

ScaleUp: middleware for intelligent environments

Daniyal Alghazzawi ^{Corresp., 1}, **Ghadah Aldabbagh** ¹, **Abdullah Saad AL-Malaise AL-Ghamdi** ¹

¹ Faculty of Computing and Information Technology, King Abdulaziz University, Jeddah, Saudi Arabia

Corresponding Author: Daniyal Alghazzawi
Email address: dghazzawi@kau.edu.sa

The development of the internet of things (IoT) expands to an ultra-large-scale, which provides numerous services across different domains and environments. The use of middleware eases application development by providing the necessary functional capability. This paper presents a new form of middleware for controlling smart devices installed in an intelligent environment. This new form of middleware functioned seamlessly with any manufacturer API or bespoke controller program. It acts as an all-encompassing top layer of middleware in an intelligent environment control system capable of handling numerous different types of devices simultaneously. This protected de-synchronization of data stored in clone devices. It showed that in this middleware, the clone devices were regularly synchronized with their original master such as locally stored representations were continuously updated with the known true state values.

ScaleUp: Middleware for Intelligent Environments

Running Title: Middleware for Intelligent Environments

Daniyal Alghazzawi (corresponding author)

(IEEE Senior member): Faculty of Computing and Information Technology, King

Abdulaziz University, Jeddah, Saudi Arabia; dghazzawi@kau.edu.sa

Dr. Ghadah Aldabagh

Computer Science Department, FCIT, King Abdulaziz University, Saudi Arabia

Dr. Abdullah Alghamdi

(IEEE Senior member): Faculty of Computing and Information Technology, King

Abdulaziz University, Jeddah, Saudi Arabia

Abstract

The development of the internet of things (IoT) expands to an ultra-large-scale, which provides numerous services across different domains and environments. The use of middleware eases application development by providing the necessary functional capability. This paper presents a

new form of middleware for controlling smart devices installed in an intelligent environment. This new form of middleware functioned seamlessly with any manufacturer API or bespoke controller program. It acts as an all-encompassing top layer of middleware in an intelligent environment control system capable of handling numerous different types of devices simultaneously. This protected de-synchronization of data stored in clone devices. It showed that in this middleware, the clone devices were regularly synchronized with their original master such as locally stored representations were continuously updated with the known true state values.

Keywords: Intelligent Environment Control Systems, Middleware, Pervasive Computing, Smart Devices

1. Introduction

Architectures for intelligent environments typically require some form of necessary middleware layer, enabling installed hardware (smart devices) and software, (agents or controller programs), to link and operate together. Various types of middleware are created to perform this crucial role, such as Universal Plug and Play (UPnP), Web Services, and bespoke mechanisms written by device manufacturers [1].

Middleware may vary in terms of functionality, with different strengths and weaknesses. These strengths and weaknesses largely depend on their deployment conditions. This is challenging for the programmers that adopt unenviable measures when designing control systems, particularly for complex smart devices or intelligent environments. Selecting a particular mechanism requires trade-off of resources or functionality to a certain extent [2]. Some bespoke control system middleware designs help recover functionality loss by reproducing components that help to extend the chosen mechanism. Alternatively, programmers could implement an architecture where multiple independent middleware layers run simultaneously within a single control system. Both of these solutions naturally require a greater coding effort to implement the system, which itself is a lengthy program processing than the default implementation.

A potential risk of data synchronization also prevails within the system [3]. For instance, the device state variability; updated using one middleware system, cannot be reflected by all the alternative mechanisms used. If their state data became desynchronized, smart devices controlled by agents may exhibit erroneous or undesirable behaviours, which could prove costly and dangerous. These issues also emerge due to hardware action request response times. If the new state data needs to be shared amongst several components within the middleware (e.g., to

synchronise the system), this additional processing could be slower to changing events within the environment. This could frustrate users when updating their device and makes operating agents more apparent within an intelligent environment. For example, walking into a dark room and needing to wait for the lights to switch on.

This paper highlights work performed by an ongoing international research project (ScaleUp project) investigating methods for scaling up intelligent environments. As a consequence of this research, a new middleware was developed, specifically designed to address the trade-off issue, when deciding between functionality provided by UPnP and Web Services implementations. The aim was to amalgamate distinctive beneficial features offered by the UPnP and Web Services approaches into a new middleware design. The new mechanism was also structured to protect intelligent environment control systems against data de-synchronization, and processing action requests in times at least on par with other existing middleware examples. This work was based on ongoing international collaboration between the University of Essex and King Abdulaziz University.

The core contribution of this study is an evaluation of both the state-of-the-art and state-of-the-practice in the middleware research offered by existing middleware products. Software engineers progressively utilize middleware for building distributed systems. Any research into ScaleUp project that avoids this trend will merely have confined effect. Therefore, this paper has analysed the effects of middleware on the software engineering research agenda. It is argued that requirements engineering techniques are required for focusing on non-functional requirements, as these impact the selection and use of middleware. Software architecture research can develop methods guiding students towards the selection of right middleware and integrating it so that it fulfils a series of non-functional requirements.

2. Literature Review

2.1 Related Work

Currently, many off-the-shelf commercial devices possess inbuilt intelligent functionality, (i.e., either contain or can be accessed in some way via a computer system) [4]. To facilitate usability of their products, hardware manufacturers include pre-installed middleware designed for computer control. Previously, control systems were device-centric; for example, the control system for washing machine cycle programming. However, with the widespread use of networking platforms such as Ethernet, Wi-Fi, and Bluetooth, the interconnectivity between different devices began to emerge. Being proponents of Pervasive Computer Science concepts, such as the Internet-of-Things [5], several hardware manufacturers chose to implement middleware for their products around established common frameworks, such as UPnP.

This was due to a belief that by adopting common middleware architecture, customers would use smart devices created by rival companies, within their proprietary personal home networks. Though, this concept is yet to be turned into a mainstream commercial reality. Various larger multinational hardware manufacturers, such as Apple or Sony are adopting more protectionist marketing strategies, creating interconnectivity between their products and services. For example, Apple iWatch smart device release will only work when paired with an Apple iPhone 5 or newer, though it is not compatible with any other older iPhone models or any rival smart-phone handset [6]. Such smart devices operate using completely bespoke proprietary middleware systems, designed solely by a particular company.

The current literature shows that the use of middleware in IoT is limited [7]. Ngu et al [8] have designed an IoT application for real-time prediction of blood alcohol content based on smartwatch sensor data. The study has conducted a survey on the competencies of the current

IoT middleware. In addition, the challenges and the enablers associated to the IoT middleware were presented in order to embrace the heterogeneity of IoT devices. Fremantle and Scott [9] have utilized a structured search approach for identifying 54 particular IoT middleware frameworks and examined the security frameworks associated to each middleware. A total of 12 requirements (integrity and confidentiality, access control, consent, policy-based security, authentication, federated identity, and device identity) were used from the first stage for validating the competencies of each system.

Elkhoder et al [10] have accounted middleware for the emerging attributes such as seamless communications, lightweight aspects, and mobile across different heterogeneous networks and domains. It involves a context-adaptive technique, which allows the user for managing the location information shared by things on the basis of policy enforcement and context-aware mechanism. This mechanism accounted both the preferences and informed consent of a user. Jyothi [11] has managed data volumes and supported semantic modeling in the open issues, specifically managing the crowd sourcing of different domains. There is a scope for research work in order to make a generic IoT-middleware system, which is relevant across all regions by making all the functional aspects reusable and can be included as enabler to the middleware system.

Razzaque, Milojevic-Jevric, and Palade [12] survey showed that the use of middleware assists in the development through its integration of heterogeneous communication and computing devices. It also states that middleware provides interoperability support for application across different devices and services. Jeon and Jung [13] have revealed that the average request rate elevated by 25 percent compared to Californium, which is a middleware for effective association in IoT environments with vigorous performance, a power consumption

reduced by up to 68% and an average response time reduced by 90% when resource management was utilized. Lastly, the latency and power consumption of IoT devices can be reduced by the proposed platform. According to da Cruz et al [14], an important role is played by middleware as it is accountable for covering the intelligence part in IoT, which make decisions, allow them to communicate, and integrate data from devices on the basis of data collected. Afterward, a reference architecture model was investigated for IoT middleware based on IoT platform requirements, which detail the effective operational approaches of each proposed module, and proposed fundamental security features for this software type. Zhao et al [15] have proposed a stack of support-communication-computing for integrating effective open-source projects in order to devise techniques for allowing sufficient uniform human-thing associations, and developing implementation foundations for cutting-edge technologies including semantic reasoning and fog computing.

Due to the lack of common frameworks adoption for building complex intelligent environments or devices, the addition of a bespoke middleware layer is critical. This layer is designed to collectively handle the respective control systems of each installed smart device type. This extra layer of middleware wraps the individual bespoke device controller programs into a single API, representing the entire intelligent environment. This is also created by different Computer Science projects, which create middleware for intelligent environments, linking their creations into bespoke device controllers. Although all might concentrate on the same area of the control system architecture, the actual functionality can have a broad design. For example, Román et al. [16] created their bespoke Gaia infrastructure to allow distributed collections of smart objects and environments, represented and accessed via an interface. Later, Roalter et al. [17] chose to integrate a control system for a robot into their intelligent environment middleware,

allowing it to seamlessly access sensors and actuators within the smart space. In their OpenCOPI middleware, Lopes et al. [18] utilized a web ontology language based around semantic web services to create a mediator regulating access between ubiquitous applications and available service providers. Stavropoulos et al. [19] also used web Services as the base for their aWESoME infrastructure, which in addition to being an intelligent environment controller, focused on promoting “energy savings” by consuming low power in its operations. Finally, Olaru et al. [20] used a context-aware multi-agent system as their middleware base for controlling ambient intelligence exhibited within an environment.

Different middleware platforms were discussed on these criteria for networked robotic systems [21]. Majority of the middleware platforms have varied objectives including reusability, development process, self-discovery, self-configuration, supporting QoS, flexibility, and integration. In addition, several middleware platforms were discussed in the study for networked robotic systems such as self-adaptation, discovery and higher-level abstractions, collaborations support, and other advanced characteristics for integration.

Rodriguez-Molina and Kammen [22] have apparently demonstrated the collection of services demonstrated by community of researchers, developers, and scientists. Furthermore, middleware solution utilizes an API that explains how services were accessed from both the applications and the hardware that has been embedded for a Smart Grid-like deployment. Moreover, the authors have indicated that boundaries exist between the hardware located and the network, which are compliant with a standard demonstrating the sub-systems of the software. Vikash et al [23] have conducted a survey on middle-wares of WSNs towards IoT in offering a comparative view of different middle-wares and how this middleware technology can be utilized for implementing several issues emerging in the development of IoT applications. Rodriguez-Molina et al [24]

have utilized different aspects with several functionalities preferred as essential for a semantic middleware architecture conjoined with maritime operations including access to the application layer, context awareness, and device and service registration. On the contrary, Rodriguez-Molina et al [24] have interweaved other technologies with middleware such as acoustic networks and wireless communications. Under such circumstances, Rodriguez-Molina et al [25] have established an approach for interchanging information at the data level among independent maritime vehicles, which is of significant importance as the required information will have to be defined, along with the size of transferred data. Rodriguez-Molina et al [25] have forwarded the Maritime Data Transfer Protocol for interchanging standardized aspects of information at the data level for maritime independent maritime vehicles, and the procedures that are needed for information interchange.

2.2 Techniques/Protocols/Paradigms used in the Development of Proposed Work

Figure 1 highlights the integration of the additional middleware layer within the control system architecture using the ‘Device Wrapper’ node. Without this layer of middleware, it would be necessary first to learn the use of interface and to control each smart device individually. This requires the use of multiple different programming languages, such as Java, Python, C, C++ or C#, and possibly OS-dependent software packages. For example, the University of Essex currently has two full-scale intelligent environment test-beds, each using a different style of additional middleware layer to amalgamate devices from various manufacturers. These include iClassroom and iSpace.

[Insert Figure 1 Here]

The iClassroom [26] is an intelligent environment customised to resemble a university or

school teaching room (Figure 2). Most smart technologies are used directly to augment presentations or other teaching strategies to enhance student learning experiences. The middleware in the environment uses Web Services to wrap diverse collection of devices into a single common interfacing mechanism. The Web Services system uses Eclipse's Jetty Web Server [27] for its operations.

[Insert Figure 2 Here]

The iClassroom uses a centralized configuration, with all the middleware running from a single server. In terms of functionality, the main user interface is in bespoke website form hosted on the environment server. This server contains numerous hyperlinks to the Web Services used to control each device. The user interface can be loaded onto any standard browser application such as smart phones and tablets. Through the user interface, smart device in the environment can be connected, controlled, and monitored to the Web Services system from any remote location capable of accessing the iClassroom network.

[Insert Figure 3 Here]

The iSpace [28, 29] is an intelligent environment test-bed customized to resemble a typical household environment. The multi-roomed space includes a full-sized lounge/kitchen, study, bathroom, and bedroom, each connected by a central hallway. Unlike the iClassroom, most of the smart technologies in the iSpace are deliberately concealed within hollow walls and ceilings (Figure 3). This provides unknowing visitors to space an initial impression that they are in a normal (i.e., non-augmented) environment. The middleware in this intelligent environment is based upon UPnP, with controllers for over sixty smart devices linked together into a single API. More specifically, Youpi UPnP stack [30] was used to implement wrappers for each smart device.

The iSpace comprises a distributed configuration, with the middleware and control code for devices which split across several different computers, each connected via a common network. The Youpi UPnP wrappers allow each smart device to broadcast its existence on the same network [31]. iSpace allows programmers to create a bespoke controller code to integrate the smart devices into their projects better. To discover them, programmers must create a Youpi UPnP control point within their code and perform a search. This can be either for a specific device or a general search that returns a set of every smart device discovered on the network. Using a returned smart device instance, a control program can then isolate specific state variables, actions and arguments to either monitor or modify the associated hardware. Using UPnP, user control programs can ‘subscribe’ to individual state variables within a smart device instance. The state of a subscribed variable changes can automatically flag an associated listener within the user’s interface program.

For a large environment like the iSpace, implementation and maintenance of UPnP middleware require significant programming effort. To operate a wrapper for each smart device instance, a significant amount of information concerning hardware and its uses are required. Initially, each device type needed every state variable, action, and argument that can be individually declared. These were associated with dedicated UPnP state listeners to allow the correct functioning of the subscription system. Finally, it was necessary to individually initialize and start each instance of a smart device in the iSpace, which included assigning unique attributes such as names and UUIDs.

3. The Environment Development Kit (EDK)

The new system is created from scratch and is not an extension or reconfiguration of a pre-existing architecture. Thereby, steps should be taken to address the synchronisation and resource

allocation issues. This led to the creation of the Environment Development Kit (EDK), which was written using only standard Java SDK, with no extensions or third-party APIs. The EDK architecture design declared smart device present within an Intelligent Environment to be;

- a) Discoverable on a network.
- b) Subscribed to using listeners, which monitored for state change events.
- c) Accessible and/or controllable via a common API.
- d) Accessible and/or controllable via a set of auto-generated Web Services.
- e) Assigned with a unique controller program instance or one collectively shared by a group.

The system operates in both a centralized or distributed context and is not OS dependent. Controllers assigned to individual or groups of devices could be updated using a ‘hot-swap’, without needing to restart any part of the system, (as long as the device is not being accessed at the time of modification). Additionally, the inventory of devices is not to be declared before the system could run, as new devices could be discovered and handled by the system at any time. The case is similar to various middleware previously designed for intelligent environments, such as UPnP.

A multicast communications system allows devices to be linked to agents and other user programs. Multicast is less dependable than alternative communication methods such as TCP. It is because there is no guarantee that intended targets would receive the message broadcasted to a network. However, it is significantly much faster than TCP and can share message across multiple targets in a subscribed group simultaneously, which determines the main reasons behind its usage in this middleware system.

Similar multicast systems are in use in computer game virtual worlds, where the environment is divided into regions, each with their assigned group containing all the players, objects and

other updatable information present in that area [32, 33]. When an entity moves from one region to another, it should also change its subscription to the corresponding multicast group while simultaneously leaving the previous one.

[Insert Figure 4 Here]

Users were able to access the EDK network, search for represented smart devices and send action requests to specific discovered instances via an inbuilt control point. Devices discovered by each control point instance were automatically ‘cloned’ and locally stored. Clone devices were exact copies of their originals, with the exception that they possessed no controller, so they could not directly interface with real hardware in the environment. Cloned devices received state updates from the original ‘master’ device. These synchronisation updates occur regularly. Typically, this requires an update several times a minute, although the interval period between messages could be increased or decreased according based on its importance for clones to return the most up-to-date state information possible.

For actuators, EDK control points bypass the clone representation and transmit a state change action request directly to the master device. Multicast communication used received response from the device indicating concerning the change of request wither its success or failure. However, if the new state was applied to the master device, then alteration of the local clone would occur during the next synchronisation update. To speed this process, the EDK was designed where changes are automatically and immediately prompted to a synchronisation update, temporarily overriding the scheduled system, (which had its internal timer reset). Figure 4 shows complete EDK system architecture for controlling a single, smart device via a third-party interface program.

Allowing users to create control points, the API provides a series of expandable classes and

methods that are used with a broad range of bespoke sophistication for devices or environments. The EDK came packaged with classes for representing several basic actuator and sensor types, used by master device representations, and subsequently by control point instances when creating their clones. Furthermore, the system was equally capable of handling encountering unknown device types in an ad-hoc manner. This was achieved by creating a cloned instance of the discovered device using a generic class appropriate for representing its declared features and variables. Several different generic classes existed depending on the device, whether the device was an actuator or a sensor and the number variables it possessed. Because of this feature, control points did not require setup or inventory of an intelligent environment before being connected. Additionally, it was possible to use the EDK with intelligent environments based around both centralized and distributed architectures, as there was no requirement to host master devices and clones on the same computer to communicate.

3.1 Implementing Master Devices

Three key components are required to exist on an EDK network;

- 1) A representation identifies the attributes of the device and declaring whether it is an actuator or sensor.
- 2) A gateway to the target EDK network from where the device will be accessible.
- 3) A controller connecting the virtual EDK representation with some counterpart hardware that exists in the real-world.

Together these components allow a “master” device to be created using the EDK API. Once deployed on the network, the virtual device representation would monitor or change the state of the real hardware-based upon received user instructions. It would also periodically broadcast details of itself and the current state of the hardware to the EDK network, which is received by

any active control points and used to create or synchronize associated “clone” devices.

Master devices can be implemented and deployed remotely on one or several networked computers. Depending upon personal preference, a master device application could also be used to generate one or several different actuators and sensors, with no requirement of device type. Multiple different control programs can be imported into the same implementation, and if desired, a group of devices can either share or each is allocated their instance of the same controller.

3.1.1 Step One: Declaring Constants

[Insert Figure 5 Here]

Firstly, several constants need to be declared to allow communication and different control systems to operate correctly. Figure (5) provide examples of these, where each attribute is explained further in the implementation process.

3.1.2 Step Two: Creating Device Instances

[Insert Figure 6 Here]

[Insert Figure 7 Here]

The EDK API contains several different classes that can be used to create new representations of individual master devices. The selection of the class for a given situation depends upon several factors, specifically, the number of state variables supported by the device along with its identification as an actuator or a sensor. For common device types, the EDK API contains several dedicated classes, many of which contain bespoke convenience methods that provide better control of specific state variables. For example, Figures 6 and 7 show the creation

of an instance of two different types of light emitting devices. Both types have a state variable “Power” declared, which controls whether the device is on or off. To support these actions, both the “BooleanLight” and “DimmableLight” classes contain two convenience methods “turnOn” and “turnoff” which allow the “Power” state variable to be set without the need to declare or process any new values directly. The “DimmableLight” class used in Figure 7 additionally contains “getBrightness” and “setBrightness” methods, as compared to the declared “Brightness” state variable, used to control the emitted light level.

Without these convenience methods, the state variable name would need to be declared along with any new state value (if applicable), in a formatted String in order to perform the same function. If the device being used is uncommon or unknown for some reason, the EDK API generic series classes can be created. This creates representations based purely on the number of declared state variables. For actuators, these would be “SingleVariableActuator” and “MultiVariableActuator”, whereas for sensors the appropriate generic classes would be “SingleVariableSensor” and “MultiVariableSensor”. Each class contains several different constructors based on the available information of the created device. For instance, in Figures 6 and 7, the constructor is provided with a name for the device instance, the supported state variables (provided in a string array for multi-variable devices, as seen in Figure 7), and a description of the device itself. In this instance, all other required variables, such as a unique UUID, are allocated to the new device by the EDK API. Figures 6 and 7 also show how a “StateChangeListener” can be attached to individual master device instances, which flag whenever the value of any supported state variable is changed either by some attached hardware or due to a user request.¹

¹ Note: When implementing master devices, it is highly recommended that each instance is given unique name value. It should always be ensured that device instances always have different uuid values, which can be generated randomly using the Java SDK UUID class.

Note: All supported state variable names must be declared when creating a new device instance. Several pre-formed device types are included within the EDK API,

3.1.3 Step Three: Adding Communications

[Insert Figure 8 Here]

An instance of the “Communications” class must be created in order for master devices to be able to connect to an EDK network. Figure 8 shows how to do this using two of the variables from Figure 5 mentioned earlier to supply values for the network address and communications port variables. The EDK uses multicast communication to allow master devices to send updates to any running control points that have joined the same group. As a consequence, it is important to ensure that the value used for the “NETWORK_ADDRESS” variable is a valid multicast address. It may also be necessary to open the value used for “COMMUNICATIONS_PORT” on firewalls, which may be blocking the sending or receiving of multicast communications packets on the network.

3.1.4 Step Four: Loading Device Controllers

[Insert Figure 9 Here]

Control system programs for individual smart devices are created independently of the EDK. Once implemented, a device control system can be uploaded into a master device via the EDK API. Figure 9 shows how to do this. “DeviceController.jar” is the filename of the controller being uploaded, from the designated “CONTROLLER_DIRECTORY”. To integrate the controller program with the master device, the EDK needs to know the package (“CONTROLLER_PACKAGE”) and the name of the main class (“CONTROLLER_CLASS”). As before, examples of these values are provided in Figure 5.

which often contain several convenience methods for performing device-specific actions. To use these methods, state variables must be declared using the specific names specified by the Javadoc information for the relevant class.

Note: If an application is intended to deploy multiple master device instances, it is recommended to create each instance at this stage before continuing.
Note: To prevent processing errors, individual state variables should never include the colon or tilde characters (i.e., : or ~) in their names or possible returnable values.

378

379 3.1.5 Step Five: Creating a Device Processing Hub

380 [Insert Figure 10 Here]

381 The next step is to create a hub to process devices. The “communications” variable of the
 382 “DeviceHub” constructor should be the same instance of the “Communications” class created
 383 back in Step Three. Once the hub is initialised, each of the smart device instances created in Step
 384 Two needs to be individually added. As these are master devices, they also need to be associated
 385 with their respective controller programs, loaded during Step Four. Figure 10 provides an
 386 example of how to implement this for light devices.

387 The purpose of the “DeviceHub” class differs slightly depending upon whether it is used
 388 within an implementation for master devices or an EDK control point. For master devices, the
 389 hub liaises with the communication system created in Step Three, to have each of its stored
 390 devices access their associated controllers and perform a live update of their recorded state
 391 variables. Typically, this would involve each device accessing the real-world hardware to which
 392 its control system connects it. Once acquired, the up-to-date state variable information is then
 393 passed back to the communications system, where it is transmitted to the EDK network and used
 394 synchronizing any listening clone device representations.

395

396 3.1.6 Step Six: Adding Web Services, (Optional)

397 [Insert Figure 11 Here]

398 Using the Web Services interface method with an EDK middleware implementation is
 399 optional, so this step can be ignored if desired. However, enabling the Web Services system only
 400 requires the code shown in Figure 11 to be added after the creation of the device hub. The
 401 variable “WEBSERVICES_PORT” is the last of the declared constants create back in Step One,

402 which in the case of this example (Figure 5).

403 Once added to an instance of the “DeviceHub” class, the EDK mechanism will auto-generate
 404 a Web Services control interface for each of the declared smart devices and add them to the
 405 internal HTTP Server. Currently, two different acceptable commands are implemented for
 406 sensors, (i.e., about and get), while three for actuators (i.e., about, get and set). The Web Services
 407 interface can be loaded using any standard Web Browser, including on most mobile devices,
 408 such as smartphones and tablet computers. The syntax for an EDK Web Service is naturally
 409 bespoke to each situation where it is used, but the basic URL structure is as follows;

**http://<Computer IP Address>:<Web Server
Port>/<Device Name>/<Command>**

411 So, based upon the “BooleanLight” and “DimmableLight” smart device examples used
 412 throughout this tutorial, some acceptable URLs would be;

413
 414 http://127.0.0.1:8000/Light1/about http://127.0.0.1:8000/Light1/get
 415 http://127.0.0.1:8000/Light1/set?Power:1 http://127.0.0.1:8000/Dimmer1/get?Brightness
 416 http://127.0.0.1:8000/Dimmer1/set?Brightness:75

417

‘127.0.0.1’	The IP address of the computer running the master device representation being accessed, (likely to be different from localhost).
‘8000’	The port number for the Web Server, as declared by the value of the <code>WEBSERVICES_PORT</code> constant created back in Step One.
‘Light1’ ‘Dimmer1’	The names of the smart devices being accessed, as declared when their representations were created in Step Two.
‘about’	A command to display general information about the specified device.

‘get’	A command to return the name of each state variable supported by the specified device, along with its currently recorded value.
‘set?Power:1’	A command to set state variable „Power“ to a value of „1“.
‘get?Brightness’	A command to return the current value of state variable “Brightness”.
‘set?Brightness:75’	A command to set state variable “Brightness” to a value of “75”.

418

419 3.1.7 Creating a Control Point

420 [Insert Figure 12 Here]

421 An instance of the “ControlPoint” class allows external client programs to access devices via
 422 an EDK network. Figure 12 shows the API code required to perform this task. The String
 423 variable used in the “ControlPoint” constructor is a bespoke name for the specific instance being
 424 created. If more than one control point is being used within a client program, (i.e., to access
 425 different EDK networks), this variable can be used to identify specific instances.

426 A control point effectively acts as a portal into the EDK middleware system, allowing users
 427 to search for groups or specific master devices on the associated network. Its associated
 428 communications system provides details of the network which is accessed by a control point. In
 429 addition to providing details of an EDK network, the communication system supplied to a
 430 control point is also responsible for processing state update messages for the master devices. The
 431 control point itself automatically generates an internal “DeviceHub” instance, which is used to
 432 create and store clone device instances based on information received by the communications
 433 system from the EDK network. The control point accesses the stored clone devices and uses their
 434 information to provide returnable results for user searches.

435 3.1.8 Searching for Devices

436 [Insert Figure 13 Here]

To search for known master devices, present on an EDK network, a “ControlPoint” instance can use its “searchForDevices” method (Figure 13). Instances of “ControlPoint” will only become aware of master devices upon receiving an update packet from them. Therefore, upon initially starting, the delay might occur before all master devices present on a network are discovered, concerning their update cycles when the control point joins the EDK multicast group. It is typically a good idea to enclose the search command in a “for”, “do-while” or “while” loop to keep the control point scanning the network until a non-empty array is returned or the desired device is found. Alternatively, this command could also be repeatedly called from within an isolated thread, which runs continuously in the background, allowing devices that start broadcasting after the control point’s initial search also to be detected.

3.1.9 Searching for Specific Devices

[Insert Figure 14 Here]

If details of the target device or devices are known ahead of time, a client program can alternatively use one of the more specific search methods of the “ControlPoint” class. Figure 14 shows three such methods, each targeting different attributes of master devices. Firstly, the “SearchForDevicesByType” method can be used to filter the known list of master devices by their type, returning an array of any matching instances stored in the control point “DeviceHub”. The topmost example in Figure 14 uses this method to search for all instances of a “BooleanLight”. The “Actuator” and “Sensor” classes in the API both contain numerous other declared variables that can be used with this method each representing one of the pre-formed smart device types included in the EDK. Alternatively, to search for a device type not included within the standard API, programs can use the “SINGLE_VARIABLE_DEVICE” and

“MULTI_VARIABLE_DEVICE” variables, (also found in the “Actuator” and „Sensor“ classes), or a bespoke device name entered as a String.

The two remaining search methods, shown in Figure 14, are each designed only to return a single device, matching either a specified name or uuid criterion. If no matching device is found, then a “null” value is returned. If, for some reason, two different master devices existed on an EDK network and both were called “Light1”, the “searchForDeviceByName” method will only return the first instance it encounters upon contents scanning of control point device hub. This also applies to the “searchForDeviceByUUID” method if both devices share the same aid value.¹

3.1.10 Processing EDK Smart Devices

The instance of “Device Hub” created by a control point is used to store clones created to represent networked master devices locally. The device hub creates the clones, which is the automatic response when information is received from the communication system that does not match any previously known representation. Clone devices are generally stored locally on the same computer as the control point used to create them. Aside from not possessing controllers for hardware, they are identical to the master devices that spawned them in every way. Even the unique attributes of the original master device are copied, including its name and quid value. As such, by using the associated classes within the EDK API, the details and states of each clone can be read, and in the case of actuators manipulated, in the same manner as if connecting directly to any master device instance.

[Insert Figure 15 Here]

The EDK API contains several different methods of reading the state of smart devices. A

¹ Note: When implementing master devices, it is highly recommended that each instance is given unique name value. It should always be ensured that device instances always have different uuid values, which can be generated randomly using the Java SDK UUID class.

selection of the possible options is listed in Figure 15. The determination of best method largely depends upon what information is required about the state of a device and how it is subsequently used.

From Figure 15 examples, the topmost method is a general “getState” command, which is common to every EDK device. When called, the “getState” command will return a single string representation of the entire device, or more specifically, its state variable values. Responses are always sent in pairs, with the name of the state variable and its current value. For instance, in the case of the “Light1” “BooleanLight” device used in the implementation examples earlier, an “etState” request could result in either of the following String responses;

Power:0
Power:1

Where, “Power” is the name of the only state variable included in a “BooleanLight” object, while zero (off) and one (on) are the current values of that state. The colon separating the two values acts as a key in a split command to allow easy separation of variable name and its value. In the case of multi-variable devices, which contain more than one state variable, an additional tilde key (“~”) is added to separate the individual attributes. So, for the “Dimmer1” DimmableLight” used in earlier examples, a “getState” request could return;

Power:0~Brightness:100

Where, “Power” is the first state variable and “Brightness” is the second, with current recorded values of zero (off) and one hundred (percentage of maximum illumination), respectively¹.

¹ Note: To prevent processing errors individual state variables should never include the colon or tilde characters (i.e.: or ~) in their names or possible returnable values.

Figure 15 shows three methods that API can be used to return only the current value of a specific state variable, rather than the name/value pairs shown above. Generally, this is achieved by specifying the name of the state variable whose value is required as a variable in the method. The second and third down examples in Figure 12 demonstrate this process for a “BooleanLight” and “DimmableLight” respectively. Typically, the returned state values are in a String format, but can easily be converted as shown with the “Brightness” variable example, where the value is converted into an integer once returned.

In many cases, the string can be avoided to integer conversion, as performed in the third example of Figure 15. This is based on EDK API device’s inclusion of convenience methods, which return state variable values in their most appropriate format automatically. This is demonstrated by the bottommost method in Figure 15, which uses a “getBrightness” method found in the “DimmableLight” class, which automatically returns the state value as an integer. All that is required to use these bespoke convenience methods is to cast the generic “Device” object returned by a control point search into the appropriate device type class, as is shown in the example.

3.1.11 Writing to Smart Devices

[Insert Figure 16 Here]

In the example provided by Figure 16, an array returned by a control point search (as described in Section 2.2), is scanned for a specific device called “Dimmer1”, which is also a “DimmableLight”. If found, the value of the state variable “Power” is requested from the device. If the value returned for “Power” is zero (off) then a command is sent to turn the light on. For

any other circumstance, the command is to turn the light off.¹

In the example, the generic “device” variable taken from the “Device” array is cast into a “DimmableLight” object. This allows the programmer to access the two convenience methods “turnoff” and “turnOn,” which are contained within the Actuator class. The EDK API contains several models for intelligent devices that can be used in place of the more generic actuator and sensor classes. Some of these classes also contain further convenience methods specific to that device type. For example, the “DimmableLight” class also contains a “setBrightness” method to allow a specific value to be entered for a “Brightness” state variable, controlling the amount of light emitted by the device.

534

3.1.12 Implementing EDK Compatible Device Controllers

[Insert Figure 17 Here]

Recalling back to Section 3.1.1, two of the constants that needed to be declared when implementing a master device (as shown in Figure 5) were “CONTROLLER_CLASS” and “CONTROLLER_PACKAGE”. The origins of the values used in the Figure 5 example, can be seen in Figure 17. The value that should be used for the “CONTROLLER_CLASS” constant is the name of the main class of the control system program, which in the example is simply “Controller”. Additionally, the value for the “CONTROLLER_PACKAGE” is the name of the package containing the declared main class, in this case “devicecontroller”.

When implementing a control program for smart devices to be used with an EDK middleware system, Figure 17 shows the minimum classes required for integration. More specifically, the “getState” and “setState” methods are both essential and should be used to

¹ Note: In the example provided in Figure 16, “Actuator.DIMMABLE_LIGHT” could also have been used instead of the String value “DimmableLight.”

directly return or update the current state of the associated hardware, respectively. If additional code is required to create a link with the associated hardware, such as using a third-party software package (e.g., RXTX or a manufacturer API), then all this code should be placed into a constructor within the main class, as shown by the “Controller” constructor in Figure 17.¹ If necessary, the constructor code should establish a connection with the smart device hardware and then maintain it as a global variable that can be accessed directly by the “getState” and “setState” methods. Alternatively, the constructor could be used to start an isolated thread, which uses locally declared variables, also accessible by the “getState” and “setState” methods, to handle state action requests. It is essential that no code necessary to directly establish a connection with hardware is included in either the “getState” or “setState” method as doing so could lead to overall instability in the EDK middleware system.

3.1.13 Controllers for Multi-Variable Devices

[Insert Figure 18 Here]

Figure 18 highlights how the controller code in Figure 17 can be extended to handle smart devices with multiple state variables. The code presented (in Figure 17), could function with a multi-variable smart device, but it is often desirable to separate certain state variables to better structure program code, or for efficiency, etc. Figure 18 shows a template for the controller used with the “DimmableLight” device “Dimmer1” example mentioned throughout this guide. Notice how the “getState” and “setState” methods have been retained (for handling the “Power” state variable), although the “Brightness” state variable has been separated and given its handling methods, namely “getBrightness” and “setBrightness.” To add additional “get” and “set”

¹ Note: Methods used in device controllers must return state values in as a single String using the expected EDK formats, as discussed in Section 3.3.1.

Note: It is recommended that the “setState” method returns the same value as an action request to “getState” once completing its operation.

methods, it is necessary to ensure that their suffix is named the same as the state variable they are expected to handle, (e.g. “getPower”, “getBrightness”, “setPower”, “setBrightness”, etc.).¹

When accessing a loaded device controller, an EDK implementation will first search through the methods of the declared controller class (i.e. “CONTROLLER_CLASS”) to see if it contains a bespoke match for the current state variable it needs to process. If no appropriate method matching the state variable can be found, the system will then automatically default to either the “getState” or “setState” method, depending upon which action is being performed.

Additionally, in Figure 15, the package name has changed, reflecting that it is a different device controller program from Figure 14 example. To be loaded correctly into the EDK and used with a master device, Figure 15 example would need to specify “devicecontrollermultiple” as the value for “CONTROLLER_PACKAGE” during the first implementation step.

4. Evaluation

4.1 Evaluation Strategy

One of the criteria used to measure the success of the EDK is its effective processing of action requests for smart devices, compared to existing middleware solutions. Since UPnP and Web Services largely inspired the new mechanism, examples of these architectures are used as a benchmark. Additionally, EDK testing with sensors and actuators is necessary, to acquire accurate results for both sending and receiving state variable data. Therefore, the middleware was integrated into the control systems of two different intelligent environments, located at the University of Essex, namely the iClassroom and the iSpace test-beds.

The intelligent environments chosen for the evaluation were selected as benchmarks to take

¹ Note: Any “set” methods should also only expect to receive a single variable containing the new state value to be processed, as shown in Figures 17 and 18. No other variables must be added for the methods to be compatible.

Note: As with “setState” it is recommended that any “set” method returns the same value as an action request to “getState” once completing its operation.

advantage of the established middleware in their existing control systems. During the evaluation, EDK and benchmark middleware systems were tested by controlling a DMX-based dimmable spotlight and a Phidgets light sensor. Primarily, it assesses the average processing speed required by the different components of the EDK, (i.e., API control and Web Services) compared to that needed by the equivalent benchmark systems. If the EDK was capable of performing the same tasks as an existing middleware system in approximately the same time frame, it is considered a viable alternative for use in an intelligent environment control system. Incidentally, it was anticipated that the EDK would actually require a significantly lower processing time than UPnP or Web Services, mostly due to the greater emphasis on locally processed variables.

4.2 Evaluation Results

To evaluate the EDK middleware system, each mechanism was tested separately, with no other programs running at the time. For each middleware system, five hundred ‘get’ or ‘set’ requests were made to the same sensor or actuator, respectively. This workload was split across five sessions of one hundred requests each. To keep the test fair, each session was performed using the same computer, which ran both timer program measuring the processing intervals, plus made the state requests to the relevant middleware implementation.

4.2.1 System Validation: Sensors

The first experiment in the project evaluation focused specifically on comparing middleware used to access sensors. In other words, devices whose functionality consisted only of sending information about their current state back to a requesting program. The Jetty-based Web Services system in the iClassroom linked directly into individual device control programs, which

automatically requested a current reading from the relevant sensor hardware when called. However, the iSpace UPnP mechanism functioned in a very similar style to the EDK middleware design. Specifically, whenever called, the ‘getState’ command would access the control system for a specific device and return the last recorded value obtained from the actual sensor itself. A ‘getState’ method call never actually resulted in the hardware being accessed directly in order to take a reading. Updates to the recorded light level value were handled by an independent thread located within the program code for the sensor, which automatically took a new reading from the hardware intermittently. In the case of the iSpace UPnP system, a new reading was recorded approximately once every ten seconds. For EDK, this interval was reduced to approximately three seconds, in order to reduce the margin of error between the recorded and actual sensor values.

[Insert Table 1 Here]

The experiment involved measuring the time required for each middleware system to return a lux value for a specific light sensor Phidget. As can be seen from the results displayed in Table 1, on average for each session run, both device interface methods offered by the EDK middleware required significantly less time to process individual action requests than both the Jetty-based Web Services or Youpi UPnP implementations. In the case of the Direct API interface method, the processing time was reduced to nanoseconds as the sensor value being returned in response to each action request was taken from a locally stored variable within the clone of the real sensor generated by the EDK. As mentioned above, the Youpi UPnP middleware used a similar method, but the recorded values for the light sensor were stored on a remote server, hence required additional processing time for the system to access and get the data values from. Unlike the API interface method, action requests made using the EDK Web Services were directed at the real

light sensor, yet they still required a fraction of the processing time of the Jetty-based system, consistently throughout the evaluation.

4.2.2 System Validation: Actuators

The second experiment performed for the evaluation focused on actuators. More specifically, this involved devices that contained one or more variables that could have their state modified, via the middleware accessing an attached control system. When setting the state of a variable, the success of the corresponding hardware being updated was not always guaranteed. This is because many devices do not provide any feedback to the controller code, concerning whether a transmitted command has been received or acted. Both, Jetty-based Web Services and Youpi UPnP middleware sent acknowledgments back to the control programs where action requests originated but always assumed a positive outcome. In contrast, the EDK did not send any acknowledgments in direct response to individual state change requests. The success or failure of an update could be determined by observing the next set of variable settings transmitted by a real device to its clones.

[Insert Table 2 Here]

This experiment involved each middleware system accessing a single DMX-controlled dimmable spotlight and alternating its brightness level between fifty and one hundred percent. Table 2 lists the average processing times required by each batch of one hundred action requests made during the experiment. As with the sensor experimentation, both EDK interface mechanisms were found on average to require significantly less processing times than the Jetty-based Web Services and Youpi UPnP systems. It is worth noting that the difference between the EDK and Jetty-based Web Services was not as profound as for the sensors. This was due to the

controller code for the DMX lights adding a delay of approximately one hundred milliseconds to the processing of each action request via a compulsory sleep command. As the EDK Web Services were directed at the representation of the actual light rather than a local clone, the delay was unavoidable.

Unlike for the sensor ‘getState’ action requests, when setting a device’s state, the EDK’s Direct API interface method also targeted the original ‘master’ representation rather than a clone. Effectively, a state change action request was forwarded to the original device by sending out a state change request using the inbuilt multicast communication system. As the system did not require any acknowledgment, the process could be ended at that point, unlike UPnP and the Web Services mechanisms, which required some kind of response, even if based upon an assumption.

4.3 Evaluation Summary

Concerning the processing times for action requests, on average, in every session of the evaluation, both interface mechanisms of the EDK were found to be significantly superior to the benchmark Jetty-based Web Services and Youpi UPnP middleware systems. This result was not entirely unexpected as the EDK utilised locally stored variables much more than either of the benchmark systems. Furthermore, in cases where remote access is necessary, (i.e. when setting the state of an actuator), the EDK does not require any confirmation of a device receiving an action state request to be returned, further reducing the processing times.

For experiment two, the additional ‘power’ state variable of the DMX spotlight was manually set to an ‘on’ state prior to each experiment session to observe the brightness changes but was otherwise unused. However, it should be noted that for each action request sent to the UPnP middleware system, it was necessary for its control point to first perform a search of all the DMX

light state variables in order to isolate the one controlling the brightness before the new state could be set. As more than one state variable existed within the device (i.e., power and brightness), this may have caused the action request processing time to be extended, depending upon which was discovered first. This additional step was not necessary for the EDK or Jetty-based Web Services, where the individual state variables were already declared within the device controller code. This may explain why the UPnP mechanism required more time than any other mechanism in each session when handling an actuator, despite being faster than Jetty-based Web Services in three out of five sessions during the earlier sensor experiment.

[Insert Figure 19 Here]

[Insert Figure 20 Here]

Finally, another observation made during the evaluation was that for each batch of one hundred action requests performed, while the average processing times generated by Jetty-based Web Services and Youpi UPnP fluctuated noticeably between sessions, the times for both EDK methods remained largely static. This effect is illustrated by the graphical representations of the results from the evaluation presented in Figures 19 and 20. These figures also clearly highlight how the EDK was capable of outperforming the two benchmark middleware mechanisms consistently throughout the entire evaluation.

5. Conclusions

This paper discussed recent work performed into creating a new middleware architecture for an intelligent environment control system. It shared insights on how control system developers often require some functionality trade-off when integrating one of the numerous existing architectures. Alternatively, an existing mechanism may be manually augmented with additional

code to overcome any missing functionality, which could result in detrimental synchronization issues within the resulting control system. This project aimed to create a new middleware, capable combining the positive features of two well-known architectures (namely UPnP and Web Services) into a single mechanism, to overcome functionality and synchronisation issues. Most importantly, it envisaged that the new mechanism would be required to process state action requests as quickly as the two benchmark middleware architectures.

Complex, intelligent environment developers must implement their additional bespoke layer of middleware to link devices and services provided by different manufacturers. iClassroom and iSpace control system architectures in two full-scale intelligent environments are also presented. It also discusses the Environment Development Kit (EDK) as the new alternative middleware as an outcome of the research performed by this project. It assesses how viable the new middleware design was as an alternative to UPnP and Web Services. This strategy involved two separate experiments testing sensors and actuators independently due to their different properties. It showed that a comparative analysis of the average action request processing times required by the EDK, plus implementations of the benchmark UPnP and Web Services systems.

The evaluation experiments both indicated overwhelmingly that the two interface methods offered by the EDK architecture implementation could process state action requests to sensors and actuators faster than either benchmark mechanism. Out of the ten sessions performed, either benchmark did not outperform the EDK in a single case. EDK and UPnP used stored variables when returning sensor state data, rather than accessing the actual device directly. These stored variables were updated periodically by an independent thread, which noted the actual sensor readings. However, in the UPnP implementation, a ten-second delay between individual readings was observed along with an update of the stored variable. For the EDK implementation, this

delay was reduced to three seconds, yet the mechanism was still capable of processing action requests faster than the UPnP benchmark. The results strongly indicated that the EDK is a viable alternative to UPnP and Web Services in a smart device or intelligent environment control system.

Another important issue for EDK was to resolve the data de-synchronised using multiple device interfacing methods. As the new mechanism was designed from the ground up, steps could be taken to allow synchronisation between the API and Web Services interface methods to be practically guaranteed. For this, both interface methods use the same sets of variables to store descriptions of individual master or clone devices. Also, clone devices were regularly synchronised with their original master; locally stored representations were constantly being updated with the known true state values. This helped defense against de-synchronization of data stored in clone devices.

5.1 Future Work

With the initial framework for the EDK established, future researches can expand the existing middleware architecture to integrate additional functionality. Firstly, an internal mechanism can be introduced to add an internal authentication system, which allows instances of smart devices represented by the mechanism to optionally be assigned with a password. Unless the correct password was supplied during communication, neither control points nor other interfaces could access the device, nor a clone of the original be created as a result of a search request or update processes.

Therefore, future study should focus on EDK development, where its variables and methods allow control points and other interfaces to identify the specific environment where an instance

of a smart device is located. Technically this could already be achieved to some extent using the existing mechanism, for example, by manipulating the name variable when creating device instances. However, with the augmentations would allow the handling of scenarios such as the device changing locations to a different intelligent environment. It would also make it easier to represent multiple different environments on the same network. This could potentially allow the creation of agents or control programs that use devices from several different environments collaboratively, as part of their functionality.

Acknowledgement

The author is highly grateful for all the associated personnel who contributed in the completion of this Scale Up Project. The author declares that the study is not funded through any source, neither it holds conflicts of interest.

References

1. S. Tan, J. Zhang, F. Sun, and S. Wang, an approach to support the interoperability of intelligent grouping and resource sharing (IGRS) and universal plug and play (UPnP) in home network environment. In 2015 IEEE International Conference on Computational Intelligence & Communication Technology, (2015), 682-686. IEEE.
2. A. Rafique, D. Van Landuyt, B. Lagaisse and W. Joosen, On the performance impact of data access middleware for nosql data stores a study of the trade-off between performance and migration cost, *IEEE Transactions on Cloud Computing* **6**(3) (2015), 843-856. <https://doi.org/10.1109/tcc.2015.2511756>
3. D. Bruneo, A. Puliafito and M. Scarpa, *Mobile Middleware, The Handbook of Mobile Middleware*, (2006), 145–166. <https://doi.org/10.1201/9781420013153.ch7>
4. B. Yuan, Z. Xiaolei and W. Yitao, Identifying Vehicle’s Steer Change via Commercial Smartphones. In *2019 International Conference on High Performance Big Data and Intelligent Systems (HPBD&IS)*, (2019), 181-185. IEEE. doi:10.1109/hpbdis.2019.8735446.
5. J. Gubbi, R. Buyya, S. Marusic, and M. Palaniswami, Internet of Things (IoT): A vision, architectural elements, and future directions, *Future Generation Computer Systems*, **29**(7) (2013), 1645–1660. <https://doi.org/10.1016/j.future.2013.01.010>
6. Z.A. Almusaylim and Z. Noor, A review on smart home present state and challenges: linked to context-awareness internet of things (IoT), *Wireless Networks*. **25**(6) (2019), 3193-3204. <https://doi.org/10.1007/s11276-018-1712-5>
7. A. Farahzadi, S. Pooyan, R. Javad and F. Reza, Middleware technologies for cloud of things: a survey. *Digital Communications and Networks (DCN)* **4**(3) (2018), 176-188.

- 787 8. A. H. H. Ngu, M. Gutierrez, V. Metsis, S. Nepal, and M. Z. Sheng, "IoT Middleware: A
788 Survey on Issues and Enabling technologies," *IEEE Internet of Things Journal*, pp. 1–1,
789 2016. Doi: <https://doi.org/10.1109/jiot.2016.2615180>
- 790 9. P. Fremantle and P. Scott, "A survey of secure middleware for the Internet of Things," *PeerJ*
791 *Computer Science*, vol. 3, p. e114, May 2017. Doi: <https://doi.org/10.7717/peerj-cs.114>
- 792 10. M. Elkhodr, S. Shahrestani and H. Cheung, "A middleware for the internet of things," arXiv
793 preprint arXiv:1604.04823. 2016 Apr 17. Doi: 10.5121/ijcnc.2016.8214
- 794 11. T. Jyothi. Open Source Middleware for Internet of Things, *International Journal of*
795 *Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering*.
796 2016. Vol. 4, Issue. 11. 71-75. Doi: 10.17148/IJIREEICE.2016.41113
- 797 12. M. A. Razzaque, M. Milojevic-Jevric, A. Palade and S. Clarke, Middleware for internet of
798 things: a survey, *IEEE Internet of things journal*, **3**(1) (2015), 70-95.
799 <https://doi.org/10.1109/jiot.2015.2498900>
- 800 13. S. Jeon and I. Jung, "MinT: Middleware for Cooperative Interaction of Things," *Sensors*, vol.
801 17, no. 6, p. 1452, Jun. 2017. Doi: <https://doi.org/10.3390/s17061452>
- 802 14. M. A. A. da Cruz, J. J. P. C. Rodrigues, J. Al-Muhtadi, V. V. Korotaev, and V. H. C. de
803 Albuquerque, "A Reference Model for Internet of Things Middleware," *IEEE Internet of*
804 *Things Journal*, vol. 5, no. 2, pp. 871–883, Apr. 2018. Doi:
805 <https://doi.org/10.1109/jiot.2018.2796561>
- 806 15. R. Zhao, L. Wang, X. Zhang, Y. Zhang, L. Wang, and H. Peng, "A OneM2M-Compliant
807 Stacked Middleware Promoting IoT Research and Development," *IEEE Access*, vol. 6, pp.
808 63546–63559, 2018. Doi: <https://doi.org/10.1109/access.2018.2876197>

16. J. Rivera and R. van der Meulen, Gartner says the internet of things installed base will grow to 26 billion units by 2020. Stamford, conn., (2013).
17. C. Perera, P.P. Jayaraman, A. Zaslavsky, P. Christen and D. Georgakopoulos, Mosden: An internet of things middleware for resource constrained mobile devices. *In 2014 47th Hawaii International Conference on System Sciences*, IEEE. (2014). 1053-1062.
<https://doi.org/10.1109/hicss.2014.137>
18. M. Bazzani, D. Conzon, A. Scalera, M.A. Spirito and C.I. Trainito, Enabling the IoT paradigm in e-health solutions through the VIRTUS middleware. *In 2012 IEEE 11th international conference on trust, security and privacy in computing and communications*, (2012), 1954-1959). IEEE. <https://doi.org/10.1109/trustcom.2012>.
19. Echelon, Available at: “<http://www.echelon.com/products/i.lon-600-lonworks-ip-852-router>”, 2015.
20. Phidgets, Available at: “<http://www.phidgets.com>”, 2015.
21. Mohamed N, Al-Jaroodi J, Jawhar I. A review of middleware for networked robots. *International Journal of Computer Science and Network Security*. 2009 May 30;9(5):139-48.
22. J. Rodriguez-Molina and D. M. Kammen, “Middleware Architectures for the Smart Grid: A Survey on the State-of-the-Art, Taxonomy and Main Open Issues,” *IEEE Communications Surveys & Tutorials*, vol. 20, no. 4, pp. 2992–3033, 2018.
23. Vikash, L. Mishra, and S. Varma, “Middleware Technologies for Smart Wireless Sensor Networks towards Internet of Things: A Comparative Review,” *Wireless Personal Communications*, Aug. 2020.

24. J. Rodríguez-Molina, S. Bilbao, B. Martínez, M. Frasher, and B. Cürüklü, “An Optimized, Data Distribution Service-Based Solution for Reliable Data Exchange Among Autonomous Underwater Vehicles,” *Sensors*, vol. 17, no. 8, p. 1802, Aug. 2017.
25. J. Rodríguez-Molina, B. Martínez, S. Bilbao, and T. Martín-Wanton, “Maritime Data Transfer Protocol (MDTP): A Proposal for a Data Transmission Protocol in Resource-Constrained Underwater Environments Involving Cyber-Physical Systems,” *Sensors*, vol. 17, no. 6, p. 1330, Jun. 2017.
26. M. Román, C. Hess, R. Cerqueira, A. Ranganathan, R.H. Campbell and K. Nahrstedt, Gaia: a middleware platform for active spaces. *ACM SIGMOBILE Mobile Computing and Communications Review*, **6**(4) (2002), 65-67.
27. L. Roalter, M. Kranz, and A. Möller, A middleware for intelligent environments and the internet of things. In *International Conference on Ubiquitous Intelligence and Computing*, Springer, Berlin, Heidelberg. 267-281.
28. F. Lopes, F.C. Delicato, T. Batista, E. Cavalcante, T. Pereira, P.F. Pires, P. Ferreira, R. Mendes, OpenCOPI: middleware integration for Ubiquitous Computing. *International Journal of Parallel, Emergent and Distributed Systems*, **29**(2) (2014), 178-212. <https://doi.org/10.1080/17445760.2013.831415>
29. T.G. Stavropoulos, K. Gottis, D. Vrakas and I. Vlahavas, aWESoME: A web service middleware for ambient intelligence. *Expert Systems with Applications*, **40**(11) (2013), 4380-4392.
30. A. Olaru, A.M. Florea, and A.E.F. Seghrouchni, A context-aware multi-agent system as a middleware for ambient intelligence. *Mobile Networks and Applications*, **18**(3) (2013), 429-443. <https://doi.org/10.1007/s11036-012-0408-9>

31. J. Dooley, M. M. Henson, V. Callaghan, H. Hagra, D. Al-Ghazzawi, A. Malibari A, M. Al-Haddad M, and A.A. Al-Ghamdi, A formal model for space based ubiquitous computing. *In 2011 Seventh International Conference on Intelligent Environments*, (2011), 74-79. IEEE.
<https://doi.org/10.1109/ie.2011.63>
32. J.-E. Simó-Ten, E. Munera, J.-L. Poza-Lujan, J.-L. Posadas-Yagüe, and F. Blanes, “CKMultipeer: Connecting Devices Without Caring about the Network,” *Advances in Intelligent Systems and Computing*, (2017), 189–196.
33. J. Eclipse, 2015, <http://www.eclipse.org/jetty>, retrieved: June 2015.

Figure Legends

Figure 1: An example of a two-tiered middleware architecture for controlling a smart device

Figure 2: The iClassroom intelligent environment and its middleware architecture

Figure 3: The iSpace intelligent environment and its middleware architecture

Figure 4: The EDK middleware architecture

Figure 5: An Example of Communication and Controller Constants.

Figure 6: Creating a New Single Variable Smart Device

Figure 7: Creating a New Multi-Variable Smart Device

Figure 8: Creating the Communications System

Figure 9: Loading device control systems

Figure 10: Creating a Hub and Adding Master Devices

Figure 11: Enabling the EDK Web Services

Figure 12: Code for creating an EDK Control Point instance

Figure 13: Searching for Known Master Devices on an EDK Network

Figure 14: Three Methods for Finding Specific Devices

Figure 15: Methods for Reading the Values of Device State Variables

Figure 16: An example of sending action state change requests to a master device.

Figure 17: A Controller Template for a Single Variable Smart Device

Figure 18: A Controller Template for a Dimmable Light, (multi-variable smart device)

Figure 19: A graph showing average processing times for a sensor request (experiment one)

Figure 20: A graph showing average processing times for setting a light (experiment two)

Table 1 (on next page)

Average middleware processing times for taking a reading from a light sensor

Table 1: Average middleware processing times for taking a reading from a light sensor

	Average Processing Time Required (milliseconds)				
	Session 1	Session 2	Session 3	Session 4	Session 5
Jetty-based Web Services	80.03	77.43	55.58	61.48	60.65
Youpi UPnP	52.51	35.51	99	22.01	106.56
Web Services (EDK)	1.99	1.66	1.61	1.64	1.40
Direct API Control (EDK)	0.00542	0.00532	0.00516	0.00517	0.00519

Table 2(on next page)

Average middleware processing times changing the brightness of a dimmable light

1 **Table 2:** Average middleware processing times changing the brightness of a dimmable light

	Average Processing Time Required (milliseconds)				
	Session 1	Session 2	Session 3	Session 4	Session 5
Jetty-based Web Services	127.09	133.46	112.57	113.92	118.29
Youp! UPnP	171.36	232.33	212.22	238.08	180.49
Web Services (EDK)	101.83	102.45	102.35	102.26	102.09
Direct API (EDK)	0.67	0.64	0.60	0.59	0.55

2

Figure 1

Figure1: An example of a two-tiered middleware architecture for controlling a smart device

Figure1: An example of a two-tiered middleware architecture for controlling a smart device

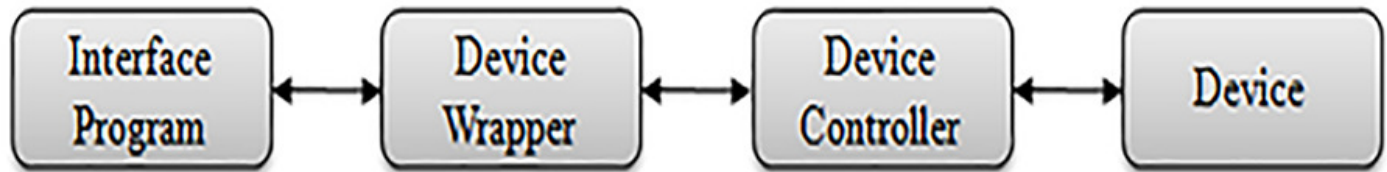


Figure 2

Figure 2: The iClassroom intelligent environment and its middleware architecture

Figure 2: The iClassroom intelligent environment and its middleware architecture

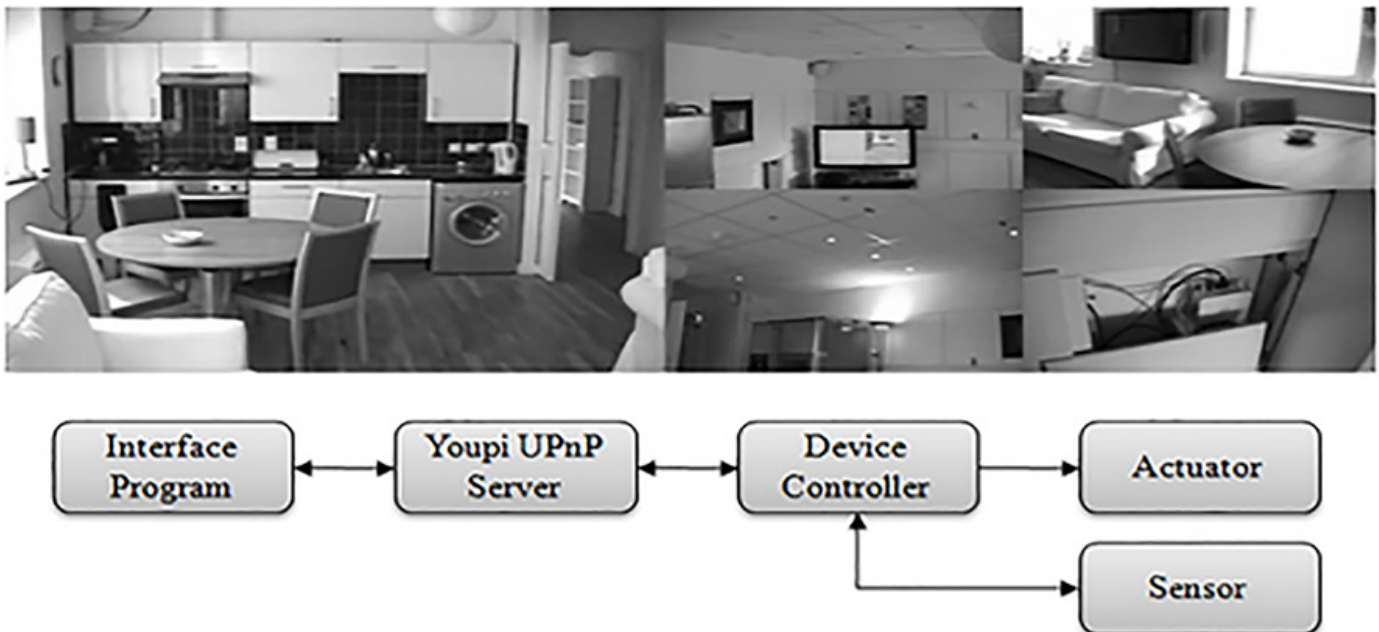


Figure 3

Figure 3: The iSpace intelligent environment and its middleware architecture

Figure 3: The iSpace intelligent environment and its middleware architecture

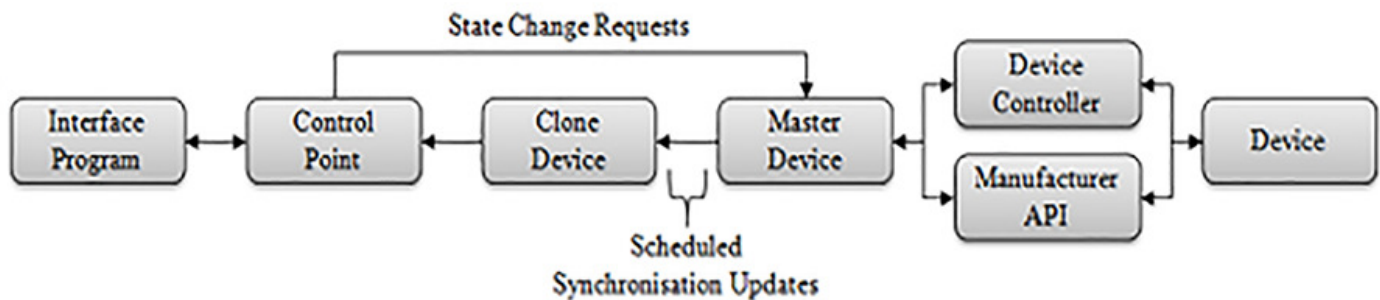


Figure 4

Figure 4: The EDK middleware architecture

Figure 4: The EDK middleware architecture

```
private final int COMMUNICATIONS_PORT = 4446;
private final int WEBSERVICES_PORT = 8000;
private final String CONTROLLER_CLASS = "Controller";
private final String CONTROLLER_PACKAGE = "devicecontroller";
private final String CONTROLLER_DIRECTORY = "controllers";
private final String NETWORK_ADDRESS = "230.0.0.1";
```

Figure 5

Figure 5: An Example of Communication and Controller Constants

Figure 5: An Example of Communication and Controller Constants

```
BooleanLight light1 =
    new BooleanLight("Light1", "Power", "Light Mk 1");
light1.addStateChangeListener(sl);
```

Figure 6

Figure 6: Creating a New Single Variable Smart Device

Figure 6: Creating a New Single Variable Smart Device

```
DimmableLight dimmer1 = new DimmableLight("Dimmer1",
    new String[] { "Power", "Brightness" }, "Dimmer Light Mk 1");
dimmer1.addStateChangeListener(sl);
```

Figure 7

Figure 7: Creating a New Multi-Variable Smart Device

Figure 7: Creating a New Multi-Variable Smart Device

```
DeviceHub lightHub = new DeviceHub(communications);
lightHub.addDevice(light1, lightController);
lightHub.addDevice(dimmer1, dimmerController);
```

Figure 8

Figure 8: Creating the Communications System

Figure 8: Creating the Communications System

```
InetAddress environmentAddress =
    InetAddress.getByName(NETWORK_ADDRESS);
Communications communications =
    new Communications(environmentAddress, COMMUNICATIONS_PORT);
```

Figure 9

Figure 9: Loading device control systems

Figure 9: Loading device control systems

```
File controllerJARFile = new File(CONTROLLER_DIRECTORY +
    File.separator + "DeviceController.jar");
Controller lightController = new Controller(
    controllerJARFile, CONTROLLER_PACKAGE, CONTROLLER_CLASS);
```


Figure 10

Figure 10: Creating a Hub and Adding Master Devices

Figure 10: Creating a Hub and Adding Master Devices

```
DeviceHub lightHub = new DeviceHub(communications);
lightHub.addDevice(light1, lightController);
lightHub.addDevice(dimmer1, dimmerController);
```

Figure 11

Figure 11: Enabling the EDK Web Services

Figure 11: Enabling the EDK Web Services

```
Webserver webServer = new Webserver(WEBSERVICES_PORT);
lightHub.addWebServices(webServer);
```

Figure 12

Figure 12: Code for creating an EDK Control Point instance

Figure 12: Code for creating an EDK Control Point instance

```
InetAddress environmentAddress =
    InetAddress.getByName(NETWORK_ADDRESS);
Communications communications =
    new Communications(environmentAddress, COMMUNICATIONS_PORT);

ControlPoint controlPoint =
    new ControlPoint(communications, "ControlPoint1");
```

Figure 13

Figure 13: Searching for Known Master Devices on an EDK Network

Figure 13: Searching for Known Master Devices on an EDK Network

```
Device[] deviceList = controlPoint.searchForDevices();
```

Figure 14

Figure 14: Three Methods for Finding Specific Devices

Figure 14: Three Methods for Finding Specific Devices

```
// Search for all instances of a specific device type
Device[] deviceList = controlPoint.
    searchForDevicesByType(Actuator.BOOLEAN_LIGHT);

// Search for a single device with the specified name
Device device = controlPoint.searchForDeviceByName("Light1");

// Search for a single device with the specified UUID
UUID uuid =
    UUID.fromString("2deb19fc-5d41-4ac9-a0ea-4e86ff7a121e");
Device device = controlPoint.searchForDeviceByUUID(uuid);
```

Figure 15

Figure 15: Methods for Reading the Values of Device State Variables

Figure 15: Methods for Reading the Values of Device State Variables

```
String deviceState = device.getState();

String powerState = device.getState("Power");

int brightnessState =
    Integer.parseInt(device.getState("Brightness"));

int brightnessState =
    ((DimmableLight) device).getBrightness();
```

Figure 16

Figure 16: An example of sending action state change requests to a master device

Figure 16: An example of sending action state change requests to a master device

```
for (Device device : deviceList) {
    if (device.getName().equalsIgnoreCase("Dimmer1") &&
        device.getType().equals("DimmableLight")) {
        if (device.getState("Power").equalsIgnoreCase("0")) {

            ((DimmableLight) device).turnOn();

        }

        else {

            ((DimmableLight) device).turnOff();

        }
    }
}
```

Figure 17

Figure 17: A Controller Template for a Single Variable Smart Device

Figure 17: A Controller Template for a Single Variable Smart Device

```
package devicecontroller;

public class Controller {

    public Controller() {

    }

    public String getState() {

    }

    public String setState(String state) {

    }

}
```


Figure 18

Figure 18: A Controller Template for a Dimmable Light, (multi-variable smart device)

Figure 18: A Controller Template for a Dimmable Light, (multi-variable smart device)

```
package devicecontrollermultiple;

public class Controller {

    public Controller() {

    }

    public String getBrightness() {

    }

    public String getState() {

    }

    public String setBrightness(String value) {

    }

    public String setState(String state) {

    }

}
```

Figure 19

Figure 19: A graph showing average processing times for setting a light (experiment one)

Figure 19: A graph showing average processing times for setting a light (experiment one)

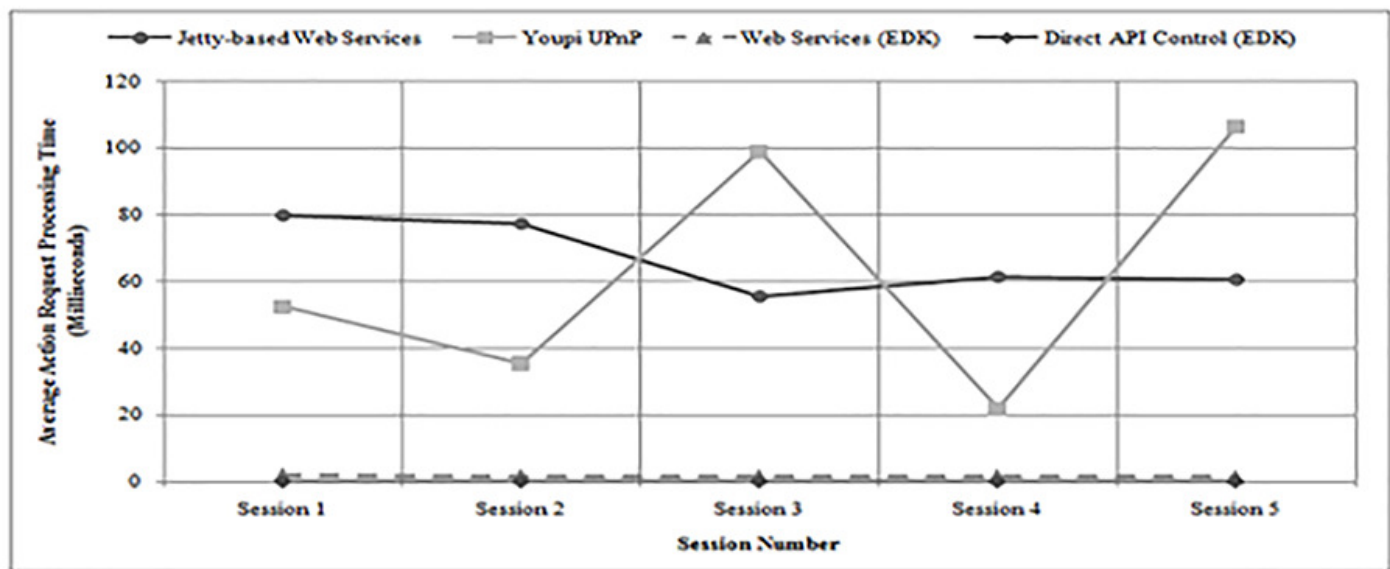


Figure 20

Figure 20: A graph showing average processing times for setting a light (experiment two)

Figure 20: A graph showing average processing times for setting a light (experiment two)

