

Cross-Domain Interoperability Using Federated Interoperable Semantic IoT/Cloud Testbeds and Applications: The FIESTA-IoT Approach

Martin Serrano, Amelie Gyrard, Michael Boniface, Paul Grace, Nikolaos Georgantas, Rachit Agarwal, Payam Barnaghi, Francois Carrez, Bruno Almeida, Tiago Teixeira, et al.

► To cite this version:

Martin Serrano, Amelie Gyrard, Michael Boniface, Paul Grace, Nikolaos Georgantas, et al.. Cross-Domain Interoperability Using Federated Interoperable Semantic IoT/Cloud Testbeds and Applications: The FIESTA-IoT Approach. Building the Future Internet through FIRE 2016 FIRE Book: A Research and Experiment based Approach, River Publishers, 2017, 978-87-93519-11-4. <http://www.riverpublishers.com/research_details.php?book_id=427>. <hal-01556289>

HAL Id: hal-01556289

<https://hal.inria.fr/hal-01556289>

Submitted on 4 Jul 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

10

Cross-Domain Interoperability Using Federated Interoperable Semantic IoT/Cloud Testbeds and Applications: The FIESTA-IoT Approach

Martin Serrano¹, Amelie Gyrard¹, Michael Boniface², Paul Grace²,
Nikolaos Georgantas³, Rachit Agarwal³, Payam Barnagui⁴, Francois
Carrez⁴, Bruno Almeida⁵, Tiago Teixeira⁵, Philippe Cousin⁶, Frank Le
Gall⁶, Martin Bauer⁷, Ernoe Kovacs⁷, Luis Munoz⁸, Luis Sanchez⁸,
John Soldatos⁹, Nikos Kefalakis⁹, Ignacio Abaitua
Fernández-Escárzaga¹⁰, Juan Echevarria Cuenca¹¹, Ronald Steinke¹²,
Manfred Hauswirth¹², Jaeho Kim¹³ and Jaeseok Yun¹³

¹Insight Centre for Data Analytics, Ireland

²IT Innovation, UK

³INRIA, France

⁴UNIS, United Kingdom

⁵Unparallel Innovation, Lda (UNPARALLEL, Portugal)

⁶Easy Global Market, France

⁷NEC Europe Ltd. UK

⁸University of Cantabria, Spain

⁹Athens Information Technology, Greece

¹⁰Sociedad para el desarrollo de Cantabria, Spain

¹¹Ayuntamiento de Santander, Spain

¹²Fraunhofer FOKUS, Germany

¹³Korea Electronics Technology Institute, Korea

10.1 Introduction

The Internet-of-Things (IoT) [61] has been identified as one of the main pillars of the world's economies and the technology enabler for the evolution of the societies and for the future developments and improvement of the Internet

[4]. A large number of research activities in Europe have been working in this direction i.e. FP7 projects in the context of Future Internet Research and Experimentation (FIRE) initiative. FIRE projects have already demonstrated the potential of IoT technologies and deployments in a number of different application areas including transport, energy, safety and healthcare. FIRE deployments and project results have also demonstrated the advantages of implementing Smart Cities testbeds (national and EU scale) both have been extensively reported in [5]. Smart City testbeds are the key places for large demonstration of IoT concepts and technology. Smart cities testbeds are prone to be large scale, highly heterogeneous and target a diverse set of application domains.

In Smart cities despite the growing number of IoT deployments, multiple installations and related testbeds, the majority of deployed IoT applications tend to be self-contained, thereby forming application silos [50]. Recent research efforts have been focused on demonstrate the capacity of IoT systems to be part of an overall arch-systems called federation (e.g., FP7 Fed4FIRE), in a federated environment it is possible the co-existence and co-operation of multiple infrastructures (including IoT testbeds). The Federation is the first step to the integration of these silos, since they provide a wide range of indispensable low-level capabilities such as resource reservation and negotiation. Nevertheless, these efforts tend to be heavyweight and do not adequately deal with the need to access diverse IoT datasets in a flexible and seamless way. In a federation one of the mayor challenges is the data centric integration and the combination of data silos that is identified as a under investigation area for IoT [4], and with a very rich potential both in terms of novel experimentation (e.g., in the scope of living labs and IoT testbeds) [49] and in terms of added-value enterprise applications. Related to data, the ability to combine and synthesize data streams and services from diverse IoT platforms and testbeds remains a challenge and multiple researches follows the promise to broaden the scope of potential data interoperability applications in size, scope and targeted business context. In the Internet of things area the ability to repurpose and reuse IoT data streams across multiple experimental applications can positively impact the Return-on-Investment (ROI) associated with the usually costly investments in IoT infrastructures and testbeds. The integration, combination and interoperability of IoT silos is fully in-line with the overall FIRE vision that makes part of the Horizon 2020 program, which aspires to allow European experimenters/researchers to investigate/develop leading-edge, ubiquitous and reliable computing services, as well as seamless and open access to global data resources.

The futuristic vision of integrating IoT platforms, testbeds and associated silo applications is related with several scientific challenges, such as the need to aggregate and ensure the interoperability of data streams stemming from different IoT platforms or testbeds, as well as the need to provide tools and techniques for building applications that horizontally integrate silo platforms and applications. The convergence of IoT with cloud computing is a key enabler for this integration and interoperability, since it allows the aggregation of multiple IoT data streams towards the development and deployment of scalable, elastic and reliable applications that are delivered on-demand according to a pay-as-you-go model. During the last 4–5 years we have witnessed several efforts towards IoT/cloud integration (e.g., [29, 39]), including open source implementations of middleware frameworks for IoT/cloud integration [23, 52] and a wide range of commercial systems (e.g., Xively (xively.com), ThingsWorx (thingsworx.com), ThingsSpeak (thingspeak.com), Sensor-Cloud (www.sensor-cloud.com)). While these cloud infrastructures provide means for aggregating data streams and services from multiple IoT platforms, they are not fully sufficient for alleviating IoT fragmentation of facilities and testbeds. This is because they emphasize on the syntactic interoperability (i.e. homogenizing data sources and formats) rather on the semantic interoperability of diverse IoT platforms, services and data streams.

Recently several IoT projects [33] have started to work on the semantic interoperability of diverse IoT platforms, services and data streams. To this end, they leverage IoT semantic models (such as the W3C Semantic Sensor Networks (SSN) ontology [16, 58]) as a means of achieving interoperable modeling and semantics of the various IoT platforms. A prominent example is the FP7 OpenIoT project, a (BlackDuck) award winner open source project in 2013, which has been developed and released as an open source blueprint infrastructure [51] addressing the need for semantic interoperability of diverse sensor networks at a large scale (see also <https://github.com/OpenIoTOrg/openiot>). The semantic interoperability of diverse sensor clusters and IoT networks is based on the virtualization of sensors in the cloud. At the heart of these virtualization mechanisms is the modeling of heterogeneous sensors and sensor networks according to a common ontology, which serves as harmonization mechanism of their semantics, but also as a mechanism for linking related data streams as part of the linked sensor data vision. This virtualization can accordingly enable the dynamic discovery of resources and their data across different/diverse IoT platforms, thereby enabling the dynamic on-demand formulation of cloud-based IoT services (such as Sensing-as-a-Service services). Relevant semantic interoperability

techniques are studied in depth as part of the fourth activity chain of the IERC cluster (IERC-AC4) (see for example [17]). Similar techniques could serve as a basis for unifying and integrating/linking geographically and administratively dispersed IoT testbeds, including those that have been established as part of FIRE projects. Such integration holds the promise of adding significant value to all of the existing IoT testbeds, through enabling the specification and conduction of large-scale on-demand experiments that involve multiple heterogeneous sensors, Internet Connected Objects (ICOs) and data sources stemming from different IoT testbeds.

Based on the above-mentioned Sensing-as-a-Service paradigm, dynamic virtualized discovery capabilities for IoT resources could give rise to a more general class of Experiment-as-a-Service (EaaS) applications for the IoT domain. EaaS services are executed over converged IoT/cloud platforms, that are developed on the basis of the technologies outlined above. EaaS services are not confined to combinations of sensor queries (such as Sensing-as-a-Service), but they would rather enable the execution of fully-fledged experimental workflows comprising actuating and configuration actions over the diverse IoT devices and testbeds. The benefits resulting from the establishment and implementation of an EaaS paradigm for the IoT domain include:

- The expansion of the scope of the potential applications/experiments that are designed and executed. Specifically, the integration of diverse testbeds for offering to the European experimenters/researchers with the possibility of executing IoT experiments that are nowadays not possible.
- The ability to repurpose IoT infrastructures, devices and data streams in order to support multiple (rather than a single) applications. This increases the ROI associated with the investment in the testbeds infrastructure and software.
- Possibility for sharing IoT data (stemming from one or more heterogeneous IoT testbeds) across multiple researchers. This can be a valuable asset for setting up and conducting added-value IoT experiments, since it enables researchers to access data in a testbed agnostic way i.e. similar to accessing a conventional large scale IoT database.
- The emergence of opportunities for innovative IoT applications, notably large scale applications that transcend multiple application platforms and domains and which are not nowadays possible.
- The avoidance of vendor lock-in, when it comes to executing IoT services over a provider's infrastructure, given that an EaaS model could boost data and applications portability across diverse testbeds.

Beyond the interconnection and interoperability of IoT and smart cities testbeds, semantic interoperability tools and techniques could also enable the wider interoperability of IoT platforms, which is a significant step towards a global IoT ecosystem.

10.2 Federated IoT Testbeds and Deployment of Experimental Facilities

Addressing the need of IoT federated infrastructures and following the interoperability need and the use of semantics IoT/cloud Testbeds and applications the FIESTA project aim to be a globally unique infrastructure for integrated IoT experimentation based on the federation of multiple interoperable IoT testbeds. FIESTA targets the main objective for defining and implementing a Blueprint IoT Experimental Infrastructure that can offer services and tools for external applications and mainly for enabling the concept Experimentation as a Service “EaaS”. FIESTA look at researching and establishing a novel blueprint infrastructure for IoT platforms/testbeds interoperability and EaaS (Experimentation-as-a-Service), which enables researchers, engineers and enterprises (including SMEs) to design and implement integrated IoT experiments/applications across diverse IoT platforms and testbeds, through a single entry point and based on a single set of credentials. The EaaS infrastructure facilitates experimenters/researchers to conduct large scale experiments that leverage data, information and services from multiple heterogeneous IoT testbeds, thereby enabling a whole new range of innovative applications and experiments.

FIESTA has implemented the testbed agnostic access to IoT datasets, providing tools and techniques enabling researchers to share and access IoT-related datasets in a seamless testbed agnostic manner i.e. similar to accessing a large scale distributed database. This also has involved the use of linking diverse IoT datasets, based mainly on the linked sensor data concept. FIESTA has implemented tools and techniques for IoT Testbeds Interoperability and Portability by providing tools and techniques (semantic models, directory services, open middleware, tools) for virtualizing and federating geographically and administratively dispersed IoT platforms and testbeds. Special emphasis was done in the specification and implementation of common standardized APIs for accessing the underlying testbeds, thereby boosting the portability of IoT experiments. FIESTA has also research and implement the meta-cloud infrastructure along with accompanying tools (i.e. portal, development,

workflow management, monitoring) facilitating the use of the EaaS infrastructure for the design, implementation, submission, monitoring and evaluation of IoT/cloud related experiments and related integrated applications.

FIESTA developed a global market confidence programme (as a Sustainability Vehicle) for enabling IoT platform/testbed providers and IoT solutions providers to test, validate and ensure the interoperability of their platforms/solutions against FIESTA standards and techniques. The programme includes a certification suite for compliance testing. As part of pursuing this objective, FIESTA ensures the development and realization of a clear sustainability path for the project's results. Furthermore, it defined ways for collaboration with other bodies and working groups, which are currently working (at EU level) towards the establishment of similar initiatives, such as the IoT forum. FIESTA is implemented in the way to be a blueprint experimental infrastructure for EaaS on the basis of the federation and virtualization of real-life IoT testbeds, but also on the basis of real-life experiments that have been designed, executed and evaluated over them. These span the areas of pollution monitoring, crisis management, crowdsensing as well as enterprise/commercial activities and emphasize portability and testbed agnostic access.

FIESTA implemented a stakeholders engagement program to guarantee the expansion in terms of experiments and testbeds by meaning of the involvement of third parties towards a global IoT experimentation ecosystem). The FIESTA ecosystem is to attract and engage stakeholders beyond the project consortium as third parties through managing an open calls process, but also through the mobilization of (third-party) research communities with a strong interest in IoT. FIESTA permanently works towards the identification and generation of reference activities to elicit and document a range of best practices facilitating IoT platform providers and testbed owners/administrators to integrate their platform/testbed within FIESTA, along with best practices addressed to researchers, engineers and organizations wishing to use the FIESTA meta-cloud EaaS infrastructure for conducting innovative applications and experimentation.

In order to validate the global and federated character of the FIESTA infrastructure, FIESTA has already established collaborations and liaisons with IoT partners in Asia (Korea) and USA. In particular, the consortium includes a Korean partner (KETI), that has also established IoT collaborations with US organizations (thanks to the Inria's collaboration with the Silicon Valley as part of the Inria@Silicon Valley programme). Note that KETI's participation in the consortium has allowed the integration/federation of a testbed located in

Asia (i.e. KETI's testbed) to the FIESTA EaaS infrastructure. At the same time, the above-listed collaborations ensures the global dissemination and outreach of the project's results, while also broadening the scope of participation in the third-party selection processes of the project (i.e. open calls) on the basis of participants from Asia and USA.

FIESTA has allocated a significant share (31%) of its foreseen budget to the introduction of third-parties (through the open calls process), notably third-parties that have started the undertaken and the conduction of new experiments and/or the blending/integration of new testbeds within the FIESTA infrastructure. Note that the stakeholders' community of the project also serves as a basis for validating the global market confidence programme of the project. The active engagement of the stakeholders in the project, but also in the third-parties selection process are boosted by FIESTA partners already animating ecosystems of researchers and enterprises (i.e. SODERCAN, Com4innov), as well as from participants from non-EU countries (i.e. KETI from Korea). Links to participants from Asia and USA are also sought (through KETI and the Inria@Silicon Valley programme). The ultimate vision of FIESTA is to provide the basis of a global IoT experimentation ecosystem.

10.3 Cross-Domain Interoperability

FIESTA project has opened new horizons in the development and deployment of IoT applications and experiments not only at a EU but also global scale, based on the interconnection and interoperability of diverse IoT platforms and testbeds FIESTA has created an ecosystem of IoT experimentation. To this end, FIESTA provides a blueprint experimental infrastructure, tools, techniques, processes and best practices enabling IoT testbed/platforms operators to interconnect their facilities in an interoperable way, while at the same time facilitating researchers and solution providers in designing and deploying large scale integrated applications (experiments) that transcend the (silo) boundaries of individual IoT platforms or testbeds. FIESTA enables researchers and experimenters to share and reuse data from diverse IoT testbeds in a seamless and flexible way that has open up new opportunities in the development and deployment of experiments and for exploiting data and capabilities from multiple testbeds. The blueprint experimental infrastructure provided by FIESTA includes a middleware for semantic interoperability, tools for developing/deploying and managing interoperable applications, processes for ensuring the operation of interoperable applications, as well as best

practices for adapting existing IoT facilities to the FIESTA interoperability infrastructure.

The FIESTA infrastructure empowers the Experimentation-as-a-Service (EaaS) paradigm for IoT experiments, while also enables experimenters to use a single EaaS API (i.e. the FIESTA EaaS API) for executing experiments over multiple IoT federated testbeds in a testbed agnostic way i.e. like accessing a single large scale virtualized testbed. Experimenters are therefore able to learn easily how to connect with the EaaS API and accordingly use it to access data and resources from any of the underlying testbeds. To this end, the underlying interconnected testbed provides common standardized semantics and interfaces (i.e. FIESTA Testbed Interfaces) enables the FIESTA EaaS infrastructure to access their data, resources and other low-level capabilities (Figure 10.1). Note that the FIESTA EaaS infrastructure is accessible through a cloud computing infrastructure (conveniently called FIESTA meta-cloud), on the basis of a cloud-based on-demand paradigm.

FIESTA also includes a directory service (conveniently called FIESTA meta-directory), where sensors and IoT resources from multiple testbeds are registered. This directory enables the dynamic discovery and use of IoT resources (e.g., sensors, services) from all the interconnected testbeds. Overall, the project's experimental infrastructure provides to the European




Figure 10.1 FIESTA interoperability model for heterogeneous IoT testbed experimentation.

experimenters in the IoT domain with the following unique capabilities (Figure 10.1):

- Access to and sharing of IoT datasets in a testbed-agnostic way. FIESTA provides researchers with tools for accessing IoT data resources (including Linked sensor data sets) independently of their source IoT platform/testbed.
- Execution of experiments across multiple IoT testbeds, based on a single API for submitting the experiment and a single set of credentials for the researcher.
- Portability of IoT experiments across different testbeds, through the provision of interoperable standards-based IoT/cloud interfaces over diverse IoT experimental facilities.

FIESTA technology leverages recent results on IoT semantic interoperability, notably results produced as part of the AC4 activity chain of the IERC, as well as within relevant projects in the IoT (e.g., FP7 OpenIoT) and FIRE (e.g., Fed4FIRE) areas. In particular, IoT projects offers the foundations of semantic interoperability at the IoT data and resources levels, while FIRE projects contribute readily available results in the area of reserving and managing resources across multiple testbeds. On the basis of these results, FIESTA research, design and deliver an open middleware infrastructure (i.e. semantics and APIs) for the virtualization and federation of IoT testbeds that enable sharing and access to a wide range of IoT-related datasets. FIESTA's infrastructure comprise semantic models enabling the virtualization, as well as middleware libraries facilitating the streaming and semantic annotation of IoT from the various testbeds in a single unified cloud infrastructure (FIESTA cloud). The FIESTA cloud therefore aggregates, manages and linked data from the various testbeds, while at the same time providing methods and tools that enables researchers to access them in a flexible and testbed-agnostic way. Therefore, the FIESTA cloud act as a meta-testbed, which integrates, linked and uses information sources from a variety of IoT/cloud testbeds.

FIESTA cloud enables European experimenters/researchers to design, implement, execute and evaluate IoT experiments based on data from various IoT testbeds all over Europe. To this end, FIESTA also offers a wide range of tools facilitating experimenters in the above tasks. These include: a) A portal infrastructure serving as a single entry point for setting up and submission of IoT experiments and the monitoring of their progress, b) Tools for designing and enacting experiments in terms of IoT/cloud services and workflows, c) Tools for sharing, linking and accessing datasets in a testbed agnostic way,

d) Tools and techniques for monitoring and managing the FIESTA cloud, including monitoring of all the necessary aspects of the underlying testbeds and e) Tools and techniques for monitoring the status of experiments and collecting data for evaluating the experiments. These tools are an integral element of the project's Experiment-as-a-Service paradigm for the IoT domain.

FIESTA establishes, implement and support a global market confidence programme, on the basis of its blueprint infrastructures and processes, that encourages and facilitate stakeholders to comply with the FIESTA interoperability guidelines and accordingly to deploy large scale innovative interoperable IoT applications. The FIESTA global market confidence programme includes a certification/compliance suite enabling platform providers and solution providers to test and ensure the level of interoperability of their platforms and services. This programme is a main vehicle for the sustainability of the project's results, as well as for impact creation at a global scale. Note that the programme is used as a vehicle for the sustainability of the project's results. During its lifetime FIESTA boost and ensure the engagement and participation of multiple platforms providers within Europe (including both consortium members and third-parties) in the FIESTA global confidence programme. Based on this engagement, FIESTA ensures the proper design, implementation, validation and fine-tuning of the programme.

FIESTA integrates diverse IoT testbeds (three in EU and one in Korea), towards providing experimenters with the possibility of designing, implementing, executing and evaluating sophisticated IoT (EaaS based) experiments that are not possible nowadays. To this end, the project leverages recent advances and results associated with semantic interoperability for IoT applications towards federating multiple IoT testbeds. FIESTA specify the scope of the IoT platforms and testbeds integration, federation and interoperability in terms of the functionalities that should be supported, the business/research actors that have access to specific functionalities of the testbeds, their EaaS model, as well as type of experiments that are enabled. FIESTA attempts to cover all aspects of IoT testbeds integration, including technology aspects (i.e. the technologies needed), business aspects (including how to run the confidence programme and ensure the longer term sustainability of the FIESTA model), organization (e.g., the processes needed to deploy/operate interoperable platforms and applications), as well as innovation aspects.

FIESTA has been validated on the basis of the federation of four existing real-life diverse IoT testbeds (provided by partners UNICAN/SDR, UNIS, Com4Innov and KETI), which include prominent European FIRE testbeds (such as SmartSantander), as well as Korean testbed (accessible through partner KETI). FIESTA first federate these testbeds and accordingly with

testbed specifications validate the federated/virtualized infrastructure on the basis of a range of EaaS experiments covering both e-science and e-business purposes. The project's experiments (which are detailed in following paragraphs) unveil the unique capabilities of the FIESTA infrastructure in terms of testbed-agnostic data sharing, execution of experiments across multiple testbeds, as well as ensuring the portability of IoT experiments across different testbeds.

In order to accomplish its goals, the project issue, manage and exploit a range of open calls towards involving third-parties in the project. The objective of the involvement of third-parties is two-fold:

- To ensure the design and integration (within FIESTA) of more innovative experiments, through the involvement of additional partners in the project (including SMEs). The additional experiments focuses on demonstrating the added-value functionalities of the FIESTA experimental infrastructure.
- To expand the FIESTA experimental infrastructure on the basis of additional testbeds. In this case the new partners undertake to contribute additional testbeds and to demonstrate their blending and interoperability with other testbeds (already adapted to FIESTA). As part of this blending, the owners of these testbeds also engage with the project's global market confidence programme, which provide them with the means to auditing the interoperability and openness of their platforms.

The involvement of third-parties therefore play an instrumental role for the large scale validation of the FIESTA experimental infrastructure, but also for the take-up of the project's global market confidence programme on IoT interoperability. It is also a critical step to the gradual evaluation of FIESTA towards an infrastructure/ecosystem for global IoT experimentation, as shown in Figure 10.2.

Beyond the validation of the FIESTA infrastructure on the basis of practical experiments and the integration of additional IoT testbeds, the project specify concrete best practices for the federation of testbeds (addressed to testbed owners/administrators) wishing to become part of the virtualized meta-cloud infrastructure of the project. Similar best practices are also produced for European researchers and enterprises (including SMEs) wishing to design and execute experiments over the FIESTA EaaS infrastructure. These best practices have been disseminated as widely as possible, as part of the project's efforts to achieve EU-wide/global outreach. The attraction and engagement of researchers and enterprises in the use of the FIESTA EaaS infrastructure is another vehicle for the sustainability and wider use of the project's results,




Figure 10.2 FIESTA evolution towards an ecosystem for global IoT experimentation.

which complement the global market confidence programme outlined above. This is overall in-line with the vision of establishing a global ecosystem for IoT experimentation (as already shown in Figure 10.2)

10.4 Experimentation as a Service

The FIESTA overall approach comprises a range of research activities that aims at setting up and validating the FIESTA EaaS model and associated blueprint experimental infrastructure, as well as a range of exploitation and sustainability activities that deals with the design and activation of the project’s global market confidence project on IoT interoperability. A set of demonstration activities have been carried out in order to showcase the capabilities of the FIESTA infrastructure on the basis of the design and execution of novel experiments.

The FIESTA project’s methodology towards researching and providing the FIESTA Experimentation as a Service (EaaS) paradigm, involves the following groups of activities, and the details are further analysed in following paragraphs:

- Analysing requirements for EaaS experimentation in the IoT domain, and specifying the detailed technical architecture of the FIESTA experimental (meta-cloud) infrastructure, including its (meta) directory of IoT resources.
- Research towards virtualizing access to the individual testbeds and their resources. This includes the provision of common standards-based interfaces and APIs (i.e. FIESTA Testbed APIs) for accessing datasets and resources in the various testbeds, according to common semantic models (ontologies).
- Research towards creating the FIESTA meta-cloud EaaS infrastructure, which enables experimenters to access data and resources from any of the underlying testbeds in a testbed agnostic way i.e. similar to accessing a single large scale virtualized testbed.

FIESTA Engineering Requirements: The FIESTA engineering requirements activities have produced the requirements associated with testbed-agnostic experimentation, as well as with the EaaS model to designing and conducting IoT experiments. They were planned early in the project's work plan and have produced the interoperability requirements and more, based on a variety of modalities for collecting and analysing requirements, including analysis of state-of-the-art, contact with stakeholders (including researchers and experimenters), analysis of the various IoT testbeds etc.

FIESTA Architecture and Technical Specifications: The FIESTA requirements have been taken into account towards producing detailed technical specifications for the EaaS model. Furthermore, a technical architecture have been established, specifying the FIESTA (meta-cloud) EaaS infrastructure, its tools, the meta-directory of IoT resources, as well as the interfaces of the above-listed components to individual FIESTA platforms and testbeds. The architecture drives the organization and integration of research tasks associated with the individual components of the FIESTA solution.

FIESTA Research on semantic interoperability for IoT (data and resources): The project's methodology includes a dedicated set of activities that aim at realizing IoT platforms/testbeds semantic interoperability at both the data and resources levels. To this end, FIESTA selects and extends the ontologies that provide the common semantics of the FIESTA interoperable infrastructure, while also working on the federation and linking of the heterogeneous data streams. As a result of the research, a set of blueprint middleware libraries enabling each testbed to adapt its data and resources to the common produced semantics.




Figure 10.3 FIESTA EaaS experimental infrastructure overview.

FIESTA Research on virtualized access to IoT/cloud infrastructures: In addition to developing the models that ensures the common semantics of resources and data across various testbeds, FIESTA project have provided a set of standards-based portable interfaces for accessing the various IoT/cloud infrastructures. The interfaces ensure that the FIESTA infrastructure can be seamlessly expanded on the basis of additional platforms/testbeds that support the specified standards-based interfaces.

FIESTA is in-line with the directions identified and prioritized as part of recent FIRE roadmaps in the areas of IoT and its convergence with cloud computing and smart city applications. FIESTA project addresses the challenges identified in recent support actions (e.g., the AmpliFIRE Support Action) for the FIRE domain. Figure 10.3 illustrates the main elements of the FIESTA EaaS infrastructure, which are further analysed in later paragraphs.

10.5 IoT Data Marketplace

FIESTA tools and techniques for accessing data in a testbed agnostic way defines a number of tools enabling submission of experiments, testbed agnostic access to (shared) data, as well as authentication and authorization of the

users are implemented and make available over the FIESTA meta-cloud infrastructure. FIESTA meta-cloud infrastructure has provided a meta-cloud infrastructure enabling access to data and resources from a wide range of underlying testbeds. This infrastructure leverages the semantics and interfaces that make FIESTA meta-cloud to serve as single entry point of the EaaS infrastructure. It also includes a (meta) directory service, which enables dynamic discovery and dynamic access to resources from any of the underlying virtualized testbeds.

The project's demonstration activities are focused on validating and demonstrating the FIESTA IoT Data Marketplace on the basis of three experiments that are designed and executed by project partners, but also based on several experiments that are executed by third-parties to be selected based on open calls processes. FIESTA demonstration of IoT Data Marketplace in a way of innovative experiments on Testbed agnostic data access and by sharing that data as a means of validating the FIESTA infrastructure is generated by using a number of innovative experiments over the FIESTA infrastructure that is being developed and demonstrated by the end of the project duration.

The focus on the IoT Data Marketplace is in three fold: a) Access to data and services from multiple IoT testbeds, b) Experiments portability across testbeds (i.e. provided that testbeds provide the sensors and/or resources needed to execute the experiment and c) Dynamic discovery of sensors and resources across multiple testbeds. A great deal of the demonstration activities is also based on new experiments to be introduced as part of the Open Calls processes of the project.

10.6 FIESTA Platform Services and Tools

FIESTA intends to become a first of a kind experimental infrastructure, which provides researchers with the capabilities of accessing data and services from multiple IoT testbeds in a seamless and testbed agnostic way. This enables researchers to design and enact more sophisticated and more innovative experiments, as part of their projects and product development processes. The realization of the FIESTA vision requires significant scientific and technological advancements in the areas of semantic interoperability of IoT testbeds, the linking of related IoT data streams, the development of IoT architectures suitable for federating multiple (cloud-based) testbeds, the provision of standards-based interfaces for accessing the various IoT/cloud testbeds, as well as the development of an on-demand EaaS model to executing experiments. The scientific and technological objectives of the project are

ground breaking since this allow researchers to experiment with data sets that stem for administratively and geographically dispersed testbeds, while at the same time ensuring the portability of the experiments across testbeds with similar/analogous capabilities. These advancements represent the scientific and technological ambition of the project. At the same time, the project has ambitious objectives associated with the sustainability and market take-up of the project's results, based on the establishment of the global market confidence programme for IoT interoperability. These ambitious targets are presented in the following paragraphs.

10.6.1 FIESTA Approach on Global Market Confidence Programme on Interoperability Service

A global market confidence programme on IoT interoperability has been designed as a vehicle for the sustainability of the project's results, but also as a means of offering these results in a structured way to many experimenters (i.e. individuals researchers and enterprises (including SMEs)). FIESTA operate the global market confidence programme on IoT interoperability, towards boosting the sustainability of the project's results, as well as towards using semantic interoperability as a vehicle for alleviating vendor lock-in and the related fragmentation of the IoT market.

The programme is designed to be validated on the basis of the auditing and certification of several IoT platforms for their interoperability against FIESTA standards and guidelines. IoT platforms/testbeds are contributed by project partners (based also on their background projects), but also by new participants joining the project following open calls processes.

The methodology of the project includes activities that aim at attracting stakeholders in the adoption and use of the project's results, based on the global market confidence programme of stakeholders. FIESTA caters for the support of these stakeholders, through providing focused training and consulting, relevant to the project's interoperability programme.

In addition to opportunities derived from the global market confidence programme on IoT interoperability, the FIESTA project conducts a wide range of dissemination and communication activities aiming at supporting the exploitation strategy and goals of the project. Likewise, all partners that are involved in exploiting the project in line with their business and research strategies, also, a set of created business plans in relation to the FIESTA exploitable products and services.

10.6.2 FIESTA Approach on Linking and Reasoning over IoT Data Streams Services

FIESTA's work on semantic interoperability of data streams is to ensure the accessibility of heterogeneous input streams in a uniform format, as well as the ability to support/implement a uniform access paradigm to these data. In addition to alleviating the complexity of the data access process, this interoperability also empowers large-scale reasoning over the multiple diverse data streams, towards linking related data streams and enabling large scale experiments, as well as experiments with richer functionality.

The most promising approach towards linking data streams is the use of Linked Open Data (LOD) standards [30] along with semantic annotations and uniform access with RESTful services (REST: REpresentational State Transfer) down to the physical sensor level. Linked Data ensures a uniform data model based on an underlying graph-based/network model (vs. a traditional relational model) capable of representing arbitrary information models in an intuitive and straightforward way. Linked Data models are used already in many domains, such as the Web, enterprise information systems, e-government (e.g., <http://data.gov.uk>), social networks (e.g., W3C Semantic Interoperability of Online Communities (SIOC) standard), sensors data (W3C Semantic Sensor Networks Incubator Group (SSN-XG), <http://www.w3.org/2005/Incubator/ssn/>), etc. with a trajectory of massive further growth. Uniform access in a RESTful way using Linked Data originated from Web-based information systems and has become the standard on Web-based systems and for accessing social media, e.g., Twitter REST API (<https://dev.twitter.com/docs/api>), as well as for many enterprise service solutions. Recently, also the IoT world has committed to RESTful access through the on-going standardization of the COntstrained Application Protocol (COAP, <https://datatracker.ietf.org/doc/draft-ietf-core-coap/>) and COntstrained RESTful Environments (CORE, <http://datatracker.ietf.org/wg/core/charter/>) by the IETF. A complete stack for Linked Data based on these abstractions has been developed by the FP7 project SPITFIRE (Semantic Service Provisioning for the Internet of Things, <http://www.spitfire-project.org/>).

Dynamic cost models and support for scalable and efficient processing are missing [60] as are query approximation and relaxation techniques for “close matches” [32]. Stream query processors for Linked Streams can already provide reasoning support up to the level of expressivity of SPARQL (<http://www.w3.org/TR/rdf-sparql-query/>). The most relevant systems are CQELS, C-SPARQL [10], and EP-SPARQL [3] among a number of research

prototypes (e.g., Sparkwave, which, however, does not have comprehensive performance evaluation results available, thus not making it comparable to the above 3 systems). These systems all share the same approach of utilizing SPARQL-like specification of continuously processed queries for streaming RDF data. If more complex reasoning is required, other approaches like nonmonotonic logic programming are required. Stream processing engines which augment stream reasoning through this kind of approach are still limited, but include those such as the use of Prova [36, 59] and Streaming IRIS [37]. Although based on logic programming, these approaches do not gain the inherent benefits of Answer Set Programming (ASP) syntax and semantics in terms of expressivity. In terms of performance, Prova is more concerned about how much background (static) knowledge can be pushed into the system, while Streaming IRIS does not test complex reasoning tasks. To the best of the consortium's knowledge, the work by Do [21] is probably the only other current stream reasoning approach for the Semantic Web that utilizes ASP. Although the work is quite recent, their approach is still much more prototypical. More importantly, this approach does not pertain to continuous and window-based reasoning over stream data.

10.6.3 FIESTA Approach on Federating IoT Stream Data Management Services

As we are heading towards a world of billions of things [26], IoT devices are expected to generate enormous amount of (dynamically distributed) data streams, which can no longer be processed in real-time by the traditional centralized solutions. Thus, IoT needs a distributed data management infrastructure to deal with heterogeneous data stream sources which autonomously generates data at high rates [9]. An early system designed to envision a world wide sensor web [11] is IrisNet, which supports distributed XML processing over a worldwide collection of multimedia sensor nodes, and addresses a number of fault-tolerance and resource-sharing issues. Along the same line, HiFi [24] also supports integrated push-based and pull-based queries over a hierarchy where the leaves are the sensor feeds and the internal nodes are arbitrary fusion, aggregator, or cleaning operators. A series of complementary database approaches aimed to provide low-latency continuous processing of data streams on a distributed infrastructure. The Aurora/Medusa [13], Borealis [1], and TelegraphCQ [12], StreamGlobe [53], StreamCloud [27] are

well-known examples of this kind. These engines provide sophisticated fault-tolerance, load-management, revision-processing, and federated-operation features for distributed data streams. A significant portion of the stream processing research merit of these systems has already made its way from university prototypes into industry products such as TIBCO StreamBase, IBM Stream InfoSphere, Microsoft Streamlight. However, such commercial products are out of reach of most IoT stream applications and there have not been any comprehensive evaluation in terms of cost effectiveness, performance and scalability. Due to this reason, there have emerged open source stream processing platforms from Apache Storm [54], S4 [55] and Spark [56] which were primarily built for some ad-hoc applications: Twitter, Yahoo!. While these platforms aim to support elasticity and fault-tolerance, they only offer simple generic stream processing primitives that require significant effort to build scalable stream-based applications.

The above systems provide steps in the right direction for managing IoT data streams in distributed settings. However, they have several federation restrictions in terms of systems of systems and system data organization. For system organization, most of distributed stream processing engines are extended from a centralized stream-processing engine to distributed system architectures. Thus, in order to enable the federation among stream processing sites, they have to follow strictly predefined configurations. However, in IoT settings, heterogeneous data stream sources are provided by autonomous infrastructures operated on different independent entities, which usually do not have any prior knowledge about federation requirements. In particular, a useful continuous federated query might need to compare or combine data from many heterogeneous data stream sources maintained by independent entities. For example, a tourist guide application might need to combine different data stream relevant to the GPS location of users, e.g., weather, bus, train location, flight updates, tourist events. Also, they might then correlate these streams with similar information from other users who have social relationships with the user via social networks such as Twitter, Facebook and also with background information like OpenStreetMap, Wikipedia. In such examples, stream data providers did not only agree how their systems are used to process those federated queries but also they did not agree on data schema/format to make data able to be queried for the federated query processing engine. Note however that the need of having uniform and predefined data schema and formats poses various difficulties for query federation on IoT applications using heterogeneous stream data sources.

In FIESTA the lack of standards has been studied as the major difficulty leading to restrictions, and the wide (and changing) variety of application requirements. Existing IoT Stream processing engines vary widely in data and query models, APIs, functionality, and optimization capabilities. This has led to some federated queries that can be executed on several IoT stream providers based on their application needs. Semantic Web addresses many of the technical challenges of enabling interoperability among data from different sources. Likewise, Linked Stream Data enables information exchange among stream processing entities, i.e., stream providers, stream-processing engines, stream consumer with computer-processable meaning (semantics) of IoT stream data. There have been a lot of efforts towards building stand-alone stream processing engine for Linked Stream Data such as C-SPARQL [10], SPARQLstream [10], EP-SPARQL, [6]. The data and query-processing model of Linked Stream data has been standard by W3C [46]. However, there are only few on-going efforts of building scalable Linked Stream Data processing engines for the cloud like Storm and S4 respectively, i.e., CQELS Cloud [31]. None of them supports federation among different/autonomous stream data providers.

10.6.4 FIESTA Approach on Semantic Interoperability for IoT/Cloud Data Streams Tools

The FIESTA EaaS approach to IoT experimentation is based on the semantic interoperability of diverse platforms. To this end the project exploits and extends recent developments in the area of semantic interoperability of IoT data streams. In the general area of data stream management for IoT, the landscape is divided between two major approaches for data stream processing [7, 48]: (i) in-network processing, which is close in essence to the Wireless Sensors and Actuators Networks (WSANs) work (peer-to-peer communication), and (ii) cloud-based processing, related to big data approaches (centralized client-server communication, where the cloud can be considered as an elastic server). With regard to (i), Data Stream Management Systems (DSMS) for WSANs may be classified into three broad families as follows:

- **Relational DSMSs** [2, 40] extend the relational model by adding concepts necessary to handle data streams and persistent queries, together with the stream-oriented version of the relational operators (selection, projection, union, etc.). State of the art DSMSs primarily

differ with respect to: the expressiveness of the query language, the associated algebra, and the assumptions made about the underlying networking architecture. More specialized proposals [20] deal with issues as diverse as blocking and non-blocking operators, windows, stream approximation, and optimizations.

- **Macroprogramming-based DSMS** [42] are oriented toward the development of applications over WSANs, as opposed to the expression of data queries over the network. The macroprograms are typically specified using a domain-specific language, and are compiled into microprograms to be run on the networked nodes, hence easing the developer's tasks who has no longer to bother with the decomposition and further distribution of the macroprograms.
- **Service-oriented DSMSs** [38] aim at integrating with classical service-oriented architectures, thereby allowing to exploit the functionalities of the infrastructure (interaction and discovery protocols, registries, service composition based on orchestration or choreography, security infrastructure, etc.).

Cloud-based approaches, on the other hand, rely on the cloud infrastructure to collect, process and store the data acquired from the environment. In contrast to DSMSs for WSANs, cloud computing offers a simple way to perform easily a wide range of heavy computations and to deal with ultra-large streams at ultra-large scales [41, 52]. These characteristics make the cloud an interesting solution for the IoT, given the expected scale. The convergence between the cloud and the IoT, referred as "Cloud of Things", is relatively recent [52] and is pioneered by emerging approaches such as Sensor Clouds [63], IoT platforms [35] and Sensing-as-a-Service [64]. Basically, all approaches share common features and follow the same global process: sensor providers (users, cities, companies, etc.) join the Cloud of Things (CoT) by registering their sensors or sensor networks. Users can send requests to the CoT, which then collects data from a set, or a representative subset, of sensors that match the requirements of the requests. These data are processed by the CoT according to the computation expressed by the request, and the results are sent back to the users.

When combining IoT data streams originating from different sources, one can leverage semantic technologies for achieving interoperability. Most of the existing (semantic interoperability) efforts to provide uniform representations for entities in the Internet of Things (IoT), i.e., Things, sensors/actuators they host, and services they provide, adopt the semantic approach and

exploit ontologies. A considerable portion of ontologies exploited in the IoT domain is inherited from efforts in the Wireless Sensor Networks domain. In the latter, the main focus is directed towards modeling sensor, actuators and their data (e.g., [15, 16, 47]). A commonly exploited ontology, to reason about sensors is SSN [16], provided by the W3C Semantic Sensor Network Incubator Group. SSN models sensing specific information from different perspectives: a) Sensor perspective to model sensors, what they sense and how they sense; b) System perspective to model systems of sensors and their deployment; c) Feature perspective to model what senses a specific property and how it is sensed; d) Observation perspective to model observation values and their metadata. Other sensor ontologies are also surveyed in [18]. Many of the ontologies surveyed therein provide a solid basis for the representation of sensors, actuators, and their data. However, those entities are only a portion of the IoT.

More efforts have been made recently to extend the ontologies with IoT-specific semantics, including Things, their functionalities, or their deployment spaces. For instance, Sense2Web [19] provides an ontology that models the following Thing-related concepts: the Entity (equivalent to a feature on interest); the Device, which is the hardware component (equivalent to a Thing); the Resource, which is a software component representing the entity; and the Service through which a resource is accessed. A resource can be a sensor, actuator, RFID tag, processor or a storage resource. Christophe et al. [14] focus more on the deployment spaces of Things rather than Things themselves, especially indoor locations. The ontologies provided by the authors provides a vocabulary to describe Objects, which are physical Things, their location, their capability, and virtual objects, which are higher level abstractions of the Things combining the above information together. Another example is the work in [62] where authors present an ontology that models services provided by Things; deployment information; Observations; Entities, which are real-world features to measure/act on, and finally Things.

10.6.5 FIESTA Approach on Semantic Interoperability for IoT/Cloud Resources Tools

FIESTA's work on semantic interoperability for IoT and Cloud resources that focuses on developing common annotation models for describing the resources and IoT data and providing validation and testing tools for semantic interoperability evaluation. The core models are constructed by investigating

the existing semantic and ontology models including the IoT-A information models (i.e. resources, service and entity models developed in the FP7 EU IoT-A project, <http://epubs.surrey.ac.uk/127271/>), W3C Semantic Sensor Network Ontology (SSN Ontology) (<http://www.w3.org/2005/Incubator/ssn/XGR-ssn-20110628/>), EF7 IoT.est models). FIESTA uses the existing concepts, namespaces and semantic models and develops a set of core models to describe IoT resources (e.g. sensor devices, gateways, actuators) and their capabilities and features and also provide semantic models to describe Cloud services and Cloud based components. The existing semantic models such as W3 SSN, IoT-A models are usually developed for specific purposes and in the domain of the projects.

10.6.6 FIESTA Approach on Testbeds Integration and Federation Tools

Federation in FIESTA is understood to be: “an organization within which smaller divisions have some internal autonomy” [43]. In terms of testbeds this considers that each testbed operates both individually and part of a larger federation in order to gain value (larger user base, potential combinations with other testbeds to support richer experimentation, etc.). Typical testbed federation functions include: resource discovery (finding the required resources for an experiment); resource provisioning (management of resources such that they are available when required); resource monitoring (monitor operation in order to collect experimental results); and finally security (ensuring authorized users can access resources, and the federation provides a trusted base to keep experiment information secure). Different federation models can then be applied to implement the federation; the FedSM project defines a number of models including lightweight federation where there is little if any central control of these functions (by the federation) through to a fully integrated model where a central federation authority implements and provides the functions.

The FIRE programme has a long standing history in developing cutting edge testbed federations. In the field of networking research: Openlab provides access to tools and testbeds including PlanetLab Europe, the NITOS wireless testbed, and other federated testbeds to support networking experimentation across heterogeneous facilities. OFELIA is an OpenFlow switching testbed in Europe federating a number of OpenFlow islands supporting research in the Software Defined Networking field. CONFINE co-ordinates unified access to a set of real-world community IP networks (wired, wireless,

ad-hoc, etc.) to openly allow research into service, protocols and applications across these edge networks. CREW federates five wireless testbeds to support experimentation with advanced spectrum sensing and cognitive radio. Finally, FLEX is a new FIRE project that works towards providing testbeds for LTE experimentation. In the field of software services, the Bonfire project created a federation of cloud facilities to support experimentation with new cloud technologies. Importantly, in terms of Internet of Things testbeds, SmartSantander provides a set of Smart City facilities through large-scale deployments of sensor networks atop which applications and services can be developed. Also, Sunrise is a federation of sensor network testbeds providing monitoring and exploration of the marine environments and in particular supporting experimentation in terms of the underwater Internet of Things. While each project typically performs federation within its own domain, the Fed4FIRE project is an initiative to bring together heterogeneous facilities across Europe so as to target experimentation across the whole Future Internet field i.e., networks, software and services, and IoT.

Many of the projects (crucially Fed4FIRE) employ OMF [45] and SFA [8] federation technologies. OMF is a control, measurement and management framework for testbeds. From an experimenter's point of view, OMF provides a set of tools to describe and instrument an experiment, execute it and collect its results. From a testbed operator's point of view, OMF provides a set of services to efficiently manage and operate the testbed resources (e.g. resetting nodes, retrieving their status information, installing new OS image). The OMF architecture is based upon Experiment Controllers that steer experiments defined in OEDL (OMF experiment Description Language), which is a declarative domain-specific language describing required resources and how they should be configured and connected. It also defines the orchestration of the experiment itself.

Outside FIRE, there have been a number of federation initiatives to support the wider Future Internet community. Two relevant ones are Helix Nebula and XIFI. XIFI is a federation of data centres connected to resources such as wireless testbeds and sensor networks; its goal is to support large-scale Future Internet trials before transfer to market. XIFI employs a federation architecture based around web technologies (e.g. OAUTH, OCCI, and open Web APIs). On the other hand, Helix Nebula – the Science Cloud is an initiative to build federated cloud services across Europe in order to underpin IT-intense scientific research while also allowing the inclusion of other stakeholders' needs (governments, businesses and citizens).

10.7 FIESTA-IoT Architecture

FIESTA deals with the federation, virtualization and interoperability of diverse IoT testbeds, notably testbeds that comply with different IoT architectures, including architectures developed by standardization bodies (e.g., OGC [44] and GS1/EPCGlobal [22]), as well as FP7 projects (such as SmartSantander [25]). These architectures serve application specific purposes and are characterized by increased penetration in specific industries. In addressing this heterogeneity, FIESTA attempt to map and describe IoT platforms complying with these architectures to a general-purpose meta-level architecture, which serves as a basis for the FIESTA virtualized architecture layer Figure 10.4. The foundation for developing this meta-architecture is the Architecture Reference Model [34], developed by the FP7 IoT-A project and the IERC cluster. The current status reached by IoT-A at the end of the project (November 2013), as far as the Architectural Reference Model (ARM) is concerned includes a set of Models, Views & Perspectives in addition to a comprehensive set of guidelines that explains how to use Model, Views and Perspectives in order




Figure 10.4 FIESTA functional blocks architecture.

to derive a concrete architecture. Part of the Guidelines is a large set of design choices that are linked to the perspectives, i.e. linked to some qualities that the system is expected to meet; part of those properties is system interoperability.

While the ARM provides some recommendations and tactics in order to achieve system interoperability, this essential quality is not guaranteed when applying the ARM to a concrete system as choices are left in the architect's hands. In order to boost the adoption of the ARM and make its usage easier the IoT Forum considers that the next step is to develop specific profiles that implement flavours of the ARM focusing on specific qualities of the system (e.g., ability to handle specific functional or non-functional requirements such as reliability, resilience, QoS awareness).

10.8 Conclusions

Fiesta has advanced the state of the art in different directions, relevant activities are on-going work but most of the progress related with design, architecting and implementation have been completed and reported, in the various public documents, in this chapter we summarize the relevant contributions in the different relevant areas where FIESTA has worked so far.

FIESTA Federation: Currently, there is no easy way to carry out experiments across a range of IoT facilities without having deep knowledge in sensor networks technology, communication technologies and platform configuration. FIESTA has opened up this space to provide a richer experimentation space that appeals to a wider range of experimenters (both in industry and research). Existing federation technologies are typically heavyweight in the effort required to add and control testbeds; in the case of OMF an experimental controller needs to be integrated with the facility so that standard conformance is achieved. While standards solve integration problems they often do so in a way that hinders long term sustainability (detracts new joiners)—new IoT testbeds must be able to quickly come and go as new technology trends emerge. FIESTA's approach to federation, built up semantic interoperability technologies and the meta-cloud infrastructure to provide novel methods to ensure that testbeds can be integrated in a lightweight manner and ensure that sustainability is not hindered.

FIESTA Architecture: Since 2014 the IoT Architecture Reference Model (IoT-ARM) sustenance and profile work is taken care of by the WG "Technology and Openness" of the IoT Forum. FIESTA has contributed to the definition of the "Semantic Interoperability" profile based on the FIESTA achievements

in that matter and envisages getting ARM/profile certification for some of the “Semantic Interoperability” enablers implemented in the course of the project. Overall, FIESTA maps several concrete architectures to the IoT-ARM, as a means to studying and ensuring their interoperability. The testbeds to be interconnected and virtualized in the project are the starting point for these mappings, while additional mappings can be realized as part of open calls that is well know ask for the federation of additional testbeds.

FIESTA Semantic Interoperability: A common limitation to all surveyed ontologies is that they still mostly lack a very important requirement: modeling the physics and mathematics, which are at the core of any sensing/actuation task, as first class entities. In more detail, it is important to relate various quantifiable and measurable (real-world) features in order to define, in a user understandable and machine-readable manner the processes behind single or combined sensing/actuation tasks. This correlation enables the system exploiting the ontologies to have a better understanding of the sensing/actuating task at hand and consequently better analyse its outcomes or substitute it more efficiently if need be, i.e., if required sensors/actuators are not available, or if the functionalities they provide do not fully satisfy the task at hand. FIESTA deals with these interoperability issues to allow researchers to design and submit interoperable experiments that are able to understand the semantics of sensing and actuating tasks and accordingly to select sensors/actuators that are suitable for executing the specified experiments. As a starting point the project leverages the W3C SSN ontology, along with mathematical formulas introduced in [28] in order to represent sensing/acting processes in a universally accepted language (i.e. algebra). FIESTA deploy middleware implementing such algebra over the federated testbeds, as well as appropriate mapping techniques for streaming tasks, in order to allow researchers to specify experiments based on combinations/compositions of sensing and actuating processes. From an implementation perspective, FIESTA deploys middleware (residing at the individual testbeds) endowing the testbeds with interoperability capabilities, along with middleware (residing at the FIESTA meta-cloud infrastructure) empowering discovery of IoT resources and compositions of sensing and actuating processes from multiple testbeds.

FIESTA has progress the state-of-the-art by introducing re-usable and common core models to describe the IoT and Cloud resources. The built models are based on the existing and common IoT models to maximize the interoperability among different providers and test-beds. At the design level, FIESTA provide semantic interoperability check and validation services

using a common portal and web services to allow service developers and test-bed providers check and evaluate the interoperability of their meta-data descriptions based on the FIESTA core models and also other existing common models. The results of semantic interoperability check and evaluation gives feedback to the semantic model designers and test-bed providers on the level of interoperability between their resource descriptions and the commonly used resource description frameworks. At the deployment level, FIESTA provides wrappers and matching services to enable translation of the resource descriptions from the test-beds to the FIESTA's core models and/or other existing common models. At the run-time level, FIESTA enables test-bed providers and Cloud service developers to publish, query and access large set of semantically annotated resource descriptions according to different semantic description models and representation frameworks (i.e. using different semantic models and also different representation formats). This enables the test-bed and Cloud service providers and developers to test and evaluate efficiency of different solutions and also to measure the level of interoperability between different schemes and also to enable the resource providers to adapt common models or use wrapper to enhance the semantic interoperability between their resource descriptions and other resources that are described within the FIESTA framework that are distributed over different test-beds.

FIESTA Linking and Reasoning: FIESTA have improved the state of the art in this area by providing highly efficient approaches for efficient processing of linked data streams typical for applications in the IoT and smart cities areas. FIESTA's work is based on CQELS. Based on this basic reasoning functionality, the project provides layered reasoning formalisms at different levels of complexity (uncertainty, nonmonotonicity, recursion) for adaptive trade-offs between scalability and expressivity as required by experimental applications in the areas addressed by the FIESTA testbeds.

FIESTA Federating IoT Data Streams: FIESTA has advanced the state-of-the-art in federated processing for IoT data through enabling semantic-based interoperability among stream processing engines using Linked Stream Data. FIESTA enables semantically-self-described stream data items to automatically travel from its point of origin (e.g., sensors) downstream to applications, through autonomously passing through many stream engines. Each of the stream engines might provide potential stream data for the targeted stream-based computation that can be expressed in a standardized continuous query language, i.e, an extension of SPARQL [57]. FIESTA also support automatic discovery of stream data at run-time based on context represented as semantic

links in stream data. This enables the federation of schema-free and semantic-based data aggregation without prior knowledge about stream data format, data schema and origins of the input stream data. FIESTA also has targeted the provisioning for a standardized RDF-based stream protocol to facilitate the semantic-based interoperability among the federation setting.

Acknowledgements

This work is funded by the European Commission under the EU-H2020 Project Grant "Federated Interoperable Semantic IoT/cloud Testbeds and Applications (FIESTA)" with the Grant Agreement No. CNECT-ICT-643943.

References

- [1] Daniel J. Abadi, Yanif Ahmad, Magdalena Balazinska, Ugur Cetintemel, Mitch Cherniack, Jeong-Hyon Hwang, Wolfgang Lindner, Anurag S. Maskey, Alexander Rasin, Esther Ryzkina, Nesime Tatbul, Ying Xing, and Stan Zdonik. "The Design of the Borealis Stream Processing Engine," Proc. 2nd Biennial Conf. Innovative Data Systems Research (CIDR 05), 2005; <http://www-db.cs.wisc.edu/cidr>.
- [2] Giuseppe Amato, Stefano Chessa, and Claudio Vairo. Mad-wise: a distributed stream management system for wireless sensor networks. *Software: Practice and Experience*, 40(5), 2010.
- [3] D. Anicic, P. Fodor, S. Rudolph, and N. Stojanovic. EP-SPARQL: a unified language for event processing and stream reasoning. In *Proceedings of the 20th International World-Wide Web Conference*, pages 635–644. ACM, 2011.
- [4] Scott Kirkpatrick (lead), Michael Boniface et al., «FUTURE INTERNET RESEARCH AND EXPERIMENTATION: VISION AND SCENARIOS 2020», AMPLIFIRE Support Action (Grant Agreement: 318550) Deliverable D1.1, June 2013.
- [5] Sathya Rao, Martin Potts, et al., «FIRE PORTFOLIO CAPABILITY ANALYSIS», AMPLIFIRE Support Action (Grant Agreement: 318550) Deliverable D2.1, October 2013.
- [6] Darko Anicic, Paul Fodor, Sebastian Rudolph, and Nenad Stojanovic. 2011. EP-SPARQL: a unified language for event processing and stream reasoning. In *Proceedings of the 20th international conference on World wide web (WWW '11)*. ACM, New York, NY, USA, 635–644.

- DOI=10.1145/1963405.1963495 <http://doi.acm.org/10.1145/1963405.1963495>
- [7] Luigi Atzori, Antonio Iera, and Giacomo Morabito. The internet of things: A survey. *Computer Networks*, 54(15), 2010.
- [8] Augé, J. and Friedman, T (2012) The Open Slice-based Facility Architecture available at <http://opensfa.info/doc/opensfa.html>
- [9] Magdalena Balazinska, Amol Deshpande, Michael J. Franklin, Phillip B. Gibbons, Jim Gray, Mark Hansen, Michael Liebhold, Suman Nath, Alexander Szalay, and Vincent Tao. 2007. Data Management in the Worldwide Sensor Web. *IEEE Pervasive Computing* 6, 2 (April 2007), 30–40. DOI=10.1109/MPRV.2007.27 <http://dx.doi.org/10.1109/MPRV.2007.27>
- [10] Davide Francesco Barbieri, Daniele Braga, Stefano Ceri, and Michael Grossniklaus. 2010. An execution environment for C-SPARQL queries. In *Proceedings of the 13th International Conference on Extending Database Technology (EDBT '10)*, Ioana Manolescu, Stefano Spaccapietra, Jens Teubner, Masaru Kitsuregawa, Alain Leger, Felix Naumann, Anastasia Ailamaki, and Fatma Ozcan (Eds.). ACM, New York, NY, USA, 441–452. DOI=10.1145/1739041.1739095 <http://doi.acm.org/10.1145/1739041.1739095>
- [11] Jason Campbell, Phillip B. Gibbons, Suman Nath, Padmanabhan Pillai, Srinivasan Seshan, and Rahul Sukthankar. 2005. IrisNet: an internet-scale architecture for multimedia sensors. In *Proceedings of the 13th annual ACM international conference on Multimedia (MULTIMEDIA '05)*. ACM, New York, NY, USA, 81–88. DOI=10.1145/1101149.1101162, <http://doi.acm.org/10.1145/1101149.1101162>
- [12] Sirish Chandrasekaran, Owen Cooper, Amol Deshpande, Michael J. Franklin, Joseph M. Hellerstein, Wei Hong, Sailesh Krishnamurthy, Samuel R. Madden, Fred Reiss, and Mehul A. Shah. 2003. TelegraphCQ: continuous dataflow processing. In *Proceedings of the 2003 ACM SIGMOD international conference on Management of data (SIGMOD '03)*. ACM, New York, NY, USA, 668–668. DOI=10.1145/872757.872857 <http://doi.acm.org/10.1145/872757.872857>
- [13] Mitch Cherniack, Hari Balakrishnan, Magdalena Balazinska, Don Carney, Ugur Cetintemel, Ying Xing, and Stan Zdonik, “Scalable Distributed Stream Processing,” *Proc. 1st Biennial Conf. Innovative Data Systems Research (CIDR 03)*, 2003, pp. 257–268; <http://www-db.cs.wisc.edu/cidr>.

- [14] Benoit Christophe. Managing massive data of the Internet of Things through cooperative semantic nodes. In IEEE Sixth International Conference on Semantic Computing, (ICSC), pages 93–100. IEEE, 2012.
- [15] Michael Compton, Holger Neuhaus, Kerry Taylor, and Khoi-Nguyen Tran. Reasoning about sensors and compositions. In SSN, pages 33–48, 2009.
- [16] M. Compton, P. Barnaghi, L. Bermudez, R. G. Castro, O. Corcho, S. Cox, et al.: “The SSN Ontology of the Semantic Sensor Networks Incubator Group”, *Journal of Web Semantics: Science, Services and Agents on the World Wide Web*, ISSN 1570–8268, Elsevier, 2012.
- [17] Philippe Cousin, Martin Serrano, John Soldatos, «Internet of Things Research on Semantic Interoperability to address Manufacturing Challenges», In the Proc. Of the 7th International Conference on Interoperability for Enterprise Systems and Applications, I-ESA 2014, Albi, France, March, 24–28, 2014.
- [18] Mathieu d’Aquin and Natalya F Noy. Where to publish and find ontologies? A survey of ontology libraries. *Web Semantics: Science, Services and Agents on the World Wide Web*, 11:96–111, 2012.
- [19] Suparna De, Tarek Elsaleh, Payam Barnaghi, and Stefan Meissner. An Internet of Things platform for real-world and digital objects. *Scalable Computing: Practice and Experience*, 13(1), 2012.
- [20] Nihal Dindar, Nesime Tatbul, Renee J. Miller, Laura M. Haas, and Irina Botan. Modeling the execution semantics of stream processing engines with secret. *The VLDB Journal*, 2012.
- [21] T. Do, S. Loke, and F. Liu. Answer set programming for stream reasoning. *Advances in Artificial Intelligence*, pages 104–109, 2011.
- [22] EPCglobal: The EPCglobal Architecture Framework, EPCglobal Final Version 1.2 Approved 10 September 2007.
- [23] Geoffrey C. Fox, Supun Kamburugamuve, Ryan Hartman Architecture and Measured Characteristics of a Cloud Based Internet of Things API Workshop 13-IoT Internet of Things, Machine to Machine and Smart Services Applications (IoT 2012) at The 2012 International Conference on Collaboration Technologies and Systems (CTS 2012) May 21–25, 2012 The Westin Westminster Hotel Denver, Colorado, USA.
- [24] Michael J. Franklin, Shawn R. Jeffery, Sailesh Krishnamurthy, Frederick Reiss, Shariq Rizvi, Eugene Wu, Owen Cooper, Anil Edakkunni, Wei Hong: Design Considerations for High Fan-In Systems: The HiFi Approach. *CIDR 2005*: 290–304

- [25] J.A. Galache, V. Gutiérrez, J.R. Santana, L. Sánchez, P. Sotres, J. Casanueva, L. Muñoz, “SmartSantander: A Joint service provision facility and experimentation-oriented tesbed, within a smart city enviroment”, Future Network & Mobile Summit, Lisbon (Portugal), July 2013.
- [26] Forecast: The Internet of Things, Worldwide, 2013. <http://www.gartner.com/document/2625419?sthkw=G00259115>
- [27] Vincenzo Gulisano, Ricardo Jimenez-Peris, Marta Patino-Martinez, Claudio Soriente, and Patrick Valduriez. 2012. StreamCloud: An Elastic and Scalable Data Streaming System. *IEEE Trans. Parallel Distrib. Syst.* 23, 12 (December 2012), 2351–2365. DOI=10.1109/TPDS.2012.24 <http://dx.doi.org/10.1109/TPDS.2012.24>
- [28] Sara Hachem, Thiago Teixeira, and Valérie Issarny. 2011. Ontologies for the internet of things. In *Proceedings of the 8th Middleware Doctoral Symposium (MDS '11)*. ACM, New York, NY, USA, Article 3, 6 pages.
- [29] Mohammad Mehedi Hassan, Biao Song, Eui-nam Huh: A framework of sensor-cloud integration opportunities and challenges. *ICUIMC 2009*: 618–626
- [30] Tom Heath and Christian Bizer (2011) *Linked Data: Evolving the Web into a Global Data Space* (1st edition). *Synthesis Lectures on the Semantic Web: Theory and Technology*, 1:1, 1–136. Morgan & Clay-pool.
- [31] Jesper Hoeksema and Spyros Kotoulas. High-performance Distributed Stream Reasoning using S4, Ordring workshop of ISWC2011. 2011.
- [32] Hogan A., Mellotte M., Powell G., and Stampouli D. (2012) Towards fuzzy query-relaxation for RDF. In *Proceedings of the 9th international conference on The Semantic Web: research and applications (ESWC'12)*, Elena Simperl, Philipp Cimiano, Axel Polleres, Oscar Corcho, and Valentina Presutti (Eds.). Springer-Verlag, Berlin, Heidelberg, 687–702.
- [33] Martin Bauer (IOT-A), Paul Chartier (CEN TC225), Klauss Moessner (IOT.est), Nechifor, Cosmin-Septimiu (iCore), Claudio Pastrone (ebbits), Josiane Xavier Parreira (GAMBAS), Richard Rees (CEN TC225), Domenico Rotondi (IoT@Work), Antonio Skarmeta (IoT6), Francesco Sottile (BUTLER), John Soldatos (OpenIoT), Harald Sundmaeker (SmartAgriFood), «Catalogue of IoT Naming, Addressing and Discovery Schemes in IERC Projects», electronically available at: <http://www.theinternetofthings.eu> <http://www.theinternetofthings.eu/sites/default/files/%5Buser-name%5D/IERC-AC2-D1-v1.7.pdf>

- [34] Andreas Nettsträter, Martin Bauer, et al. Internet-of-Things Architecture Project (IoT-A), Deliverable D1.3, «Updated reference model for IoT v1.5», July 2012.
- [35] M. Kovatsch, M. Lanter, and S. Duquennoy. Actinium: A RESTful runtime container for scriptable Internet of Things applications. In Proc. of the 3rd International Conference on the Internet of Things, IOT '12, 2012.
- [36] A. Kozlenkov and M. Schroeder. Prova: Rule-based java-scripting for a bioinformatics semantic web. In Data Integration in the Life Sciences, pages 17–30, Springer, 2004.
- [37] N. Lanzanasto, S. Komazec, and I. Toma. Reasoning over real time data streams. <http://www.envision-project.eu/wpcontent/uploads/2012/11/D4-8v1-0.pdf>, 2012.
- [38] Le-Phuoc, D., Dao-Tran, M., Xavier Parreira, J., & Hauswirth, M. (2011). A native and adaptive approach for unified processing of linked streams and linked data. The Semantic Web–ISWC 2011, 370–388.
- [39] Kevin Lee, «Extending Sensor Networks into the Cloud using Amazon Web Services», IEEE International Conference on Networked Embedded Systems for Enterprise Applications 2010, 25th November, 2010.
- [40] Samuel R. Madden, Michael J. Franklin, Joseph M. Hellerstein, and Wei Hong. Tinydb: an acquisitional query processing system for sensor networks. ACM Trans. Database Syst., 30(1), March 2005.
- [41] Samuel Madden. From databases to big data. IEEE Internet Computing, 16, 2012.
- [42] Ryan Newton, Greg Morrisett, and Matt Welsh. The regiment macroprogramming system. In Proceedings of the 6th international conference on Information processing in sensor networks, IPSN '07, pages 489–498, New York, NY, USA, 2007. ACM.
- [43] “federation, n.”. OED Online. March 2014. Oxford University Press. <http://www.oed.com/view/Entry/68930?redirectedFrom=federation> (accessed March 28, 2014).
- [44] Open Geospatial Consortium, OpenGIS Sensor Observation Service version 1.0.0, Arthur Na, Mark Priest, 2008-02-13.
- [45] Thierry Rakotoarivelo, Max Ott, Guillaume Jourjon, Ivan Seskar, “OMF: a control and management framework for networking testbeds”, in ACM SIGOPS Operating Systems Review 43 (4), 54–59, Jan. 2010.
- [46] RDF Stream Processing Community Group, <http://www.w3.org/community/rsp/>

- [47] David J Russomanno, Cartik R Kothari, and Omoju A Thomas. Building a sensor ontology: A practical approach leveraging ISO and OGC models. In IC-AI, pages 637–643, 2005.
- [48] S. Sagioglu and D. Sinanc. Big data: A review. In Collaboration Technologies and Systems (CTS), 2013 International Conference on, 2013.
- [49] H. Schaffers, A. Sällström, M. Pallot, J. Hernandez-Munoz, R. Santoro, B. Trousse, “Integrating Living Labs with Future Internet Experimental Platforms for Co-creating Services within Smart Cities”. Proceedings of the 17th International Conference on Concurrent Enterprising, ICE’2011, Aachen, Germany, June 2011.
- [50] Gregor Schiele, John Soldatos, Paul Lefrere, Nathalie Mitton, Kahina Hamadache, Manfred Hauswirth ‘Towards Interoperable IoT Deployments in Smart Cities’, submitted for publication to the European Conference on Networks and Communications EuCNC, Bologna, Italy, June 23–26, 2014 (EuCNC 2014).
- [51] Martin Serrano, Manfred Hauswirth, John Soldatos, “Design Principles for Utility-Driven Services and Cloud-Based Computing Modelling for the Internet of Things”, International Journal of Web and Grid Services (to appear), 2014.
- [52] J. Soldatos, M. Serrano and M. Hauswirth. “Convergence of Utility Computing with the Internet-of-Things”, International Workshop on Extending Seamlessly to the Internet of Things (esIoT), collocated at the IMIS-2012 International Conference, 4th6th July, 2012, Palermo, Italy.
- [53] Bernhard Stegmaier, Richard Kuntschke, and Alfons Kemper. 2004. StreamGlobe: adaptive query processing and optimization in streaming P2P environments. In Proceedings of the 1st international workshop on Data management for sensor networks: in conjunction with VLDB 2004 (DMSN ’04). ACM, New York, NY, USA, 88–97. DOI=10.1145/1052199.1052214 <http://doi.acm.org/10.1145/1052199.1052214>
- [54] Distributed and fault-tolerant realtime computation. <http://storm.incubator.apache.org/>
- [55] Distributed Stream Computing Platform. <http://incubator.apache.org/s4/>
- [56] Lightning-fast cluster computing. <http://spark.apache.org/>
- [57] SPARQL 1.1 Query Language. <http://www.w3.org/TR/sparql11-query/>
- [58] Kerry Taylor, «Semantic Sensor Networks: The W3C SSN-XG Ontology and How to Semantically Enable Real Time Sensor Feeds», 2011 Semantic Technology Conference, June 5–9, San Francisco CA, USA.

- [59] K. Teymourian, M. Rohde, and A. Paschke. Fusion of background knowledge and streams of events. In Proceedings of the 6th ACM International Conference on Distributed Event-Based Systems, DEBS '12, pages 302–313, New York, NY, USA, 2012, ACM.
- [60] Umbrich J., Karnstedt M., Hogan A., Parreira J. (2012) Hybrid SPARQL queries: fresh vs. fast results. In Proceedings of the 10th International Semantic Web Conference, 2012.
- [61] Ovidiu Vermesan & Peter Friess, «Internet of Things - Global Technological and Societal Trends», The River Publishers Series in Communications, May 2011, ISBN: 9788792329738.
- [62] Wei Wang, Suparna De, Ralf Toenjes, Eike Reetz, and Klaus Moessner. A comprehensive ontology for knowledge representation in the Internet of Things. In Trust, Security and Privacy in Computing and Communications (TrustCom), 2012 IEEE 11th International Conference on, pages 1793–1798. IEEE, 2012.
- [63] M. Yuriyama and T. Kushida. Sensor-cloud infrastructure – physical sensor management with virtualized sensors on cloud computing. In Network- Based Information Systems (NBIS), 2010 13th International Conference on, 2010.
- [64] Jia Zhang, Bob Iannucci, Mark Hennessey, Kaushik Gopal, Sean Xiao, Sumeet Kumar, David Pfeffer, Basmah Aljedia, Yuan Ren, Martin Griss, et al. Sensor data as a service – a federated platform for mobile data-centric service development and sharing. In Services Computing (SCC), 2013 IEEE International Conference on, 2013.

