Distributed Software Infrastructure for General Purpose Services in Smart Grid

Original

Distributed Software Infrastructure for General Purpose Services in Smart Grid / Patti, Edoardo; Angeliki Lydia Antonia, Syrri; Marco, Jahn; Pierluigi, Mancarella; Acquaviva, Andrea; Macii, Enrico. - In: IEEE TRANSACTIONS ON SMART GRID. - ISSN 1949-3053. - 7:2(2016), pp. 1156-1163. [10.1109/TSG.2014.2375197]

Availability:

This version is available at: 11583/2577139 since: 2016-04-10T17:16:41Z

Publisher:

IEEE

Published

DOI:10.1109/TSG.2014.2375197

Terms of use:

openAccess

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2016 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Distributed Software Infrastructure for General Purpose Services in Smart Grid

Edoardo Patti, Angeliki Lydia Antonia Syrri, Marco Jahn, Pierluigi Mancarella, *Senior Member, IEEE*,
Andrea Acquaviva, *Member, IEEE*, Enrico Macii, *Fellow, IEEE*

Abstract—In this paper, the design of an event-driven middleware for general purpose services in Smart Grid is presented. The main purpose is to provide a peer-to-peer (P2P) distributed software infrastructure to allow the access of new multiple and authorized actors to Smart Grid's information in order to provide new services. To achieve this, the proposed middleware has been designed to be i) event-based, ii) reliable, and ii) secure from malicious Information and Communication Technology (ICT) attacks, as well as iv) to enable hardware independent interoperability between heterogeneous technologies. To demonstrate practical deployment, a numerical case study applied to the whole UK distribution network is presented and the capabilities of the proposed infrastructure are discussed.

Index Terms—Middleware, Distributed Systems, Pervasive Computing, Demand Response, Distribution network, Aggregation

I. INTRODUCTION

THE concept of the Smart Grid is pervading all levels of the power system chain with the aim of facilitating the pathway towards more sustainable, economical and reliable networks by deploying low carbon technologies and advanced ICT options. However, this requires rethinking the entire control approach to power systems, particularly in distribution networks, where many of the major changes are likely to happen and many renewable energy sources, electric vehicles, storage, and so forth, will be connected. Also, new commercial structures will be needed to enable new actors such as aggregators, virtual power plants, energy service companies, etc, to participate in a fast-evolving distributed marketplace. In this context, research is needed to develop optimal ICT infrastructures that could facilitate interactions among all the relevant actors and different controllable network devices and technologies for provision of different services. In particular, while recent development of Ubiquitous Computing [1] and Internet of Things [2] concepts and relevant technologies could help address this challenge by providing means to seamlessly interact with distributed sensors and actuators, a

key open point remains as to how to achieve true interoperability between heterogeneous devices and facilitate access to data and in case controls to multiple parties. Middleware technologies and service-oriented architectures seem to be promising options along this direction.

On these premises, the aim of this paper is to introduce a comprehensive framework for the development of a distributed real-time event-based software infrastructure that could involve different actors in a Smart Grid (SG) context. More specifically, a novel design of an event-driven service-oriented middleware is proposed, whose main objectives are to i) provide easy integration of heterogeneous technologies, both wireless and wired; ii) enable hardware-independent interoperability across these technologies; iii) facilitate the access of multiple actors to both control technologies and relevant data to foster competition in the (distributed) marketplace to provide various power systems services; and iv) enable interoperability with also third-party software exploiting a web services approach which could facilitate further general purpose services and business cases. At the same time, the proposed solution intrinsically features secure and trusted communication between different actors. Also, scalability is guaranteed thanks to a publish/subscribe approach [3] so that different actors can access the same information coming from the middleware for different purposes without affecting others.

Exploiting this paradigm, and to illustrate an application of the proposed concepts, Demand Response (DR) could be put forward as an example of interaction between end-users, system operators, retailers and so forth. In fact, DR could be *utility driven*, for instance contributing to distribution network capacity support, reducing operational costs, and improving system reliability [4]; or *customer/market driven* [5], [6], where customers may adapt their load level in response to real-time pricing. Altogether, DR could be a useful controllable product for wholesale market and transmission/distribution system operators [7], including for minimisation of the spinning reserve from partially loaded generators [8], real-time balancing [9], and corrective control [10]. In this context, different DR providers in different or similar geographical areas may interact with different market actors for provision of different services. On the other hand, different parties may need access to data from the same DR provider for different applications (for instance, aggregation of reserve services and consumption measurements). Additionally, DR dispatch notices could vary from minutes (balancing services) to day ahead (preventive constraint management), depending on the service it is called to provide [11], and these notice times

This work was supported by the DIMMER FP7 European research project and by the UK Engineering and Physical Science Research Council, grant reference EP/I031/650/1, as a part of the Autonomic Power System Project.

E. Patti, A. Acquaviva and E. Macii are with the Dept. of Control and Computer Engineering, Politecnico di Torino, Italy. Emails: {edoardo.patti, andrea.acquaviva, enrico.macii}@polito.it

A.L.A. Syrri and P. Mancarella are with the Dept. of Electrical and Electronic Engineering, The University of Manchester. Emails: {angelikilydiaanton.syrri, p.mancarella}@manchester.ac.uk

M. Jahn is with the Dept. of User Centered Computing, Fraunhofer Institute for Applied Information Technology FIT, Germany. Email: marco.jahn@fit.fraunhofer.de

need to be properly considered in the design of the supporting ICT infrastructure. It is in the attempt to manage all this complexity of multiple parties and services that the benefits of the proposed middleware platform can fully emerge.

The rest of the paper is organized as follows. Section II reviews relevant background literature. Section III introduces the proposed middleware for general purpose services. Section IV presents an example of deployment in distribution networks. Section V discusses the capabilities of the ICT infrastructure that supports the service provided, quantifying the appropriate number of actors and middleware subsystems involved. Section VI provides the concluding remarks.

II. RELATED WORK

Recent development of Ubiquitous Computing (UbiComp) [1] and Internet of Things (IoT) [2] technologies can help address the challenge of moving towards a fully operated SG by providing means to seamlessly interact with distributed sensors and actuators. In this context, a key challenge that remains is achieve true interoperability across heterogeneous devices and between different applications. Service Oriented Architectures (SOA) seem to be promising along this direction [12], [13]. In addition middleware can be useful for developing SG solutions and services that exploit new data sources [14], [15].

At the building level, middleware solutions have been developed in order to achieve the interoperability across heterogeneous technologies [16], [17] also exploiting event-driven and user-centric approaches [18]. They also provide to authorized entities a set of API (Application Programming Interfaces) in order to integrate buildings in the Smart Grid system and enable the communication between them.

Kim et al. [19] present a data-centric middleware to allow decentralized monitoring and control, exploiting a publish/subscribe model [3], which is appropriate for delivering information but is not yet sufficient to have data access that is independent of this model. Indeed, other communication approaches, like SOA, are needed in order to provide new SG services that can easily retrieve information without having to wait for new events.

CoSGrid [20] is a middleware for measuring and controlling the electrical power of heterogeneous SG infrastructures. The communication across the entities in the grid is enabled by exploiting a remote method invocation and an event notification approach. However, the communication flows are not protected from malicious threats. Indeed CoSGrid does not implement any feature to make secure the communication channel.

GridStat [21], [22], [23] is another example of middleware for Smart Grids. Like the proposed solution, it exploits the event-subscribe approach. Moreover, it provides support to its application for QoS (Quality of Service), which is the ability to provide, in the communication, different priority to the data flows. However, the middleware works with its own closed and dedicated network infrastructure [24], which is incompatible with the existing IP-based infrastructures (Internet), so new routers and devices must be deployed.

Finally in [25], Salvadori et al. propose an ICT infrastructure for Smart Grid, which integrates a set of smart sensors and

Fig. 1. Architectural scheme for the proposed middleware

communication systems for different applications. It consists of a hardware platform that receives data from the sensors via wireless network or through physical network and then forward them to a control systems through Ethernet or RS232. However, following the vision of UbiComp and IoT, in large Smart Grid applications this is not sufficient because it must be open to any kind of commercial technologies both wireless and wired.

With respect to the presented solutions, the proposed middleware enables true interoperability between heterogeneous protocols and devices, both wireless and wired, providing real hardware abstraction. Moreover, exploiting the existing IP networks, it enables a peer-to-peer (P2P) [26] software infrastructure based on both publish/subscribe model and SOA. Furthermore, it provides features to enable secure and trusted communication between the peers. Finally, it provides a Rule Framework to easily develop control policies.

III. MIDDLEWARE FOR GENERAL PURPOSE SERVICES IN SMART GRID

In the world of Ubiquitous Computing [1] and Internet of Things [2], one of the main issues concerns the coexistence of several heterogeneous technologies and consequently their interoperability. Future SG systems will be UbiComp and IoT environments that have to deal with multiple and different actors (such as devices, applications and technologies) to provide services. To cope with these issues and to be open to future developments, we employ a middleware approach. Starting from the open source LinkSmart middleware [27], which is a generic service-oriented middleware for UbiComp, we propose here the design of a middleware for general purpose services in Smart Grid. As shown in Figure 1, it consists of a three-layered architecture with i) an *Integration Layer*, ii) a *Services Layer* and iii) an *Application Layer*. The middleware provides developers with a set of components, called managers. They are designed exploiting a SOA approach, and each manager exposes its functionalities as Web Services. Hence, the proposed middleware is a service-oriented distributed infrastructure consisting of a collection of software components which aims to i) allow interoperability across het-

erogeneous technologies and ii) provide tools, Web Services and APIs for the development of distributed applications. It stands between the user application and the heterogeneous devices and technologies. The rest of this section describes each layer of the proposed infrastructure in more detail.

A. The Integration Layer

The proposed infrastructure leverages upon an ICT infrastructure made of heterogeneous monitoring and actuation devices, both wireless and wired, which exploit different communication protocols and standards, such as ZigBee, EnOcean or BACnet. The *Integration Layer* exploits the concept of *Integration Proxy* to enable interoperability across heterogeneous technologies. This is the milestone to develop systems which are suitable for Smart Grid. More specifically, the *Integration Proxy* is a middleware-based software component that acts as a bridge between the middleware network and the underlying technologies, devices or subsystems. Each technology needs its own Integration Proxy to export its functionalities as Web Services. Hence, the Integration Proxy is the key to ensure communication between heterogeneous devices and allows us to use each low-level technology transparently inside the middleware network. Specifically, the Integration Proxy is a software component that runs on a PC and communicates directly with the heterogeneous networks receiving real-time information from various devices, regardless of the adopted communication protocols, hardware or network topology. Once the information is received and interpreted by the Integration Proxy, this is immediately sent to the middleware network exploiting the publish/subscribe approach provided by the Event Manager (see Section III-B3).

In a nutshell, the *Integration Layer* of the proposed middleware consists of several Integration Proxies, one for each technology. We developed Integration Proxies to manage Wireless Sensor and Actuator Networks (WSAN) which exploit the following protocol stacks: i) IEEE 802.15.4, ii) ZigBee and iii) EnOcean. In addition, we developed an Integration Proxy to allow the interoperability with the OPC Unified Architecture [28], which incorporates all the functionalities provided by different standards, such as BACnet. Hence, the backwards compatibility with wired technologies is enabled and integrated into our middleware.

B. The Services Layer

The *Services Layer* provides components specifically designed for general purpose services in SG, which should support the management of reoccurring tasks.

1) *The Network Manager*: The middleware, through the LinkSmart Network Manager, allows direct communication between all the applications inside its network, even if they are behind a firewall or NAT (Network Address Translator). Web Service calls are routed through the Network Manager, which creates a SOAP (Simple Object Access Protocol) tunnel to the requested service endpoint [29].

2) *The Trust and Crypto Managers*: Ardito et al. [30] emphasize how ICT in Smart Grid is "a decentralized network, where intelligence is distributed across several devices" and/or actors. It also introduces relevant issues related to security, which must not be neglected. The proposed middleware already comes with features to enable a secure and trusted communication between different actors [31]. The Trust Manager controls whether a device or service in the P2P LinkSmart network can be trusted or not. Therefore, it enables *mutual authentication* between actors by providing the means to create a Public Key Infrastructure (PKI). Hence, malicious peers cannot call services in the middleware network and cannot receive any kind of data. The Crypto Manager allows cryptographic operations used for message protection exploiting symmetric and asymmetric encryption in order to guarantee the *confidentiality* between the parties. In addition, it can sign each message with digital certificates providing *integrity of data*.

3) *The Event Manager*: In an event-based communication approach, the Event Manager provides a data centric model based on the publish/subscribe service [3] for the middleware Web Services. This allows the development of loosely-coupled event-based systems. This approach decouples the production and consumption of the information by removing all the explicit dependencies between the interacting entities, which increases scalability. In Smart Grids, where we deal a lot with events coming from both devices and distributed software, this mechanism is a key requirement to develop systems and applications.

The Event Manager provides us with the functionality of a topic-based publish/subscribe mechanism [3] for LinkSmart Web Services. Hence, each event contains both measurement and timestamp and it is published under a certain topic. The event topic has a hierarchical format, which also provides some basic semantic information about the type of event. An event topic for publishing a simple power consumption measurement would look like this:

MEASUREMENT/SENSOR/1234/PowerConsumption

where *MEASUREMENT/SENSOR* is an identifier for the type of event, 1234 is the sensor id and *PowerConsumption* is the type of measurement. It is worth noting that, following this approach also other middleware software components, and not only sensors, can publish events just by changing the event topic.

Using this kind of event topic format, software components interested in certain events can subscribe for them. Moreover, wildcards can be used for subscription to groups of events. For example, an application that would be interested in all sensor events (like a central persistence application) could subscribe for the topic *MEASUREMENT/SENSOR/**.

4) *The Semantics Framework*: It enables semantic interoperability across heterogeneous devices and technologies. The middleware provides managers to store, access, and update semantic knowledge about an application domain (or even across different domains) and the implemented system. In our case we modelled knowledge and meta-data about sensors and actuators as well as their relation to domain model objects

such as appliances, grid substations or buildings. These data are modelled and managed by well-known semantic web technologies, adhering to existing standards, namely Web Ontology Language (OWL). Knowledge is made available to application developers, allowing them to query any kind of information from a rich domain model. This could be the location or capabilities of a sensor but also, for instance, a list of all sensors in a specific grid substation, or an actuator with a certain control capability.

5) *The Discovery Manager*: It is responsible for discovering locally available devices that are connected via an Integration Proxy using WS-Discovery protocol. Once an Integration Proxy is discovered, semantic information is extracted and used by the Discovery Manager to update its knowledge base, which contains the global knowledge about devices in the network (utilizing the aforementioned Semantics Framework). When several Discovery Managers are available in the P2P middleware network they synchronize their knowledge about devices so that all of them have the same knowledge about the available devices connected to the same middleware network.

6) *The Rule Framework*: Typical power system management functions can be expressed in rules: the system listens to certain events, processes them based on given knowledge and algorithms and performs a resulting action. Hence, a specific control strategy can be developed by putting together different basic rules. The Rule Framework allows a fully flexible implementation of any kind of rule-based system. The framework provides standard interfaces as a basis for specific rule implementations. These rule implementations can be combined in a rule engine that executes the rules on incoming events. Rule logic and contextual information needed to execute a rule are kept separately, following the principle of the separation of concerns. This allows designers to reuse rule implementations in different contexts, e.g. to apply the same energy control policy in different subsystem, but with different settings, depending on the peculiarities of the subsystem itself.

C. The Application Layer

The *Application Layer* represents the highest layer of the proposed infrastructure. It provides a set of API to develop distributed event-based applications in order to manage the grid and post process data coming from the lower layers. At that level, interoperability is enabled between different devices as well as, thanks to the Web Service approach, between third party software. Hence, different applications for several actors (such as aggregators, energy suppliers and system operators) can be developed to provide general purpose services down to the single appliance in a house. Furthermore, in order to avoid huge ICT network overheads due to transmission of such fine grained information, data aggregation applications can also be developed to aggregate information about some subsystems [8]. Similarly, exploiting the functionalities provided by the Rule Framework, control policies could be designed and deployed across the SG in order to optimize the demand response process. Finally, each component of the proposed solution can be duplicated in the middleware network providing reliability from the software side. Hence, these

Fig. 2. Distribution Network under middleware deployment

properties, jointly with the ICT security features described in Section III-B2, allow the development of robust applications for monitoring and management in a SG context.

IV. AN EXAMPLE OF P2P COMMUNICATION PARADIGM IN DISTRIBUTION NETWORKS

Existing Distribution Networks (DNs) could be characterized by their strict hierarchical infrastructure, with a centralized control system and mainly one-way communication. In order to move forward to the Smart Grid vision, a distributed control approach is needed. Moreover, ICT technologies must be taken into account to enable two-way communication not only between DN entities but also between various entities or actors (e.g. energy suppliers, aggregators, etc.) that can be involved to provide new services.

Figure 2 shows a distributed approach for moving towards a fully operated Smart Grid. It shows the DN, operated by the Distribution System Operator (DSO), and the other components involved, from the substation and the transformer level, down to the customer's meters at building or home level. In order to exchange information across the SG, a P2P [26] communication network topology shall be available. The P2P communication paradigm is a "self-organizing of equal, autonomous entities (peers) which aims to shared usage of distributed resources in a networked environment avoiding central services" [26]. Hence, each peer acts simultaneously as supplier and consumer of resources enabling the communication directly with another peer. We propose to exploit the middleware introduced in Section III to enable a P2P communication network across different entities and actors that are the peers for the proposed infrastructure. Moreover, the proposed solution enables a distributed management of the grid taking into account also other factors, such as renewable or other distributed energy resources.

Following this approach, the Integration Proxy and the Event Manager become the two main middleware components. In fact, they ensure interoperability across heterogeneous devices and enable data centric communication between different actors, respectively. In addition, two P2P middleware-enabled applications can be developed and deployed to manage each subsystem:

- The *Data Aggregator*. It provides an aggregation of real-time consumption information coming from heterogeneous sources both hardware, thanks to the

Fig. 3. Example of Peer To Peer Communication

Integration Proxy, and software (e.g. other Data Aggregators in the infrastructure).

- The *Control Policy Manager*. It implements the control strategies to manage its subsystems. It receives and process real-time information from heterogeneous devices, Data Aggregators and/or other applications before taking decisions and sending the corresponding actuation commands. Furthermore, it can also receive or send action commands from/to other Control Policy Managers, again exploiting the Event Manager.

As shown in Figure 2, we propose to introduce these new components at each level of the DN also including buildings and homes. At DSO-, Substation- and Transformer-level, a Data Aggregator and a Control Policy Manager will be deployed in addition to an Integration Proxy and an Event Manager. Thanks to them, the DN can be divided in various P2P inter-connected subsystems, which will be able to exchange information with other distributed services or entities managed by different actors (see Figure 3). Moreover, thanks also to the integration of protocols for Building Management Systems [18] and heterogeneous commercial off-the-shelf devices, as depicted in Section III-A, fine grained monitoring and actuation up to building-, home- or appliance-level can be reached by replicating at the customer's meter-level the described software components. Hence at home- or building-level, Data Aggregator, Control Policy Manager, Integration Proxy and Event Manager should be deployed. In addition, scalability is guaranteed thanks to the publish/subscribe approach [3] adopted by the Event Manager, as described in Section III-B.

It is worth noting that the proposed infrastructure provides a system to enable communication also between different actors. So it does not matter who owns a certain subsystem or a certain application, because by exploiting the proposed solution the information can be easily sent to the middleware network and can be easily consumed by other actors, if authorized, to provide services.

A. P2P communication reliability in the proposed middleware solution

Our middleware enables the set-up of a P2P network where each peer is an actor and/or an entity of the SG. In this

Fig. 4. Example of P2P communication flow reliability

TABLE I
MAXIMUM ALLOWED NUMBER OF PEERS PER BANDWIDTH

Tech.	Bandwidth		Max number of peers		Max Coverage
	Download	Upload	Pubs	Subs	
Local Area Network					
Ethernet	10 Mbps	10 Mbps	1.5 k	1.5 k	100 m
	100 Mbps	100 Mbps	15.6 k	15.6 k	100 m
	1 Gbs	1 Gbs	156 k	156 k	100 m
WiFi	54 Mbps		8.5 k		300 m
	600 Mbps		93.5 k		1 km
Wide Area Network					
Optical Fiber	100 Mbps	100 Mbps	15.6 k	15.6 k	10 km
	662 Mbps	662 Mbps	103.4k	103.4k	60 km
	2448Mbps	2448Mbps	382.5k	382.5k	60 km
	1 Gbps	1 Gbps	156 k	156 k	20 km
DSL	8 Mbps	1.3 Mbps	1.2 k	203	4 km
	12 Mbps	3.5 Mbps	1.8 k	546	7 km
	24 Mbps	3.3 Mbps	3.7 k	515	7 km
	85 Mbps	85 Mbps	13.2 k	13.2 k	1.2 km
	200 Mbps	200 Mbps	31.2 k	31.2 k	1 km
WiMAX	128 Mbps	28 Mbps	20 k	4.3 k	10 km
	1 Gbps	1 Gbps	156 k	156 k	100 km
3G/4G	14.4 Mbps	5.75 Mbps	2.2 k	898	5 km
	84 Mbps	22 Mbps	13 k	3.4 k	
	326 Mbps	86 Mbps	50.9 k	13.4 k	
	1 Gbps	500 Mbps	156 k	78 k	100 km
Satellite	28 kbps		4		Depend on number of satellites
	128 kbps		20		
	450 kbps		70		

scenario, the Event Manager can be considered as a bottleneck for the whole information flow, as it is in charge of forwarding the data from publishers to subscribers. Indeed, if the EM crashes for any reason the information flow is interrupted. However, in our middleware network, more Event Managers can coexist together and each of them can handle different information flows. So, the EM is not a unique entity in the whole network topology. Moreover, the Network Manager (see Section III-B1) also provides features to make each middleware component reliable; so it can be duplicated and deployed in different servers. Therefore, multiple entities of the same EM that manages a specific information flow can be deployed in the network. Thus, as shown in Figure 4, if an EM fails for any reason, another duplicated entity will be automatically and transparently selected to ensure the communication without breaking the information flow between actors.

Another possible bottleneck could be the physical link to connect the EM to the Internet backbone. Indeed, depending on the bandwidth of the adopted technology, the amount of data that it can manage in terms of bytes changes. As shown in Table I, for each link technology we calculated the number of maximum allowed theoretical peers dividing its bandwidth for

our middleware event message size, which is almost 800 bytes (6400 bits). We also assumed that each peer can send/receive a message per second. Moreover, each technology exploits two different channels for download and upload. So, publishers (pubs) exploit the download channel and subscribers (subs) the upload one. On the contrary, WiFi and satellite technologies use a single channel for both download and upload.

In the Wide Area Network (WAN) the prevailing wired technologies are optical-fiber and DSL (Digital Subscriber Line), while the wireless are WiMAX, 3G/4G and satellite. So, if the Event Manager is connected to the Internet backbone via fiber optics, theoretically the link can manage from about 1.5k to 156k peers in download and from about 1.5k to 156k peers in upload. With DSL, it handles from almost 1.2k up to 31.2k publishers and from almost 203 to 31.2k subscribers. The WiMAX link manages from about 20k to 156k peers in download and from about 4.3k to 156k in upload. With a 3G/4G link, the EM can handle from almost 2.2k to 156k pubs and from almost 868 up to 78k subs. Finally, exploiting the communication based on satellite technologies, the Event Manager can manage from almost 4 to 70 peers both publishers and subscribers. Moreover, for wireless WAN technologies the number of sent/received events can decrease due to weather conditions that influence the performance of the link itself [32], [33].

If publishers and subscribers are in the same Local Area Network (LAN), the prevailing technologies are Ethernet and WiFi. Ethernet can manage from almost 1.5k up to 156k peers in download and from almost 1.5k up to 156k peers in upload. Finally, WiFi can handle in the same channel from almost 8.5k to 93.5k peers both pubs and subs.

V. CASE STUDY APPLICATION: THE MIDDLEWARE AS A PLATFORM FOR DR SERVICES

In order to illustrate the above concepts with a case study example, the middleware platform described in the previous sections will be used as the ICT support to provide real-time DR services. First, the case study will be presented and subsequently the feasibility of the proposed platform and the prerequisites for an extensive deployment throughout the UK will be studied. More specifically, this will be done by quantifying the features of the ICT support which could affect negatively the P2P communication reliability. However, it is worth noting that the proposed solution exploits the already existing Internet backbone and its deployment does not affect its correct functioning. In addition, it does not require major changes to the Distribution Network, except for the deployment of the middleware software components.

A. Description of the case study: DR for corrective control

In the following UK based case study, DR provides corrective control actions to DSOs when they would need to manage the network constraints following a fault. Depending on the type of DR programme and the size of the responsive load, DR could be activated directly by the DSO, or via a DR aggregator (upon receipt of a load control request). In this context, it is expected that all UK customers are

Fig. 5. Interactions between actors for DR corrective control services

equipped either with a smart meter, providing information about load consumption and available amount of DR, or a building management system, interfacing with a group of responsive appliances. Without loss of generality, hereinafter it will be assumed that for residential and commercial customers the service is provided through DR aggregators, whereas for industrial customers the service is provided directly to the DSO. Moreover, DR aggregators activate DR within three minutes upon request of the DSO, in order to contribute to corrective control in case of a system disturbance. For the residential customers participating in the scheme, the smart meter and the smart appliances are integrated in the middleware through specific Integration Proxies (as described in Section III-A). This makes it possible to control specific appliances at the home level independently, and to aggregate the measurements of all the controlled appliances. For each participating building a unique message is sent through the smart meter to the interested actor. This message includes the total load consumption of the customer, as well as available DR, which is the actual flexible load. The information flow is bidirectional i) between smart meters and DR aggregators that send the control commands and ii) between DR aggregators and a DSO that requests the service. The nature of the DR service implies that the information exchange between smart meters and DR aggregators should have per-minute frequency in order to respond efficiently to the system operators' instructions and have an accurate perception of the actual DR. All these interactions are depicted in Figure 5.

B. Dealing with the bottleneck of the physical link

The key factor that affects the middleware's efficiency is the capability of the physical link to send/receive data to/from the Internet backbone. Under the above assumptions, a DR aggregator may have to manage thousands of customers belonging to a DSO area. The same also applies to a DSO, which may have to interact with various aggregators as well as industrial customers. As a consequence, the bottleneck would appear when a DR aggregator or a DSO would like to receive data coming from the smart meters. For that reason in the

rest of the section we will identify the maximum number of customers that can be managed by an actor without violating the download bandwidth of the physical link. We will also identify the number of subsystems an actor should build when it has to manage a large number of customers or large amount of information. Since the DR service is provided to each DSO, UK is divided in 14 regions corresponding to the 14 UK DSO areas. Then, the appropriate number of actors operating in a specific region will be calculated. These actors could be the DR aggregators, but also any other actor providing a different service in the region. In any case, the information exchange between the customers and any actor that manages them would require the exploitation of WAN technologies. Hence, only the WAN technologies of Table I will be discussed, with the exclusion of satellites as their small bandwidth is not adequate for the case study application.

1) *Calculating the maximum number of customers per actor:* To estimate the number of customers in each region, it is assumed that the house density in the UK is 109 houses per Km^2 . Furthermore, the data packet sent from each house is about 800 bytes per minute, as mentioned previously. The maximum number of customers that an actor can subscribe to receive data from also depends on the technology the actor deploys to connect to the Internet backbone. This maximum number can be calculated as the ratio of the technology bandwidth to the data packet size sent by each customer. The results are shown in Table II, which shows the amount of contracted industrial customers that a DSO can control directly through the Event Manager. The same applies for any other actor who wants to deploy P2P communication with customers, which could be either distributed around the UK or located in a specific region.

2) *Calculating the required number of subsystems, as the intermediate layer between the customers and an actor:* In those circumstances when an actor would like to interconnect with a number of customers higher than the one in Table II, or when DR aggregators would need to manage a large amount of information coming from millions of smart meters, an appropriate number of P2P interconnected subsystems needs to be calculated. These P2P subsystems aim to support the information exchange and service provision on behalf of the DR aggregators. These subsystems act as an intermediate layer between the customer and the actor. Therefore, the so-called subsystems are also actors/customers of the DR services under study. They could be subsystems owned by the DR aggregators, but they could also be owned by other companies playing in the market and providing a service to the DR aggregators. Each subsystem manages its specific area, receiving signals from the DR aggregators and controlling buildings under its supervision. The above can be facilitated by the middleware-based application as follows: i) the data aggregator application is used to aggregate data coming from the area, and then send the aggregated packet to the DR aggregator; and ii) the control policy manager facilitates each subsystem to act independently, sending control commands to the customers it manages.

Assuming that all UK customers participate in the services presented, the appropriate number of subsystems is calculated

for each DSO region (for the case that local independent services need to be provided), and for the whole UK area. The number of subsystems is calculated as the ratio of the data packet size of the area to the bandwidth of the technology. All the above is summarized in Table III, providing the required number of subsystems (as the average number calculated for the bandwidth of each different WAN technology) that an actor needs to deploy in order to manage and receive the data coming from a specific region. For instance, if North

TABLE II
MAXIMUM NUMBER OF CUSTOMERS PER ACTOR FOR DIFFERENT TECHNOLOGIES

Technology	Download Bandwidth (Mbps)	Max Number of Customers
Optical-Fiber	100	983040
	662	6507725
	2448	24064819
	1000	9830400
DSL	8	78643
	12	117965
	24	235930
	85	835584
	200	1966080
WiMAX	128	1258291
	1000	9830400
3G/4G	14.4	141558
	84	825754
	326	3204710
	1000	9830400

West England wished to contract DR aggregators to manage the customers in the region: i) four DR aggregators should be contracted, or equivalently ii) a DR aggregator should deploy four subsystems. In the second case, the subsystems control all loads under their supervisory area and report to the DR aggregator, who in turn receives load request signals from the DSO. Needless to say, for a different number of houses/customers involved, the number of subsystems changes since it depends on the data size sent by the area.

C. Further discussions

Besides the presented services provided by the existing actors (DSO, DR aggregators), at any time a new actor could appear and subscribe to the information sent by the smart meters. Each new actor receives the events through the Event Manager, utilizing its own physical link, without any conflict with the other actors and without even knowing of their existence. For instance, energy suppliers in every region could exploit their own subsystems, using the data aggregator application, to aggregate electricity meter readings for billing purposes. As a second example, DSOs could subscribe to receive load consumption data from the DR aggregators throughout the year for quality of supply and network security purposes. Meanwhile, for actors interested in home-level services, they could subscribe to receive the original data published by the smart meters and also control heterogeneous home devices thanks to the middleware's Integration Layer.

VI. CONCLUSIONS

In this paper, a distributed software infrastructure for developing general purpose services in Smart Grid has been intro-

TABLE III

REQUIRED NUMBER OF SUBSYSTEMS FOR EACH DSO REGION AND THE WHOLE UK FOR AN AVERAGE COMBINATION OF WAN TECHNOLOGIES

DSO region	Number of DR customers	MBytes per minute	Number of Subsystems (for a combination of technologies)
East England	2084080	1590.03	6
East Midlands	1703343	1299.55	5
London	1274319	972.23	4
North Wales, Merseyside and Cheshire	325692	248.48	1
West Midlands	547180	417.47	2
North East England	936528	714.51	3
North West England	1543985	1177.97	4
North Scotland	4272092	3259.35	11
South Scotland	4272092	3259.35	11
South East England	2081355	1587.95	6
Southern England	6762578	5159.44	17
South Wales	2264911	1727.99	6
South West England	2597361	1981.63	7
Yorkshire	1297427	989.86	4
TOTAL	31962942	24385.79	80

duced. It aims to enable hardware-independent interoperability across heterogeneous devices. Moreover, it exploits a P2P communication paradigm to facilitate the access of multiple actors to the Smart Grid, thus allowing provision of services and information exchange between them.

The various middleware layers and its components have been introduced. Moreover, an example of deployment in distribution networks has been presented where it was shown that different buildings could be associated to different aggregators that provide different services to different actors, which is indeed the main strength of the proposed middleware concept. The actors could receive data directly from specific customers or aggregated data coming from subsystems, exploited throughout the network, interacting independently. In addition, the capabilities of the proposed distributed infrastructure in terms of maximum number of customers an actor can manage and relevant number of subsystems that are required were calculated for the DR service presented. The same method could be applied to any other service provided by another actor, ensuring scalability and reliability for all the reoccurring services and therefore truly enabling the development of a distributed market place in a Smart Grid context.

REFERENCES

- [1] M. Weiser, "The computer for the 21st century," *Scientific American*, 1991.
- [2] H. Kopetz, "Internet of things," *Real-Time Systems. Design Principles for Distributed Embedded Applications*, 2011.
- [3] P. T. Eugster, P. A. Felber, R. Guerraoui, and A. M. Kermarrec, "The many faces of publish/subscribe," *ACM Computing Surveys*, June 2003.
- [4] E. Martinez Cesena and P. Mancarella, "Distribution network reinforcement planning considering demand response support," in *Proc. of PSSC*, 2014.
- [5] A. Conejo, J. Morales, and L. Baringo, "Real-time demand response model," *IEEE Trans. on Smart Grid*, Dec. 2010.
- [6] C.-K. Woo, I. Horowitz, and I. Sulyma, "Relative kw response to residential time-varying pricing in british columbia," *IEEE Trans. on Smart Grid*, Dec. 2013.
- [7] F. Rahimi and A. Ipakchi, "Demand response as a market resource under the smart grid paradigm," *IEEE Trans. on Smart Grid*, June 2010.
- [8] K. Samarakoon, J. Ekanayake, and N. Jenkins, "Reporting available demand response," *IEEE Trans. on Smart Grid*, Dec. 2013.
- [9] P. Mancarella and G. Chicco, "Real-time demand response from energy shifting in distributed multi-generation," *IEEE Trans. on Smart Grid*, Dec. 2013.
- [10] A. L. A. Syri and P. Mancarella, "Reliability evaluation of demand response to increase distribution network utilisation," in *Proc. of PMAPS*, 2014.
- [11] National Grid, "Electricity ten year statement," Nov. 2012.
- [12] G. A. Pagani and M. Aiello, "Service orientation and the smart grid state and trends," *Service Oriented Computing and Applications*, 2012.
- [13] S. Karnouskos, "The cooperative internet of things enabled smart grid," in *Proc. IEEE ISCE*, Nov. 2009.
- [14] D. S. Markovic, D. Zivkovic, I. Branovic, R. Popovic, and D. Cvetkovic, "Smart power grid and cloud computing," *Renewable and Sustainable Energy Reviews*, 2013.
- [15] X. Fang, D. Yang, and G. Xue, "Evolving smart grid information management cloudward: A cloud optimization perspective," *IEEE Trans. on Smart Grid*, Mar. 2013.
- [16] T. G. Stavropoulos, K. Gottis, D. Vrakas, and I. Vlahavas, "awesome: A web service middleware for ambient intelligence," *Expert Systems with Applications*, Sept. 2013.
- [17] G. Candido, A. Colombo, J. Barata, and F. Jammes, "Service-oriented infrastructure to support the deployment of evolvable production systems," *IEEE Trans. on Industrial Informatics*, Nov. 2009.
- [18] E. Patti, A. Acquaviva, M. Jahn, F. Pramudianto, R. Tomasi, D. Rabourdin, J. Virgone, and E. Macii, "Event-driven user-centric middleware for energy-efficient buildings and public spaces," *IEEE Systems Journal*, 2014.
- [19] Y.-J. Kim, M. Thottan, V. Kolesnikov, and W. Lee, "A secure decentralized data-centric information infrastructure for smart grid," *IEEE Communications Magazine*, Nov. 2010.
- [20] D. Villa, C. Martin, F. Villanueva, F. Moya, and J. Lopez, "A dynamically reconfigurable architecture for smart grids," *IEEE Trans. on Consumer Electronics*, May 2011.
- [21] K. Tomovic, D. Bakken, V. Venkatasubramanian, and A. Bose, "Designing the next generation of real-time control, communication, and computations for large power systems," *Proc. of the IEEE*, May 2005.
- [22] C. Hauser, D. Bakken, and A. Bose, "A failure to communicate: next generation communication requirements, technologies, and architecture for the electric power grid," *IEEE Power and Energy Magazine*, Mar. 2005.
- [23] H. Gjermundrod, H. Gjermundrod, D. Bakken, C. Hauser, and A. Bose, "Gridstat: A flexible qos-managed data dissemination framework for the power grid," *IEEE Trans. on Power Delivery*, Jan. 2009.
- [24] D. Germanus, I. Dionysiou, H. Gjermundrod, A. Khelil, N. Suri, D. Bakken, and C. Hauser, "Leveraging the next-generation power grid: Data sharing and associated partnerships," in *Proc. of IEEE PES*, Oct. 2010.
- [25] F. Salvadori, C. Gehrke, A. de Oliveira, M. de Campos, and P. Sausen, "Smart grid infrastructure using a hybrid network architecture," *IEEE Trans. on Smart Grid*, Sept. 2013.
- [26] R. Steinmetz and K. Wehrle, "What is this peer-to-peer about?" in *Peer-to-Peer Systems and Applications*, ser. LNCS. Springer, 2005.
- [27] M. Eisenhauer, P. Rosengren, and P. Antolin, "A development platform for integrating wireless devices and sensors into ambient intelligence systems," in *IEEE SECON*, June 2009.
- [28] [Online]. Available: <https://opcfoundation.org/>
- [29] F. Milagro, P. Antolin, P. Kool, P. Rosengren, and M. Ahlsén, "SOAP tunnel through a P2P network of physical devices," in *Internet of Things Workshop*, Sept. 2008.
- [30] L. Ardito, G. Procaccianti, G. Menga, and M. Morisio, "Smart grid technologies in europe: An overview," *Energies*, 2013.
- [31] M. Hoffmann, A. Badii, S. Engberg, R. Nair, D. Thiemert, M. Matthess, and J. Schuette, "Towards semantic resolution of security in ambient environments," in *Developing Ambient Intelligence*. Springer, 2008.
- [32] N. Nor, I. Rafiqul, W. Al-Khateeb, and S. Zabidi, "Environmental effects on free space earth-to-satellite optical link based on measurement data in malaysia," in *Proc. of ICCCE*, July 2012.
- [33] S. Enoch and I. Otung, "Propagation effects in wimax systems," in *Proc. of NGMAST*, Sept. 2008.