



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

IoT software infrastructure for Energy Management and Simulation in Smart Cities

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

IoT software infrastructure for Energy Management and Simulation in Smart Cities / BRUNDU, FRANCESCO GAVINO; PATTI, EDOARDO; OSELLO, Anna; Giudice, Matteo Del; RAPETTI, NICCOLO'; Krylovskiy, Alexandr; Jahn, Marco; VERDA, Vittorio; GUELPA, ELISA; RIETTO, LAURA; ACQUAVIVA, ANDREA. - In: IEEE TRANSACTIONS ON INDUSTRIAL INFORMATICS. - ISSN 1551-3203. - STAMPA. - 13:2(2017), pp. 7740051.832-7740051.840. [10.1109/TII.2016.2627479]

Availability:

This version is available at: <https://hdl.handle.net/11585/818304> since: 2021-04-08

Published:

DOI: <http://doi.org/10.1109/TII.2016.2627479>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

F. G. Brundu et al., "IoT Software Infrastructure for Energy Management and Simulation in Smart Cities" in IEEE Transactions on Industrial Informatics, vol. 13, no. 2, pp. 832-840, April 2017

The final published version is available online at:

<https://doi.org/10.1109/TII.2016.2627479>

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

IoT software infrastructure for Energy Management and Simulation in Smart Cities

Francesco G. Brundu, Edoardo Patti, Anna Osello, Matteo Del Giudice, Niccolò Rapetti, Alexandr Krylovskiy, Marco Jahn, Vittorio Verda, Elisa Guelpa, Laura Rietto, Andrea Acquaviva

Abstract—This paper presents an IoT software infrastructure that enables energy management and simulation of new control policies in a city district. The proposed platform enables the interoperability and the correlation of (near-)real-time building energy profiles with environmental data from sensors as well as building and grid models. In a smart city context, this platform fulfills i) the integration of heterogeneous data sources at building and district level, and ii) the simulation of novel energy policies at district level aimed at the optimization of the energy usage accounting also for its impact on building comfort. The platform has been deployed in a real world district and a novel control policy for the heating distribution network has been developed and tested. Results are presented and discussed in the paper.

Index Terms—Internet-of-Things, middleware, energy saving, energy flow simulation, Smart City, distributed software infrastructure

I. INTRODUCTION

GLOBAL urban population reached more than 54% of the total global population [1]. For this reason, in an energy-intensive world, a key requirement consists on promoting and developing new control policies to optimize energy consumption in cities that include buildings and distribution networks (e.g. power and heating). For instance, thermal energy peaks smoothing requires information about building energy profiles as well as grid topology and characteristics, to evaluate their actual impact on the network infrastructures. This information can be made available through pervasive sensors installed both in buildings and in the grid. In case of district heating, building energy profiles have to be detected as well as grid topology and characteristics, with the purpose of reshaping thermal energy to reduce the peaks. This reshaping has an impact on building air temperature profile, that can be carefully evaluated to keep user comfort at an acceptable

level.

Furthermore, thanks to Internet-of-Things (IoT) [2] communication paradigms [3], monitoring data can be made available at the utility control centers so that complex policies can be implemented on cloud or cluster infrastructures.

Another relevant impact of IoT is the ease of integration of heterogeneous data sources that can be exploited to develop smarter policies. In this way, sensor data can be integrated with Building Information Models (BIMs) [4], grid models and Geographical Information Systems (GISs) [5]. BIMs are 3D models with construction and energy characteristics and they allow detailed energy simulations, that can be used to evaluate the impact of thermal reshaping on indoor air temperature. From the other side, building simulations can be compared with temperatures provided by indoor sensors to tune energy models, that can be used to evaluate policy impact on buildings where environmental sensors are not deployed.

In the last decade, several frameworks have been proposed to exploit IoT technologies at building and house level [6] [7] [8], and software solutions have been proposed to enable interoperability among various data formats and protocols [9] [10] [11]. However, the integration of district models with sensor data remains a challenge.

In this paper a distributed IoT platform able to collect, process and analyze energy consumption data and structural features of systems and buildings in a district is presented. In particular, data from Building Information Models [4] (BIMs), System Information Models [12] (SIMs) and Geographical Information Systems [5] (GISs) are interrelated and enriched with historical and (near-) real-time collection of data from heterogeneous IoT devices. Such devices are deployed across the district to monitor and manage the energy distribution systems (both power and heating).

To evaluate the software infrastructure, a real-world case study has been considered, with hundreds of devices connected to monitor and manage the heating distribution network. In particular, a policy aimed at rescheduling buildings energy requests to smooth peaks of district heating network has been simulated and evaluated. The IoT platform exploits data from sensors installed in heat exchangers (i.e. the connected “things”) to determine building energy profile, while network models are used to evaluate the peaks accounting for grid characteristics and topology. On the other side, user comfort impact due to rescheduling is measured with temperature sensors inside building. Sensor data are also

This work was supported in part by EU, FP7 SMARTCITIES 2013 DIMMER project (District Information Modelling and Management for Energy Reduction).

F.G. Brundu., E. Patti, A. Osello, M. Del Giudice, N. Rapetti, V. Verda, E. Guelpa, L. Rietto and A. Acquaviva are with Politecnico di Torino, Torino, 10129 Italy (e-mail: name.surname@polito.it).

A. Krylovskiy and M. Jahn are with Fraunhofer Institute for Applied Information Technology FIT, Sankt Augustin, 53754 Germany (e-mail: name.surname@fit.fraunhofer.de).

exploited to tune building models. This is relevant when models have to be used to evaluate the impact on inhabitants' comfort on buildings that are not instrumented with environmental sensors.

To summarize, the presented infrastructure acts as a district information management system, in which different stakeholders playing in a smart city scenario can collaborate and provide novel energy management policies and simulations.

This article is organized as follows. Section II reviews relevant background literature. Section III introduces the proposed IoT software infrastructure for energy management and simulation. Section IV draws the real-world case study. Section V introduces the energetic policies, which may be exploited thanks to the proposed solution, and discusses the experimental results. Finally, Section VI provides the concluding remarks.

II. RELATED WORK

Recently, different ICT solutions have been proposed to cope with the need of energy consumption optimization in Smart City domains. In the following the most relevant in this field are presented.

In a Smart Grid context, Kim et al. [13] presented a data-centric middleware to allow decentralized monitoring and control, exploiting a publish/subscribe model [14], which is appropriate for delivering information but is not yet sufficient to have data access that is independent of this model. Indeed, the request/response communication approach is also needed to provide novel services that can easily retrieve data without having to wait for new events.

In [15], a distributed software infrastructure for general purpose services in power systems is presented. The software architecture enables the interoperability across heterogeneous devices by creating a secure peer-to-peer network.

Differently from these solutions, the IoT platform presented in this paper aims at creating a virtual model of a city district for providing Smart City services. Hence, considering data coming from IoT devices deployed across the energy distribution networks is not yet enough. Indeed, such data have to be integrated and correlated together with information, often geo-referenced, about buildings.

In the Smart City scenario, several middleware solutions have been proposed to integrate heterogeneous data sources. The ReActOR system, presented in [16], is characterized by a tiny footprint and can be deployed as a service. It is composed by three layers: a facade layer (Web Service), a core layer, and an extensions layer (providing support for different technologies). It enforces user authorization, and each user is mapped to set of devices that he or she can manage, control, and sense. Finally, ReActOR provides only identifiers for the different agents in the system, hiding user and device data. This middleware is limited to hardware devices and it does not support the integration of different data sources, such as Database Management Systems. Cândido et al. [17] presents an evolvable and customizable infrastructure, focused on interoperability, modularity, uncomplicated management and adaptation, and composed by a set of different components.

Each component exposes its services to the network using open web standards, enabling the composition of interoperable modules. However, these approach is limited to industrial automation scenarios.

SensorGrid4Env [18] provides a service-oriented architecture that eases the design of open large-scale semantic-based sensor network applications for environmental management. Further, it enables the rapid development of thin applications (e.g. mashups) and it allows the integration of real-time data with historical data from other heterogeneous data-sources. This solution is tailored to environmental management and cannot be applied seamlessly to city district setting.

For what concerns the interoperability across various application domains, different research projects and initiatives contributed to the definition of models and guidelines. For instance, the IoT-A project [19] provides an IoT reference model, allowing the description of an IoT solution by using shared building blocks, a reference architecture and general advices to IoT architects. On the other hand, the OneM2M alliance [20] aims to develop detailed technical specifications, to address the need for a common M2M Service Layer, using existing IoT and Web standards. However, it does not cover many aspects of IoT platforms, such as scalability, availability and deployment. FI-WARE [21], another research project, is devoted to design a service infrastructure for the Future Internet vision. Such infrastructure would be composed by reusable components, which could be selected and complemented with additional specific components.

The novel contribution of the following work is the design and development of a distributed IoT platform, based on open standards of the Web, which creates a single virtual parametric district model. It is in charge of providing a standard access to heterogeneous IoT devices and District Information Systems deployed across the city. Information is correlated, enriched, and provided by the presented platform as a shared open repository. This IoT platform is instrumental for energy simulations on alternative control policies for the energy distribution network.

III. IOT SOFTWARE INFRASTRUCTURE

Building a distributed IoT platform to create a virtual parametric district model for energy management and simulation is a challenging task. Indeed, such platform needs to transparently integrate hundreds of heterogeneous data sources and IoT devices, that may be exploited to monitor and manage energy consumption. Furthermore, it has to be scalable and reliable. For this reason, a suitable model is needed to step back and see the districts at a lower level of detail or to focus on specific entities (e.g. buildings, IoT devices). The proposed infrastructure implements two communication paradigms: i) non-real-time request/response based on REST [22] and ii) (near-) real-time publish/subscribe based on MQTT [23].

This IoT software infrastructure is based on the LinkSmart[®] OpenSource Middleware [24], adopting and extending it to address the requirements of a Smart City context. The tasks of the platform include: i) modeling of real-world IoT devices and ICT systems, ii) providing their search and discovery by




Figure 1. Architectural schema for the proposed IoT software infrastructure for Energy Management and Simulation

applications, and iii) exposing their data via common application protocols. To implement these tasks, the platform defines several abstraction models and APIs implemented by a number of integration components and services described in this section. Figure 1 shows the architectural schema of the proposed software infrastructure. It consists of three layers: i) *IoT Devices and Technologies Integration layer*, ii) *Services layer* and iii) *Applications layer*. The rest of this section describes each layer in more detail.

A. IoT Devices and Technologies Integration Layer

The proposed IoT platform leverages upon ICT infrastructure made of heterogeneous devices and technologies, which exploit different communication protocols and standards. The *IoT Devices and Integration Layer* (bottom layer in Figure 1) takes advantage of Device Connectors to enable the interoperability among heterogeneous devices.

Device Connectors are integration components that use a common abstraction model to describe real-world IoT Devices and ICT systems integrating them in the middleware. The Device Connector is a middleware-based software component that acts as a bridge between the middleware network and the underlying technologies, devices, or subsystems [25]. It exposes Web Service APIs for discovering and querying ICT data sources for each integrated system allowing seamless data access by the applications and other services. Instead of implementing system-specific and proprietary APIs of the integrated systems, the middleware clients use the same Web Service API to access data from all integrated systems. The Device Connector works also as publisher and/or subscriber for the integrated IoT devices and ICT systems, thus providing a (near-) real-time communication (see Section B).

Finally, it exploits the SenML (Sensor Markup Language) data format, defined as draft by *The Internet Engineering Task Force* (IETF) [26], to transmit sensor measurements.

B. Services Layer

The *Services Layer* is the core of the proposed IoT software infrastructure. It provides components specifically designed for accessing and managing information, coming from




Figure 2. DIM Services schema

heterogeneous IoT devices and technologies, by exploiting a Web Service approach. It is also in charge of creating a virtual District Information Model correlating data from different data-sources. Finally, it offers features for simulating control policies aiming at optimizing energy consumption (e.g. electric, heating). The main components of this layer are described in the following.

1) *Service Catalog*: It describes the available services in the network, by exposing a JSON-based RESTful API. It is the entry point for applications, used to discover available services in the network. It contains entries of services such as the Resource Catalog and the MQTT Broker.

2) *Resource Catalog*: It provides a constantly updated registry of endpoints for IoT devices and resources which are available through the infrastructure. It exposes a JSON-based RESTful API to: i) Device Connectors, which register the endpoints to access devices and resources, and ii) Applications, which discover such devices and their access endpoints to access information. The Resource Catalog provides flexibility to the whole infrastructure because new IoT devices and ICT systems can be transparently added, removed or replaced in the system.

3) *Message Broker*: It provides (near-) real-time asynchronous communication for the different components of the infrastructure, by means of the MQTT (Message Queue Telemetry Transport) [23] communication protocol. This protocol follows the publish/subscribe [14] paradigm, which allows the development of loosely-coupled event-based systems. Each device or component in the platform can i) publish data, and ii) subscribe to specific event notifications. Thus, the *Message Broker* is in charge of receiving data from publishers and forwarding data to subscribers. This module can be exploited by either internal components or client applications. This approach increases the infrastructure scalability, by removing the dependencies between interacting entities.

4) *Historical Datastore*: Querying historical data is one of the basic features of the platform required by IoT applications and services. Considering the heterogeneity of the data sources in a typical Smart City platform, it is essential to provide a unified API for accessing historical sensor data in such systems. The platform defines a Historical Datastore API that

is implemented in a standalone service, as well as for each integrated ICT system that already provides a database with historical data and an API to access it. Hence, the standalone service implementation can be used for integration of new ICT or already existing data sources. In addition, it provides an optional aggregation functionality that can be used to down-sample high-frequency sensor measurements and calculate basic aggregates on them. The integrated system-specific implementations of the Historical Datastore API allow to provide a unified API for accessing historical sensor data across all integrated systems without data duplication. Integrating historical data from a large ICT system can be indeed problematic to manage due to both its large volume, as well as the storage and access policies. By providing system-specific implementations of the Historical Datastore API, both of these issues can be tackled by forwarding data requests to the integrated system and enforcing access policies between the middleware and the data owners. The Historical Datastore is able to store and manage huge amount of data. Thanks to its API, it is ready to provide such data to Big Data Analytics tools, that can help energy managers in developing control policies based historical data analysis.

5) *Semantic Metastore*: Semantic interoperability is another important feature of the platform, which enables the use of Semantic Web technologies to annotate and interlink integrated data sources and query them using the semantic attributes. In addition to that, such annotations can also be used for describing higher-level services and applications, as well as to model real-world entities they operate with that are not directly managed by the platform. This feature is implemented by the Semantic Metastore service, which builds on the off-the-shelf Semantic Web technologies such as Apache Jena [27] to provide higher-level, developer-friendly APIs for populating and querying the semantic knowledge base through REST Web Services. The Semantic Metastore may be queried by either internal components or client applications.

6) *DIM Services*: A District Information Model (DIM) is composed by different entities. In particular, it is possible to define three specific data sources, which are represented using relational databases, and can be integrated to get a comprehensive view of a district:

- **Building Information Models** (BIMs [4]) are parametric 3-Dimensional models, where each model describes a building - both structurally and semantically (e.g. by defining materials and costs [28]). BIM supports decisions for the whole building lifecycle;
- **Geographical Information Systems** (GISs [5]) map the geographical location and topology of district entities, such as district buildings or energy distribution networks: GIS provides data management and modelling for advanced cartography;
- **System Information Models** (SIMs [12]) describe size and structure of energy distribution networks. SIM is built to provide information for both visualization and simulation by exploiting parametric and topological data.

Despite the fact that all these data-sources could be represented using a relational database, no commonly agreed method is available to integrate and correlate them from the

district point-of-view. Because of such fragmentation, a single database for all heterogeneous data-sources cannot be deployed, because i) Database Management System (DBMS) technologies are mostly incompatible one another, and ii) the same value may have conflicting semantics between different databases. Instead, the following subsystem is heavily based on RESTful Web Services and distributed deployment. The use of RESTful Web Services provides a uniform interface to each component of the system. In particular, the same response format is returned, independently from the actual queried data source. This also means that, if a component technology (e.g. a DBMS) changes, the client application receives the same response as a result of the same query. A bird's-eye view of the DIM Services subsystem is shown in Figure 2.

The subsystem is built in a micro-service fashion [29], i.e. it consists of a suite of small, autonomous, lightweight services. From a logical point-of-view, each component in DIM services consists of two independent sub-layers: i) *Data layer* and ii) *Interface layer*. In the following, the structure of the BIM Service Provider is presented, which is easily extended to GIS and SIM Service Providers. The Data layer is a set of several databases, where each database represents a building in the district. The databases are managed by using a relational DBMS, which is able to store the internal representation of the BIM models (which have been created with proprietary software Autodesk Revit [30]). The Interface layer provides a RESTful Web Service interface to the underlying databases. It interacts with the Data layer by using SQL queries and preparing JSON results. Furthermore, it also provides building resources, such as proprietary (Revit files) and public BIM standards (IFC and gbXML).

As depicted in Figure 2 within the District Information Model other kinds of data sources are needed, to complement the BIM information. In particular, GIS and SIM have been identified. Both these data-sources need a Web Service interface, to be integrated into DIM Web Services by means of DIM context layer. Both the GIS and the SIM Service provider return JSON results through the REST API. On top of each Service provider, there is the DIM Context layer, which is the main interface to district information. It creates a virtual parametric model of the district that can be queried as a whole by client applications. It acts as a translator from the district context (district-based queries) to a specific data-source context (e.g. building-based queries). Conversely, it also integrates back responses from the different Service Providers to a single response. Moreover, the DIM Context layer is independent from each Service Provider Interface layer, and it behaves as a proxy to the latter. By using inter-provider queries, it is able to compose complex queries, which may span different information domains (e.g. building and geographical information). For instance, by querying both the BIM and the GIS Service Providers, it gives the capability of retrieving information about buildings composed by a specific number of storeys.

7) *Simulation engine*: It is a specific component of the proposed IoT software infrastructure used to help in developing new control policies and to perform energy simulations in a district environment. It consists of a Web Service enabled interface, which provides a standard entry-




Figure 3. Procedure for the analysis of alternative thermal request profiles

point to the district simulator. In addition, the Simulation engine provides features to develop control policies, able to retrieve information from all the other platform components exploiting both publish/subscribe and request/response communication approaches (MQTT and REST respectively). Hence, such simulations can take into account both District Information Models and historical or information collected (near-) real-time from the heterogeneous IoT devices to optimize the energy consumption. For example, this module can gather environmental information from the Historical Dastore that collects sampled data sent by IoT devices. The Simulation engine can work also as a subscriber to receive (near-) real-time data from IoT devices through the Message Broker. Information about district as a whole can be retrieved by DIM Services, which correlates information about BIM, GIS and SIM. To perform such correlations, the DIM Services also exploits semantic information retrieved from the Semantic Metastore.

It is worth noting that the Simulation engine module is able to perform simulations for the different energy flows in the district (e.g. electrical and heating). Section V provides a more in-depth description of control policies for the district heating network.

C. Applications Layer

The *Applications Layer* (the highest layer in Figure 1) provides a set of APIs and tools to develop applications to manage and post-process data coming from the underlying layers of the proposed IoT infrastructure. Such applications may be web-based (e.g. online dashboard for visual management of the district using cartographic maps), mobile (e.g. mobile app for facility managers), desktop and semantic web-clients. At this level, the interoperability among different IoT devices and technologies is enabled thanks to the Web Services approach.

IV. CASE STUDY

This IoT platform has been developed to provide smart city

services for energy management in a smart city context. As main case study, this platform has been deployed and tested in a real district. In such district, the heating distribution network (DH) provides thermal energy to about 50% of private and public buildings. Furthermore, Wireless Sensor Networks (WSN) have been deployed to collect indoor data about air temperature and relative humidity on seven representative buildings.

The heating network in the city under investigation is one of the largest in Europe, being composed by about 580 km of double pipeline, with more than 5700 buildings connected (over 55% of the town). The thermal request is highly variable during the day and during the heating season. In particularly daily request presents a morning peak of about 1300-1400 MW while the afternoon request ranges from about 850 MW in typical winter days to about 400 MW in typical middle season days. Such request is fulfilled mainly through three efficient cogeneration systems, able to produce a total thermal power of about 740 MW, while the exceeding request is covered using thermal storage units and auxiliary boilers.

Currently, in the selected district, a large number of IoT devices (about 4000) are deployed for monitoring and managing the DH. They collect data (e.g. instantaneous power, cumulative energy consumption, water flows and temperatures) every five minutes. This is possible for about 70% of the buildings connected with the network. The remaining buildings have been characterized assuming a behaviour equal to the average profile of the buildings in the network. This approach has been validated through comparison of the calculated and measured thermal request profiles at the plants. Data is then sent to remote Historical Dastore unit by means of embedded gateways, which monitor the heat exchanger of each building. Each gateway is accompanied by a GPRS modem and executes the following tasks: i) scheduling of data collection, ii) management of IoT sensors, iii) communication via GPRS, and iv) data transfer to remote *Historical Dastore*. The *Historical Dastore* for DH consists of four dispatchers, to improve reliability and scalability. Each measurement is sent to a different dispatcher, which forwards it to a database.

Differently, a significant part of the infrastructure consists of heterogeneous data sources, e.g. Building Information Models. Such models have been prepared for selected buildings and, the corresponding SIM, for the district heating network.

This paper assesses how this infrastructure is beneficial: i) to optimize the energy demand, ii) to manage a large number of IoT devices, and iii) to exploit of parametric models for energy simulations. In particular, BIMs allow fine-grained energy simulations, down to the room level, and are useful to plan retrofitting of existing buildings.

The most popular construction types in the selected district are load-bearing masonry (mostly used until the '50s) and reinforced concrete. Many cases present a mixed solution between the two technologies. For this reason, seven representative buildings have been considered and WSN sensors for indoor air temperature and relative humidity have been deployed on them. They exploit the Device Connector component to be integrated into the proposed software infrastructure. In particular, such sensors communicate using




Figure 4. Thermal loads in a typical winter day obtained by using the current request profiles and the modified request profiles

MQTT protocol [14] [23] and send collected data to a Historical Datastore.

For this case study, a Primary School in the selected district has been chosen. This primary school was originally built in 1902 and almost totally rebuilt at the end of the World War II. The building structure is a three storey load-based masonry with a heavy massive envelope. The main heating supply is based on radiators. Such rooms have been selected considering the orientation. Three representative rooms have been selected to better describe the air temperatures trends of the case study and the building shape: the first room (S) is oriented to South, the second room (E) to East and the third room (N) to North-West. Further, the room N is also located in the highest floor of the North-West corner of the primary school.

V. ENERGY OPTIMIZATION RESULTS

A. Methods and Policies

In order to reduce the primary energy consumption to supply heating to the buildings connected with the district heating network, the use of cogeneration should be maximized. This can be achieved through thermal storage [31] but further improvements can be achieved through optimization of the thermal request profiles of the buildings acting at district level. Such action is particularly important considering that the district heating network is in continuous evolution, with increasing number of buildings connected to the network. Instead, the built of new storage units has to face the issues related with land occupation and also the investment costs. The availability of the IoT infrastructure previously described allows the implementation of an innovative approach for the analysis of alternative thermal request profiles, which are able to produce similar effects than thermal storage units. In particular, by means of the Simulation Engine presented in Section III.B, it is possible to run thermal request simulations following the approach described in Figure 3, that involves all the buildings connected to the heating network in the district. Red boxes indicate the external input for the procedure, while

blue boxes perform the various calculations.

In **Block 1** (Strategy Definition), an alternative thermal request profile for a specific building is defined on the basis of weather forecasts for the following day. In **Block 2** (Historical Datastore - measurements) the measurements, sent by the IoT-enabled heat exchangers in the buildings through Message Broker, are retrieved from the Historical Datastore of the infrastructure. Subsequently, the algorithm computes the daily thermal consumption associated with the current request profile, as well as: i) the evaluation of the status of the heat exchanger, ii) the values of parameters characterizing the heat losses and the dynamic behavior of the building [32]. The latter quantities are calculated every day, taking advantage of almost steady state behavior reached in the afternoon and transient evolution when the heating system is switched off or attenuated. **Block 3** (Building compact model) assesses the feasibility of the proposed strategy for the building, as it compares the expected average internal temperatures resulting from the proposed rescheduling with the original scheduling. The acceptability of the internal temperature is imposed as a constraint for the new scheduling, which can be adjusted in the case it is not acceptable.

Block 4 (DIM (SIM + GIS) model of the network) represents an additional input for this tool. It retrieves from the DIM Services component of the infrastructure the SIM model of the network correlated with geo-referenced information (GIS) of the various sub-networks. This model defines the physical structure of sub-networks, and it is updated every year to properly consider possible expansions of the district heating to new areas or possible connection of additional buildings. **Block 5** (Network model) performs the thermo-fluid dynamic simulation of the district heating network in order to transform the thermal request of the users into thermal load for the full district and then for the thermal plants. The physical model of the network is fully described in [33].

Finally, **Block 6** (Plant model) performs calculation of the KPIs referred to energy and environmental impacts resulting from comparison of the current request profiles and the new profiles. Main KPIs are the reductions in primary energy consumption and CO₂ emission. In order to proceed with the simulation at building level, the thermal request simulation produces set points and heat scheduling for each heat exchanger. Such control policy is feasible thanks to the fine-grained monitoring provided by the proposed platform, and by the IoT devices deployed in the buildings. Detailed information about energy consumption of buildings served by the heating network are available. Indeed, thanks to this pervasive IoT platform, it is possible to quantify the energy request profile of each building and not only the cumulative heating request profile at the energy power plant.

To evaluate the impact of the new district heating control policy, simulations were also performed at building level. These simulations take as input the BIM for the building developed with Autodesk Revit 2015. Taking into account the building typology, the BIM is compiled with each building feature. By means of the DIM Services component of the infrastructure, the BIM model is then imported into Energy+ environment in order to calculate the internal temperature evolutions. The Building Simulation Engine is compiled setting the room occupancy and using heating set points and

TABLE I
ENERGY SCHEDULES FOR SELECTED BUILDING

	Start	Pause	Restart	Stop
Current	5:30 am	---	---	10:00 pm
Proposed	5:00 am	3:00 pm	3:30 pm	10:00 pm

TABLE II
DIFFERENCE BETWEEN SIMULATED AND MEASURED TEMPERATURE AND HUMIDITY (CURRENT SCHEDULE)

ΔT sim. - meas.		Room S	Room E	Room N
Max	[°C]	0.44	0.71	0.38
Min	[°C]	-0.09	-0.03	-0.25
Average	[°C]	0.37	0.23	0.32
St. Dev.	[°C]	0.98	0.46	0.69
$\Delta \phi$ sim. - meas.		Room S	Room E	Room N
Max	[%]	6.18	5.32	5.28
Min	[%]	-6.02	-5.01	-4.89
Average	[%]	1.28	0.76	0.52
St. Dev.	[%]	4.76	4.02	3.56

TABLE III
PMV CURRENT SCHEDULE VS PROPOSED SCHEDULE

		Room S	Room E	Room N
Current	Min	-0.06	-0.32	-0.80
	Max	-0.34	0.09	-0.30
	Avg	0.16	-0.09	-0.48
Proposed	Min	-0.09	-0.37	-0.82
	Max	0.31	0.04	-0.36
	Avg	0.13	-0.12	-0.52

schedules provided by the district thermal request simulation.

Subsequently, the Building Simulation Engine estimates thermal parameters such as the indoor temperature and thermal comfort for the building. Furthermore, for buildings under analysis, temperatures are then compared with the real measurements operated in selected rooms, where IoT sensors have been deployed. Every 15 minutes, these devices send monitoring information through the Message Broker to the Historical Datastore. In particular, this paper presents the results of simulations performed for the selected building on working days, when the primary school is occupied by users.

B. Experimental Results

In order to evaluate possible benefits that can be achieved though proper change of the thermal request profiles of the buildings, two scenarios corresponding with a typical winter day for the city under investigation are considered. Figure 4 reports the thermal loads of the plants: in the reference scenario (solid line), the current request profile of the buildings is considered. In a second scenario (dashed line), the request profiles of 50% of the buildings is modified by anticipating or postponing the time the heating system is switched on. The alteration is within a period smaller than 30 minutes, but the request profile curve of each building is unmodified.

Optimization of the strategy mainly involves an anticipation of the start-up time for a large percentage of the buildings.

This causes an increase in the total thermal load from about 4:30 a.m. to 6:00 a.m. Figure 4 shows that peak request is reduced of more than 120 MW and, more important, about 0.14 GWh of heat is moved from the periods exceeding the total capacity of the cogenerators to periods characterized by a total load smaller than the capacity of cogenerators. This means that this amount of heat can be produced using the cogenerators instead of auxiliary boilers. The primary energy factor of heat produced through boilers is about 1.11, which means that 1.11 units of fuel are requested to produce a unit of heat. In contrast, the primary energy factor of heat produced through the combined cycles operating in cogeneration mode is about 0.36. This small value is justified by the fact that the electrical efficiency of these combined cycles is about 0.58, when they operate in full electricity production. When they operate in cogeneration mode, for each 10 MW of heat produced, the electricity production decreases of about 2.1 MW [34]. A corresponding reduction of 3% in the primary energy consumption is achieved. This reduction depends in a non-linear way on the percentage of buildings which heating schedule can be modified: if the percentage of buildings is reduced to 25%, the primary energy saving is reduced to about 2%, while if the percentage of buildings is increased to 75%, the primary energy savings is increased to about 3.5%.

It is worth highlighting the fact that application of a strategy, aiming at optimizing the shape of the thermal request profiles of the buildings, would result in an even larger reduction in the primary energy consumption.

Once the new schedule is defined, the building model is applied to the evaluation of the acceptability of the changes in the thermal profiles in the various buildings. This means that when the thermal request profile of a building is modified, internal temperatures are checked so that the new comfort conditions are acceptable for the end-users. The current and the proposed (anticipated) schedules for the primary school are shown in TABLE I. At building level, the comfort has been estimated using the temperature trends and by exploiting the Predicted Mean Vote methodology (PMV) [35], which is a measure of thermal comfort of the building. PMV has been estimated for i) the current energy distribution schedule and for ii) the schedule proposed as output from the district model.

First, simulated and measured temperatures from IoT devices have been compared to assess whether the model was able to simulate correctly a real world building. Indeed, the use of IoT devices, which can be exploited through the infrastructure, allows an unprecedented fine-grained measurement of indoor conditions that can be used to validate simulation models.

Figure 5 and 6 explain how the infrastructure, and most importantly the Simulation Engine, can be used to assess the impact on comfort of the scheduling policy shown in Figure 4. We show here that a simulated evaluation is reliable, by comparing with real data taken from installed sensors. Also, we envision a potential usage of the infrastructure where sensor data can be used to tune the simulation model.




Figure 5. Simulated and actual temperature trends for selected rooms

Figure 5 depicts the actual and simulated (with the current schedule) temperature trends for the three selected rooms (S, E and N). The trends are estimated in four working days and during the weekend. A brief description of temperatures and relative humidity (ϕ) evolutions is also provided in TABLE II. In all cases, simulation tends to overestimate the maximum temperatures and slightly underestimate the minimum temperatures. South room (S) is constantly exposed to the solar radiation during the day, therefore it experiences the maximum temperatures. Deviations are larger in the central part of the day and are mainly due to unforeseeable factors, such as the such as ventilation, which is usually operated manually opening the windows, the variable window shading and the occupancy of rooms. These effects have a significant impact on the temperatures measured by the IoT devices but are not considered in simulations, as these should rely only on predictable data at a district level. In this case, the average difference between simulated and measured temperature is $0.37\text{ }^{\circ}\text{C}$ with a standard deviation of $0.98\text{ }^{\circ}\text{C}$. Differently, the east room (E) and north room (N) do not fully benefit of solar exposure during the day. In the case of room E the average temperature difference is $0.23\text{ }^{\circ}\text{C}$, while the standard deviation is $0.46\text{ }^{\circ}\text{C}$, which is the best model prediction. Room N is located in the north-west corner at the top floor of the primary school; the average temperature difference $0.32\text{ }^{\circ}\text{C}$ with a standard deviation of $0.69\text{ }^{\circ}\text{C}$. Relative humidity is slightly overestimated in simulations, but this does not significantly affect the evaluation of internal comfort conditions.

It is worth highlighting that the model is applied with the goal of checking the effects of variations in schedules using meteorological forecasts for the following day, therefore the aspect it is requested to capture is the dynamic behavior of the building. This can be analyzed by comparing the average temperature decreases when the heating system is switched off (i.e. between 8:30 p.m. and 4:30 a.m.). The measured values for the three rooms are $0.05\text{ }^{\circ}\text{C}/\text{hour}$, $0.08\text{ }^{\circ}\text{C}/\text{hour}$ and $0.10\text{ }^{\circ}\text{C}/\text{hour}$, respectively, while the simulated values are $0.06\text{ }^{\circ}\text{C}/\text{hour}$, $0.10\text{ }^{\circ}\text{C}/\text{hour}$ and $0.10\text{ }^{\circ}\text{C}/\text{hour}$. These differences are sufficiently small; therefore, the model is suitable for the




Figure 6. PMV methodology for proposed schedule for selected rooms

analysis.

Finally, to better assess the degree of comfort, the PMV methodology has been used to compare the current energy distribution schedules and the schedules proposed by the output of the district simulation. This evaluation has been performed using the simulation results. As shown in Figure 6, the expected PMV for the proposed schedule is only slightly below that corresponding with the current schedule. PMV varies from about 0.3 in the central hours of the day and -1.3 at night, being the last value representative of slightly cold comfort conditions. Values below -1 are only experienced in the north room and outside the working hours. A comparison between the current and proposed schedules is presented in TABLE III, where only the working hours (8 am - 5 pm) are considered. On average, the proposed schedule causes a slightly smaller PMV in the selected rooms, with a reduction of about 0.03 during daytime, due to the 20 min pause of the heating system which is introduced between 5:50 and 6:10. The internal temperatures and the PMV corresponding with the current schedule are recovered in the early afternoon for South and East rooms and in the late afternoon for the North room through slightly longer operation of the heating system in the afternoon.

In the presented case, the building simulation supports the adoption of a new energy distribution schedule by the energy provider, that promotes energy saving (at the district level) without affecting the user comfort (which has been estimated using the building model).

VI. CONCLUSION AND FUTURE WORK

This work presented a novel IoT platform for city district data management and energy flow simulations. In particular, this platform is instrumental i) to integrate heterogeneous IoT devices for monitoring and management of a whole city district; ii) to share building and energy network resources, both for visualization and simulation of energy policies, at building and district level; and iii) to assess the quality of the energy model of buildings.

The presented platform is designed and developed following the micro-service paradigm, and it is able to scale up reliably. Further, the use of shared open standards of the

Web makes it easy to integrate into existing systems. The platform is ready to send actuation commands to the devices deployed in the district. For instance, we are currently testing it on the DH scheduling control system. In this case, scheduled on-off heating switching times are sent to IoT devices to control heat exchangers. The results of this experimentation are preliminary and will be discussed in future work.

Finally, this platform satisfies the need for an energy management system for a city district, which is currently unmet in state-of-art applications. As shown in the real world case study, this platform is suitable to run energy distribution policies simulations from the district level down to the building room level, with a special focus on both district energy savings and end-user comfort level. It is worth noting that such policies are feasible thanks to the fine-grained monitoring provided by the proposed IoT platform.

In future works, the infrastructure will be extended and complemented with a Big Data Analytics module, which will be used to carry historical data analysis on the data collected from the infrastructure.

REFERENCES

- [1] World Health Organization, "Urban population growth," [Online]. Available: http://www.who.int/gho/urban_health/situation_trends/urban_population_growth_text/en/. [Accessed February 2016].
- [2] K. Ashton, "That 'Internet-of-Things' Thing," [Online]. Available: <http://www.rfidjournal.com/articles/view?4986>. [Accessed September 2015].
- [3] D. Giusto, A. Iera, G. Morabito and L. Atzori, The internet of things: 20th Tyrrhenian workshop on digital communications, Springer Science & Business Media, 2010.
- [4] C. Eastman, C. M. Eastman, P. Teicholz and R. Sacks, BIM handbook: A guide to building information modeling for owners, managers, designers, engineers and contractors, John Wiley & Sons, 2011.
- [5] R. F. Tomlinson, "A geographic information system for regional planning," in *In GA Stewart, (ed.): Symposium on Land Evaluation, Commonwealth Scientific and Industrial Research Organization*, Melbourne, 1968.
- [6] A. Kamilaris, A. Pitsillides and V. Trifa, "The smart home meets the web of things," *International Journal of Ad Hoc and Ubiquitous Computing*, vol. 7, no. 3, pp. 145-154, 2011.
- [7] 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 Syst. J.*, DOI: 10.1109/JSYST.2014.2302750.
- [8] D. Bonino, E. Castellina and F. Corno, "The DOG gateway: enabling ontology-based intelligent domotic environments," *IEEE Trans. Consum. Electron.*, vol. 54, no. 4, pp. 1656-1664, 2008.
- [9] L. Da Xu, W. He and S. Li, "Internet of Things in industries: A survey," *IEEE Trans. Ind. Informat.*, vol. 10, no. 4, 2014.
- [10] K. P. Athanasios, J. V. Gialelis, C. E. Alexakos, M. J. Georgoudakis and S. Koubias, "Vertical integration of enterprise industrial systems utilizing web services," *IEEE Trans. Ind. Informat.*, vol. 2, no. 2, pp. 120-128, 2006.
- [11] W. He and L. Da Xu, "Integration of distributed enterprise applications: a survey," *IEEE Trans. Ind. Informat.*, vol. 10, no. 1, pp. 35-42, 2014.
- [12] F. G. Brundu, E. Patti, M. Del Giudice, A. Osello, E. Macii and A. Acquaviva, "DIMCloud: a distributed framework for district energy simulation and management," in *Internet of Things. User-Centric IoT*, Springer International Publishing, 2015, pp. 331-338.
- [13] Y.-J. Kim, M. Thottan, V. Kolesnikov and W. Lee, "A secure decentralized data-centric information infrastructure for smart grid," *IEEE Commun. Mag.*, vol. 48, no. 11, pp. 58-65, 2010.
- [14] P. T. Eugster, P. A. Felber, R. Guerraoui and A.-M. Kermarrec, "The many faces of publish/subscribe," *ACM Computing Surveys (CSUR)*, vol. 35, no. 2, pp. 114-131, 2003.
- [15] E. Patti, A. L. A. Syri, M. Jahn, P. Mancarella, A. Acquaviva and E. Macii, "Distributed Software Infrastructure for General Purpose Services in Smart Grid," *IEEE Trans. Smart Grid*, DOI: 10.1109/TSG.2014.2375197.
- [16] D. T. Neves, M. Santos and M. Pinto, "ReActOR: A middleware as a service to interact with objects remotely," in *Proc. ICIT 2015*, 2015.
- [17] G. Cândido, A. W. Colombo, J. Barata and F. Jammes, "Service-oriented infrastructure to support the deployment of evolvable production systems," *IEEE Trans. Ind. Informat.*, vol. 7, no. 4, pp. 759-767, 2011.
- [18] "SensorGrid4Env," [Online]. Available: <http://www.sensorgrid4env.eu/>. [Accessed September 2015].
- [19] A. Nettstraeter, "Architectural reference model for IoT," *EC FP7 IoT-A (257521) D*, vol. 1, p. 2, 2012.
- [20] "OneM2M," [Online]. Available: <http://onem2m.org/>. [Accessed September 2015].
- [21] "FI-WARE," [Online]. Available: <https://www.fiware.org/>. [Accessed September 2015].
- [22] R. T. Fielding, *Architectural styles and the design of network-based software architectures*, Irvine: University of California, 2000.
- [23] "MQTT," [Online]. Available: <http://mqtt.org>. [Accessed November 2015].
- [24] "Linksmart," [Online]. Available: <http://linksmart.eu>. [Accessed December 2015].
- [25] E. Patti, A. Acquaviva and E. Macii, "Enable sensor networks interoperability in smart public spaces through a service oriented approach," in *IEEE International Workshop on Advances in Sensors and Interfaces (IWASI)*, 2013.
- [26] "SENML Draft," [Online]. Available: <https://tools.ietf.org/html/draft-jennings-senml-10>.
- [27] A. Jena. [Online]. Available: <https://jena.apache.org>. [Accessed January 2016].
- [28] A. Osello, *The Future of Drawing with BIM for Engineers and Architects*, Palermo: Dario Flaccovio Editore, 2012, pp. 1-323.
- [29] M. Fowler, "Microservices," [Online]. Available: <http://martinfowler.com/articles/microservices.html>. [Accessed September 2015].
- [30] "Autodesk Revit," [Online]. Available: <http://www.autodesk.com/products/revit-family/overview>. [Accessed November 2015].
- [31] V. Verda and F. Colella, "Primary energy savings through thermal storage in district heating networks," *Energy*, vol. 36, pp. 4278-4286, 2011.
- [32] V. Verda and G. Baccino, "Primary energy reductions in district heating networks through variation of the thermal load profile of the users," in *Proceedings of ECOS*, 2014.
- [33] A. Sciacovelli, V. Verda and R. Borchellini, *Numerical design of thermal systems*, Torino: Clut, 2013.
- [34] V. Verda, M. Chirio and A. Ponta, "A Thermoeconomic Based Tool for the Real-time Management of a Cogenerative Combined Cycle.," in *ASME 2006 International Mechanical Engineering Congress and Exposition*, 2006.
- [35] P. O. Fanger, *Thermal comfort. Analysis and applications in environmental engineering*, Copenhagen: Danish Technical Press, 1970.