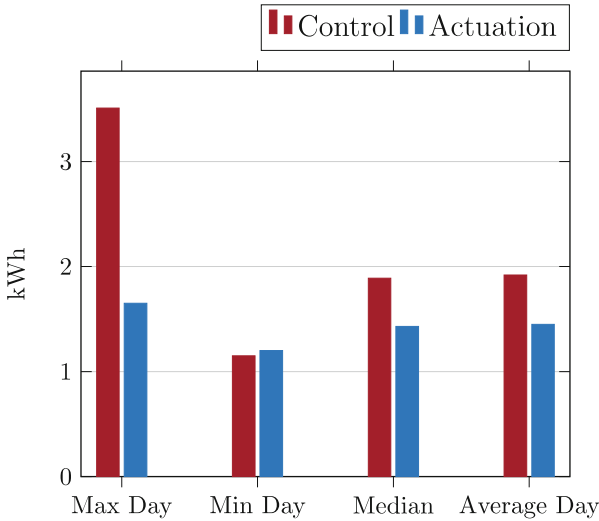


Fig. 16.9 Daily energy usage—the average, max, min, and median of the control period and actuation period [220]

16.5 Experiences and Lessons Learnt

Based on a reflection of our experience of using the RLD in the pilot environments, the following lessons were identified as key learnings to inform the design of future digital twins and intelligent applications for smart environments using the RLD [4].

Developer Education Across the pilots, we worked with a diverse set of development teams with different backgrounds, from embedded devices to web front-ends. The dataspace concept was new to most of them, and they were accustomed to working in an environment where they have full-control with the expectation of exact results. Also, the store-and-query culture is more common among developers and users. The processing of data on-the-fly and detecting only data that is of interest in real time, without storage in most cases, can be challenging (aka. event processing) for some developers to understand. Embracing the dataspace took time and required us to demonstrate both the benefits and limitations of the paradigm. Developer education was critical to the adoption of the dataspace. Workshops and tutorials held at pilot sites proved to be an effective mechanism of engaging developers to educate them on the capabilities of the platform and the dataspace data management approach.

Incremental Data Management Can Support Agile Software Development The project teams for each pilot operated using an agile software development methodology. The incremental approach of the dataspace and the use of the event-based paradigm were efficient during the design and development phase. The RLD enabled the teams to work at a pace suitable to the stakeholders and data owners involved. The RLD allowed the project team to include new data sources during a development iteration, or to increase the level of integration of an existing source. The decoupling achieved via the catalog and the use of events and streams removed dependencies between parties. It enabled the project teams to work with participants in the pilots in an incremental manner where we could quickly demonstrate value with a low upfront investment in data integration. As the pilots progressed, more and more data became available in the RLD enabling the creation of sophisticated data-intensive intelligent applications, digital twins, and analytics.

Build the Business Case for Data-Driven Innovation It is important to clearly articulate the business case for the RLD to justify the necessary investment in data infrastructure. Within our pilots, we discovered a strong business case for data-driven innovation by justifying the investment based on the resulting cost savings achieved due to improving resource efficiency (e.g. energy and water savings). A key challenge was to bring together the different stakeholders in the pilots to support and deliver the project. For example, the IT organisations had the data, but the savings resulting from the system benefit the operations teams of the organisations (e.g. water and energy). Thus, operations have a clear motivation to invest, but IT does not. By bringing these stakeholders together, we were able to build a holistic business case.

Integration with Legacy Data Is a Significant Cost in Smart Environments While sensors and connected devices are an essential source of data in a smart environment, they are not the only source of data necessary to make an environment “smart”. In our pilots, a considerable number of different legacy data sources needed to be integrated to collect the information necessary to make informed and intelligent decisions. While the RLD provided an effective incremental approach that integrates legacy data at a minimum cost, it is not a silver bullet to data integration costs in smart environments and the cost of integrating with legacy data should not be underestimated. This is of relevance within enterprise settings where the non-technical challenges (e.g. sharing data among departments) can be as significant as the technical ones. See Chap. 2 for further discussion on these challenges.

The 5 Star Pay-As-You-Go Model Simplified Communication with Non-technical Users The 5 star pay-as-you-go model for data management (see Chap. 4 for more details) was particularly useful regarding communicating both enhanced functionality and the additional costs of tighter integration with the RLD support services. Within the pilots, it was common to integrate data to the 3 star level on most services. The investment to bring a source to 4 and 5 stars was only made for core datasets within a pilot, and not for each service. Interestingly, many datasets that were initially identified in the early design phases as of high importance (e.g. sensor specifications, detailed infrastructure schematics) remained at the 1 star level as they were not needed by the final applications developed. This resulted in significant savings by avoiding unnecessary integration costs. Within the commercial pilots where more legacy data was available, the 5 star model supported the articulation of the business case for the investments necessary to include data sources and the level of their integration in the dataspace.

A Secure Canonical Source for Entity Data Simplifies Application Development Programmable access to the catalog by enabling queries over the machine-readable metadata and entities was crucial to facilitate application development in the dataspace. The role of the catalog and EMS as a canonical source for identifiers for entities was critical to managing the entities in the dataspace. Demonstrating the secure query capability of the access control service was essential to get “buy-in” and build trust with the pilot data owners. For example, the sensor data within the domestic pilot was sensitive, and we needed to assure the residents it was secured so that only privileged users could access their sensor data.

Data Quality with Things and Sensors Is Challenging in an Operational Environment Data quality challenges are further complicated as participating data sources, and things within the RLD are not under its full control. Data quality issues included incorrect file formats, incorrect timestamps, unusual sensor usage values, multiple and conflicting values, and missing data. Specifically, concerning the timestamps, the different time zones of pilot sites in different countries posed a challenge, as well as the time changes due to Daylight Saving Time. Keeping raw data where possible, allowed these issues to be addressed and for the analysis to be rerun with the data quality issues resolved. Finally, physical access to the

infrastructure can be a significant challenge within operational environments. Within the Linate airport pilot, the infrastructure was often underground within secured parts of the airport. One cannot rely on having physical access to restart or update infrastructure. As a result, the system design must be fault tolerant and adapt to operating conditions.

Working with Three Pipelines Adds Overhead The complexity of maintaining the RLDs' three different processing pipelines (the batch, real time, entity layers of the entity-centric query services, see Chap. 10) was challenging concerning the engineering and operational overhead involved. Diagnosing problems and faults required the workflow of all pipelines to be checked for issues, and this can increase the time needed to resolve a problem. A possible future direction is to look at end-to-end exactly-once stream processing technologies (Kappa Architectures). However, the highly decentralised nature of a smart environment and the lack of end-to-end control within dataspace may not be suitable to the additional coordination/control overhead of exactly-once stream processing approaches. This is an area of future work.

16.6 Summary

In this chapter, we reflect on the experience of developing different IoT-based intelligent applications and digital twins within five different smart environments, from Airport to Schools, where the goal has been to engage users within IoT-based intelligent systems to increase water and energy awareness, management, and conservation. The overall design philosophy has been guided using Boyd's "OODA Loop" for decision-making. The chapter detailed the role of a Real-time Linked Dataspace and its support services to enable the creation of intelligent applications and digital twins. The effectiveness of intelligent applications and digital twins within the pilots is evaluated to determine the level of savings achievable. The evaluation identified significant savings within the evaluation period at all the pilot sites. Finally, we reflected on our experiences using the RLD and captured these as a set of lessons learnt.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

