Open Access

Open Journal of Internet of Things (OJIOT)
Volume 6, Issue 1, 2020
http://www.ronpub.com/ojiot
ISSN 2364-7108

Information-Centric Semantic Web of Things

Michele Ruta, Floriano Scioscia

Department of Electrical and Information Engineering, Polytechnic University of Bari,
Via Orabona 4, I-70125, Bari, Italy, {michele.ruta, floriano.scioscia}@poliba.it

ABSTRACT

In the Semantic Web of Things (SWoT) paradigm, a plethora of micro-devices permeates an environment. Storage and information processing are decentralized: each component conveys and even processes a (very) small amount of annotated metadata. In this perspective, the node-centric Internet networking model is inadequate. This paper presents a framework for resource discovery in semantic-enhanced pervasive environments leveraging an information-centric networking approach. Information gathered through different Internet of Things (IoT) technologies can be exploited by both ubiquitous and Web-based semantic-aware applications through a uniform set of operations. Experimental results and a case study support sustainability and effectiveness of the proposal.

TYPE OF PAPER AND KEYWORDS

Regular research paper: *Semantic Web, Pervasive computing, Service discovery, Information-centric networking, Peer-to-peer protocols*

1 INTRODUCTION

The *Semantic Web of Things* (SWoT) should enhance the *Internet of Things* (IoT) capabilities and applications leveraging the *Semantic Web* theory and technology. It aims to embed semantically rich and always-available information fragments into the physical world, by enabling storage/retrieval of annotations in/from tiny smart objects. Traditional Internet-based applications (email, Web browsing, VoIP) adopt a node-centric conversational model since the TCP/IP protocol stack is based on a request/response approach. This model is not sustainable in pervasive environments, where user agents may interact simultaneously with many surrounding micro-devices. Retrieved information should be automatically processed to better support current activity of a user in an unobtrusive fashion, s/he even being not necessarily aware of individual device interactions. Hence, the SWoT vision should be favored by pervasive knowledge-based systems achieving high degrees of autonomic capability in information storage, management and discovery, also providing transparent access to information sources. From this standpoint, *information-centric* (*a.k.a.* data-centric, content-centric) network infrastructures are more suitable, as they are centered on information resources, and not on computer hosts. While research in sensor networks has already shown [18] that a content-centric approach is more effective than a node-centric one for information exchange in resource-constrained and volatile contexts, this is increasingly true even for the Internet itself. Trends in Internet and Web usage show that *what* is being exchanged is becoming more important than *who* are exchanging information [46]. An exhaustive technical comparison between node-centric and content centric networking is in [20]. In information-centric networks, and particularly in dynamic environments such as the SWoT, resource discovery is a pivotal feature.

This paper is accepted at the *International Workshop on Very Large Internet of Things (VLIoT 2020)* in conjunction with the VLDB 2020 conference in Tokyo, Japan. The proceedings of VLIoT@VLDB 2020 are published in the Open Journal of Internet of Things (OJIOT) as special issue.

Nevertheless, current paradigms employ elementary "string matching" to compare requests and encoded resource attributes. This is too simplistic to be useful for advanced applications. Similarly, in most information-centric networking approaches in literature, retrieval is based on structured "content addressing" schemes that are rather arbitrary and need case-by-case agreements among involved actors to allow practical interoperability. Ideas and technologies borrowed from the Semantic Web vision may allow to overcome these limitations, since the formal Knowledge Representation (KR) foundations provide common vocabularies to describe resources and express requests.

In this paper a semantic content-centric discovery framework is proposed, where autonomous objects can be retrieved, queried and inventoried without requiring a central control and coordination. A proper dissemination protocol allows to exactly locate suitable annotated descriptions directly on embedded micro-devices attached to items of interest. The proposed framework includes a peer-to-peer distributed application-layer protocol allowing dissemination and discovery of knowledge. Information is gathered through different identification and sensing technologies to be exploited by inference engines and semantic-aware applications, in either pervasive Mobile Ad-hoc Networks (MANETs) or Web contexts, through a uniform set of operations. Mobile nodes capable of extracting information from embedded micro-devices work as cluster-heads w.r.t. resources in their range. IEEE 802.11, along with IP (Internet Protocol) and UDP (User Datagram Protocol), is adopted as reference network infrastructure. Resource Description Framework [40] is the resource annotation language w.r.t. RDF Schema (RDFS) [6] vocabularies, so allowing semantic-based applications to leverage querying, reasoning and matchmaking tools, based on formal logics. Noteworthy features of the proposed framework are: (i) peer-to-peer cooperation for managing environments populated by autonomous objects in a distributed fashion; (ii) semantic-enabled content-centric discovery; (iii) reactive dissemination to quickly locate semantic annotations on micro-devices attached to items in the field; (iv) backward compatibility w.r.t. standard identification and sensing technologies, thus allowing legacy applications to co-exist with novel ones; (v) dissemination and discovery integrated at the application level on top of IP and UDP, in order to preserve compatibility with standard protocols and equipment for end-to-end routing in the Internet [20]. The feasibility of the approach has been demonstrated by a simulation campaign exploiting a network simulator package assessing effectiveness of metadata dissemination and discovery protocol architecture. Benefits of the proposed approach are illustrated in a case study for an IoT scenario about disaster recovery.

The remaining of the paper is organized as follows: in the next section motivation and possible application scenarios are provided. In Section 3 relevant related work is discussed. Architectural aspects are detailed in Section 4, along with the data propagation and service retrieval protocol. The subsequent case study clarifies the framework through an illustrative example. Simulation methods and results are presented and commented in Section 5 before conclusion.

2 MOTIVATION

Each resource in the semantic-enabled Web -not only documents and services, but any entity of interest such as people, institutions, common knowledge topics- can be annotated with metadata, using RDF, w.r.t. an RDFS or Web Ontology Language (OWL 2) [29] vocabulary. Information interlinking produces *knowledge graphs* whose basic unit is the *(subject, predicate, object)* statement. URIs (Uniform Resource Identifiers) are used to unambiguously identify both resources and predicates that relate them. On top of RDF, RDFS and OWL enable KR based on formal model-theoretic semantics, which allows the use of existing reasoning engines to infer new information from the one stated in the semantic annotations. In the majority of current applications -see Section 1.5 of [2] for a survey- Knowledge Representation Systems (KRSs) play a role very much like Database Management Systems (DBMSs). Both are used as single fixed entities which are immediately available for information queries and updates. This approach is effective only as long as large computing resources and a stable network infrastructure are granted. A different strategy is needed to adapt KR tools and technologies to functional and non-functional requirements of pervasive computing applications as the Internet of Things. They are characterized by user and device mobility and dependency on context. Furthermore, severe resource limits affect processing, storage, link bandwidth and power consumption. As a consequence, knowledge-based systems conceived for wired networks are hardly adaptable, due to architectural differences and performance issues. Anyway, several emerging identification, monitoring and sensing technologies are now suitable to connect physical and digital world also overcoming heterogeneity and evanescence of mobile ad-hoc contexts. *Radio Frequency IDentification* (RFID) is the most widespread technology for product/object identification and tracking, but also for monitoring human-objects interaction [32]. *Wireless sensor networks* (WSNs) monitor environmental parameters, supporting both queries and automatic alerts triggered by application-defined events [5]. Both technologies are characterized by the dissemination of unobtrusive, inexpensive and disposable micro-devices in a given environment. Due to space, power and cost constraints, they usually have very low storage, little or no processing capabilities and short-range, low-throughput wireless links. Each mobile host in the area can access information only on micro-devices in its communication range. Consequently, approaches based on centralized control and information storage are utterly impractical. The goal of the framework proposed here is to allow objects, places, events and phenomena to be easily and thoroughly described by means of semantic annotations stored within an associated micro-device. Information about an object becomes structured and complete; it can follow accurately the object history, being progressively built or updated during its life cycle. Tangible benefits can be so provided *e.g.*, to the management of supply chains and of the life cycle of industrial products, by improving the traceability of production and distribution. Manufacturing and quality control can exploit accurate descriptions of raw materials, components and processes; supply chain management can benefit from improved item tracking; the verification of multi-factor service level agreements between commercial partners can be automated. Furthermore, sale depots could obtain easier inventory management and could provide *u-commerce* (ubiquitous commerce) capabilities [48]. Furthermore, smart post-sale services can be provided to purchasers, by integrating knowledge discovery and reasoning capabilities in home and office appliances [35]. In addition, asset management is greatly improved in those scenarios where retrieval should be based on relevant object properties and purposes, rather than mere identification codes. In healthcare applications, equipment, drugs and patients can be thoroughly and formally described and tracked, not only to ensure that appropriate treatments are given, but also to provide decision support in therapy management. This is an evident improvement w.r.t. infrastructures lacking support for formal semantics, such as [53]. Likewise, in museums, libraries and archaeological sites, smart semantic-based content fruition can be granted to local visitors as well as to remote clients connected through the Web, leveraging the lightweight infrastructure already deployed for internal inventory and research. Finally, Wireless Semantic Sensor Networks are an emerging and challenging paradigm. Advanced solutions can be built for environmental monitoring, precision agriculture and disaster recovery, by means of semantic sensory data dissemination and resource discovery features that are provided by the framework.

To clarify the rationale behind the proposed approach and its benefits with respect to classical resource discovery paradigms, this paper adopts a disaster recovery mission planning case study as a running example. When a disaster occurs in urban settings, hostile environments can make *search and rescue* life-threatening even for experienced human operators. Robot teleoperation has several issues as well: environmental conditions (*e.g.*, low visibility) may make piloting unfeasible and continuous communication requires robust connectivity, which is often unavailable in a calamity. For this reason a team of autonomous robots is dispatched as a wireless sensor and actor network with support for service/resource discovery, comprising both ground and aerial units [12]. The team is coordinated by a robotic supervisor in contact with the control station where human operators decide mission goals and monitor progress.

3 RELATED WORK

Building the SWoT requires decentralized and collaborative middleware, suitable for pervasive and ubiquitous computing. Several proposals for semantic mobile middleware infrastructures can be found in literature. Earlier solutions inherited designs from common stable networked infrastructures, relying on centralized brokers for management and discovery of devices, services and information [7, 28, 44, 14]. The infrastructure in [47] for ubiquitous semantic-based resource discovery is related to the proposal presented here due to the devised decentralized collaborative paradigm, but it was based on a direct reuse of traditional Semantic Web technologies, particularly SPARQL [43] language for queries and HTTP (HyperText Transfer Protocol) for resource transfer. Significant performance overheads and high complexity of mobile node implementation are open issues when adopting protocols not optimized for pervasive computing environments.

More recent service-oriented architectures (SOA) for the IoT exploited Semantic Web technologies to annotate data, devices and services through standard vocabularies such as the Semantic Sensor Network ontology of the World Wide Web Consortium [9] and share produced information as Linked (Open) Data (LOD) [16]. In order to facilitate interoperability and solution design in SWoT contexts, ontology catalogs have been receiving increasing focus: the *Linked Open Vocabularies for Internet of Things* (*LOV4IoT:* `https://lov4iot.appspot.com/`) repository is among the most comprehensive and well-structured efforts [15]. *Sense2Web* [4], *Linked Stream Middleware* (LSM) [24] and *Ztreamy* [13] are LOD platforms to publish sensor data, combine them and link them to existing resources on the Semantic Web. In [45] KR was exploited to support automatic sensor composition: functional and non-functional properties of sensors, as well as users' goals, were described w.r.t. an OWL ontology and an orchestrator combined sensors and processes to satisfy a request. The proposal, however, did not outline a complete system for IoT service discovery, focusing only on sensor description and composition methods. In [30], requests expressed in terms of device characteristics undergo a matchmaking process with ontology-based sensor descriptions: quantitative attributes and semantic-based reasoning are combined to improve the discovery and select appropriate sensors through an exploratory search. The user-centric task-oriented IoT service framework in [23] adopts an ontology-based model of activities, user profiles and context. Elementary services are pre-filtered w.r.t. context and then composed in processes to satisfy user requests; processes are finally ranked via quality of service (QoS) estimation. The approach, however, is inherently centralized (requiring an on-premise or cloud-based support infrastructure), possibly limiting performance and manageability in large-scale IoT applications. Moreover, the role of semantics appears quite limited, with support for exact match of functional service attributes only. The framework in [51] adopts a probabilistic approach instead, building user-user, thing-thing and user-thing correlation models to support recommendation of the most interesting things to interact with for a given user and context.

The above proposals require a dependable Internet connection and/or a support infrastructure to enable service/resource discovery features. This impairs their application to ad-hoc networks of resource-constrained things. Peer-to-peer overlay network techniques are used to enable large-scale discovery within heterogeneous networks. In such cases, most approaches adopt only a resource name resolution scheme, not allowing to use articulate features for discovery, selection and ranking.

In order to improve upon the state of the art, the present paper shares core ideas with *information-centric networking* (*a.k.a. named data networking, NDN*) approaches for general-purpose computer networks. Basically, those approaches move to the network level some basic functions of Distributed Hash Tables (DHT), used in overlay networks (*e.g.*, for file sharing) or of IP-based resource discovery protocols such as Service Location Protocol. Among other proposals, Content-Centric Networking (CCN) [20] stood out for its elegant, practical approach for progressive transition of the Internet toward the CCN model. In CCN, only two packet types exist: *Interest* and *Data*. They are characterized only by a *content name* field, entirely replacing source/destination host address fields. A node, interested in a data chunk, propagates an Interest packet in the network; a Data packet is transmitted only in response to an Interest (following the reverse route) and "consumes" that Interest. Further notable features are end-to-end security and the exploitation of the broadcast property of wireless channels, which makes the protocol suitable to wireless ad-hoc networks. Lee *et al.* [25] analyzed benefits of content-centric frameworks from the energy efficiency standpoint, a relevant aspect for mobile and pervasive computing. Nevertheless, as pointed out in [46], some open issues remain. Primarily, the unique name of each Data chunk is a binary string, so that only bitwise longest-prefix matching can be used for resource discovery. Furthermore, defining such structured names relies on application-specific shared policies, which appears as a limitation to broad interoperability. Finally, experimental results in [20] and in subsequent work [19] concerned only networks with very few nodes. Approaches proposed in latest years explored different design decisions concerning architecture, caching, mobility and security in larger experimental settings, highlighting their benefits and issues [50, 10]. In particular, scalability cannot be neglected, as experiments reported in Section 5 show. Content caching is the primary mechanism to reduce network load, improve availability and increase network scale [52]. The IoT exacerbates scalability and mobility issues significantly w.r.t. conventional computer networks, requiring novel approaches [1].

Efficient peer-to-peer data dissemination is a key technological issue to make information-centric paradigms successful. In latest years, *epidemic protocols*, *a.k.a.* gossip protocols [21], have received significant interest, as they require no network configuration and provide a good trade-off between algorithmic complexity and performance guarantees [22, 49]. They are also robust against high communication failure rates and data loss events, a crucial aspect in IoT scenarios [26].

Table 1 summarizes the main features of the most relevant related works w.r.t. the proposed approach. The proposed approach is the only one combining a distributed, peer-to-peer architecture for information dissemination and discovery based on low-overhead transport protocols with semantic-based service discovery and composition.

Table 1: Comparative table of information-centric IoT service discovery approaches

Reference	Architecture	Service discovery	Service registration / advertisement	Transport protocol	Notes
GSD [7] – 2006	Distributed, P2P	Hybrid (services group IDs and ontology-based service descriptions)	Controlled broadcast dissemination	UDP	
FIUF [28]– 2007	Distributed, client/server	Syntactic (XML)	Registration to zone directory service (“resource index” node)	UDP	Multi-agent system
mRDP [47] – 2007	Distributed, client/server	Semantic (SPARQL-equivalent query language over RDF)	None (request multicast only)	UDP for requests, TCP (HTTP) for replies	
AIDAS [44] – 2008	Centralized	Semantic (matching of OWL/RDF descriptions)	Registration to zone directory service (“context manager”)	UDP	Ranking based on preference priority
CCN [20] – 2009	Distributed, client/server	Syntactic (longest-prefix string matching)	None (request broadcast and reply caching only)	UDP	Integrability into current IP routers
SOCRADES [14] – 2010	Distributed, client/server	Syntactic (keyword-based)	WS-Discovery (multicast)	UDP	RESTful service interfaces
Sense2Web [4] – 2010	Centralized (web application)	Semantic (SPARQL on RDF descriptions)	Linked Data	TCP (HTTP)	Spatial, temporal and ontology-based service attributes
Colitti *et al.* [8] – 2011	Centralized (gateway)	Syntactic (CoAP)	None (resource changes sent in unicast in observe mode)	UDP (CoAP)	HTTP-CoAP bridging
LSM [24] – 2012	Centralized platform (Web application)	Semantic (SPARQL on RDF descriptions)	Registration to platform	TCP (HTTP)	Both data sources and continuous queries supported as services
NetInf [10] – 2013	Distributed, P2P	Syntactic (name-based, hierarchical)	Publish to name resolution server	Agnostic (TCP, UDP or other)	Embedded PKI-based security
Ztreamy [13] – 2014	Distributed, client/server	Hybrid (URI prefix, SPARQL on RDF descriptions, custom filters)	Publish to server	TCP (HTTP)	Data stream oriented
CASSARAM [30] – 2014	Distributed, client/server	Hybrid (SPARQL queries, relational filters on context attributes)	Unspecified	TCP (HTTP)	Attribute weights based on user priorities
SoIoT [23] – 2016	Centralized platform	Hybrid (ontology-based goal and context model, quantitative QoS attributes)	Register with the platform	TCP (HTTP)	Semantic service composition
Yao *et al.* [51] – 2016	Centralized (Web application)	None (probabilistic proactive recommendation)	Web-based interface	TCP (HTTP)	Probabilistic service-context correlation, RESTful Web services
This work	Distributed, P2P	*Hybrid (context attributes and semantic matchmaking of OWL descriptions)*	*Controlled broadcast dissemination*	*UDP*	Semantic service composition

4 FRAMEWORK

Figure 1 shows the conceptualization of the proposed knowledge dissemination and discovery framework. IP is leveraged for basic addressing and routing in local (typically wireless and ad-hoc) networks and internetworking between autonomous networks (including wide area networks and the Internet). The *semantic information-centric internetworking* layer grants common access to information provided by semantic-enhanced embedded devices and sensors populating a smart environment.

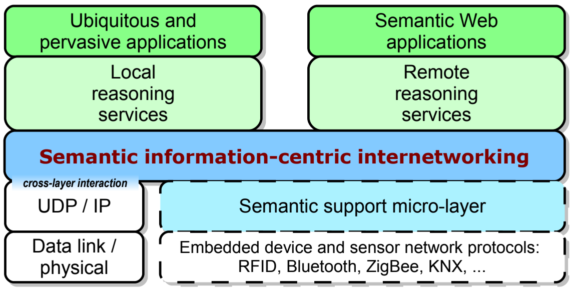

Figure 1: Semantic information-centric framework

The framework provides interoperability with mobile ad-hoc networks protocols of embedded devices and sensors to allow the extraction of information resources from the environment. In order to support both annotated information exchange and backward compatibility towards current application-level protocols and interactions, each mobile identification and sensing technology requires a *semantic support micro-layer* for adapting in the framework. Previous research work introduced semantic support to widespread technologies such as Bluetooth [33], EPCglobal RFID [11], ZigBee [36], KNX protocol for home and building automation [34], Constrained Application Protocol (CoAP) [39] and

Eddystone for Bluetooth Low Energy (BLE) beacons [38]. Further embedded device domains could be integrated into the framework in a similar fashion.

Applications use the content-centric network to discover resources/services, upon which they can execute logic-based queries and reasoning services. In MANET environments, mobile hosts with embedded reasoning capabilities [37] enable ubiquitous and pervasive semantic-aware applications. Furthermore, by means of the same protocol primitives, a gateway node can expose semantic annotations towards remote hosts as well as it can forward remote requests inside the local network. In this way, Semantic Web applications, powered by traditional query and inference engines for the Web, are allowed access, reporting and monitoring capabilities, so enabling the integration of pervasive information into Linked Data on the Web [16].

With respect to the search and rescue case study, a reference example can be considered to understand the proposed framework. *An explosion has occurred in a campus laboratory. Environmental hazards may include fire, toxic gases, debris, liquid pools and narrow tunnels. A robot rescue team is coordinated by a mobile headquarter deployed near the disaster area. Communication is provided by an ad-hoc IEEE 802.11 network.* Robot units are highly heterogeneous w.r.t. movement mechanisms, provided capabilities/resources and functional requirements. The team must be able to self-coordinate in the field, interacting with human controllers only for general mission directives. A team leader robot will interface with the nearby headquarter and orchestrate features of team members dynamically. It will use semantic-based discovery upon wireless IEEE 802.11 connectivity to select the best team configuration for achieving the mission goal.

In the following subsections, the overall architecture is explained, as well as the dissemination and discovery interaction stages. Details about relevant protocol primitives are provided.

4.1 Architecture

In KR approaches adopted by the Semantic Web, two kinds of knowledge are modeled:
– *conceptual* knowledge, or general knowledge about the problem domain;
– *factual* knowledge, which is specific to a particular problem.

Conceptual knowledge is represented in the form of an *ontology*, describing general properties of concepts and relationships among them. Factual knowledge is specific to the individuals in the domain of discourse: current Semantic Web technologies such as RDF allow to describe individual resources and their existing relationships. An ontology and a set of asserted facts form a Knowledge Base (KB) from which further entailed knowledge can be derived.

In the proposed approach, the classical KB (intended as a fixed and centralized component) evolves toward a *ubiquitous* entity: ontology files can be managed by one or more hosts, while individual resources are scattered within a smart environment, because they are physically tied to micro devices deployed in the field. For example, in RFID-based scenarios, each individual resource is a semantically annotated object/product description, stored within its RFID tag. Since several object classes, described w.r.t. different ontologies, can co-exist within a physical environment, they share the system infrastructure. Nevertheless, each individual resource annotation refers to one ontology providing the conceptual knowledge for the particular domain. Ontology Universally Unique Identifier (OUUID) codes [33] are adopted to mark ontologies unambiguously and to associate each individual to the ontology w.r.t. which it is described. In the Semantic Web technological stack, URIs are used to identify ontologies; furthermore, according to Linked Data best practices, URIs should be used as URLs (Uniform Resource Locators) pointing to the actual ontology document via HTTP. Nevertheless, in this framework OUUIDs are preferred because:
– URIs have variable length and are generally much longer than OUUIDs. That introduces overhead in network protocol fields, particularly in bandwidth-constrained mobile ad-hoc networks targeted by our framework.
– OUUID is easily mapped to data types for resource class identifiers adopted by most standard mobile discovery protocols, e.g., UUID in Bluetooth Service Discovery Protocol [33] and data types in messages exchanged by ZigBee application objects [36].
– URIs can be used to locate ontologies only as long as Internet connectivity is available. In many of the pervasive and ubiquitous contexts listed in Section 2, mobile devices cannot be connected to the Internet due to cost, power and environmental constraints. In all those cases, using URIs rather than OUUIDs provides no practical advantage while introducing overhead.
– Whenever Internet connection is available, ontology access is still granted by means of OUUID-to-URL mapping mechanisms. In [11], the ONS (Object Naming Service) facility of the EPCglobal technological stack for RFID, based on standard DNS (Domain Naming Service) protocol, was exploited for this purpose. In pervasive contexts not based on RFID, similar decoding mechanisms based on DNS or on HTTP redirection - already widely adopted on the Web, *e.g.*, by popular URL *shortening* services- can be trivially integrated into the architecture in the same way. OUUID-to-URL

mapping also allows to handle the case of composed ontologies, consisting of several modules connected by means of *import* relations. Hence, without loss of generality, hereafter it can be assumed that each resource refers to only one OUUID (which is the one of the "master" ontology in case of imports). Resources belonging to the same domain will likely be described by means of the same ontology, while objects of different categories may refer to different ontologies. In detail, in the proposed framework each resource is featured by:
– 96-bit ID, globally unique item identifier (*e.g.*, the 96-bit EPC code for an RFID tag or the 64-bit MAC address -padded to fit the space and to allow different MAC protocols to be distinguished- for a ZigBee sensor);
– 64-bit OUUID;
– a set of data-oriented attributes, which allow to integrate and extend logic-based reasoning services with application-specific and context-aware information processing;
– semantic annotation, stored as a compressed OWL document fragment.

For the compression of OWL annotations, a *homomorphic* encoding scheme for XML-based documents introduced in [41] is adopted. Homomorphism preserves XML document structure during compression, so enabling query processing directly on encoded annotations, without requiring preliminary decompression. Details about the compression format and algorithm are not reported here, in order to keep the paper short and focused; the reader is referred to the above-mentioned work for a thorough coverage.

In the proposed framework, the overall network can be seen as a two-level infrastructure, as depicted in Figure 2. Pervasive identification and sensing technologies are exploited at the *field layer* (interconnecting embedded micro-devices dipped in the environment and hosts able to receive the transmitted data) whereas the *discovery layer* is related to the inter-host communication. Each network host –marked as (1) in Figure 2– acts as a cluster head (CH) for field devices in its direct range (2), using available communication interfaces. Resources acquired at the field layer through different protocols are exposed at the discovery layer in a uniform fashion (3), according to the structure described above. Interaction among hosts is performed by means of the IP-based information dissemination and resource retrieval framework outlined hereafter. If two nodes are in wireless range or share the same wired transmission medium, the interaction between them can happen directly (4), otherwise it is necessary to adopt a multi-hop routing (5) by exploiting other nodes as intermediate links (6). However, a node does not depend on some other ones to advertise/register object descriptions.

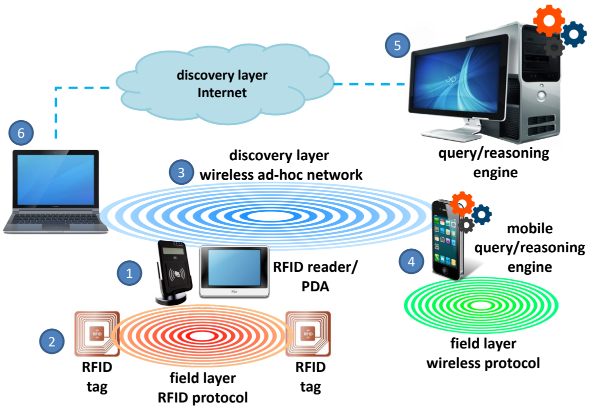

Figure 2: Field and Discovery layers in the proposed framework architecture

Resources are autonomously acquired from the field layer and exposed. At the same time, nodes are able to discover them thanks to the preliminary propagation of data each cluster head has seen in its range. In short, the information-centric framework is based on four interaction stages:
1. extraction of resource parameters (for carrying object characteristics from field layer to discovery one);
2. resource information dissemination (to make nearby nodes fully aware of the "network content");
3. resource discovery based on a peer-to-peer collaborative protocol (see later on);
4. extraction of selected resource annotations (for carrying semantic-based descriptions from field level to the discovery one) to allow semantic-based queries and reasoning.

It should be pointed out that the proposed approach is fully decentralized. Address and main characteristics of each resource/object are autonomously advertised by the related cluster head, using small-sized messages throughout the network. Care has to be taken in the use of broadcasting mechanism to advertise object features, because an uncontrolled flooding could become largely inefficient in terms of bandwidth usage and power consumption (both fundamental and precious resources in pervasive environments). In the proposed approach, only resource parameters are advertised in broadcast throughout the network in order to unambiguously identify both the location and the category of a resource. Then, if a node needs a specific resource, it will send an explicit unicast request for its semantic annotation. Before starting whatever RDF query or reasoning task, a requester has to recompose also the ontology in a consistent fashion. For this purpose, typical techniques of hybrid peer-to-peer file sharing are exploited. In particular, the OUUID of the reference ontology will be used to retrieve ontology chunks from network hosts. When all the chunks become available, the ontology is ready for use. Summarizing, the discovery procedure occurs in two steps: the first one is syntax-based, the second one is semantic-based. The first phase aims to select resource descriptions potentially interesting for the requester via the OUUID matching and the contextual parameters evaluation. The second one aims to select the best available resources. In this second stage the requester directly queries semantic annotations of resources from the provider, so preparing the further reasoning. Semantic matchmaking allows ranking of service/resource descriptions w.r.t. a request as concept expressions w.r.t. a common ontology. The proposed approach leverages two nonstandard inference services to provide fine-grained resource classification, logic-based ranking and outcome explanation (see [37] for algorithms and further details):

– *Concept Contraction*: if a request $\mathcal{R}$ and a supplied resource $\mathcal{S}$ are in contrast, Contraction determines which part of $\mathcal{R}$ is conflicting with $\mathcal{S}$. By retracting this part, denoted as G (for *Give up*), a concept expression K (for *Keep*) is obtained, representing a conflict-free (contracted) version of the original request;

– *Concept Abduction*: if request and resource are not in contrast, but $\mathcal{S}$ does not completely satisfy $\mathcal{R}$, Abduction determines what should be hypothesized in $\mathcal{S}$ in order to obtain a full match. The solution H (for *Hypothesis*) to Abduction can be interpreted as what is requested in $\mathcal{R}$ and not specified in $\mathcal{S}$.

Both Contraction and Abduction have associated *penalty functions* to measure the semantic distance of $\mathcal{S}$ from $\mathcal{R}$. This enables a logic-based relevance ranking of a set of resources w.r.t. a request;

– *Concept Covering*: in IoT scenarios it is often useful to aggregate several low-complexity devices and services in order to satisfy a specific request. The Abduction-based *Concept Covering Problem* (CCoP) aims to cover (*i.e.*, satisfy) constraints expressed in a request as much as possible, by means of the conjunction of service instances, and to provide explanation of the possible uncovered part $\mathcal{H}$ of the request.

The proposed hybrid, on-demand approach has been chosen considering that semantic descriptions are needed only in the last discovery phase, whereas the preliminary ontology-based selection procedure tackles ontology-related interoperability issues. In this way, a significant reduction of the induced traffic can be obtained and hardware/software node complexity can be decreased with respect to other semantic-based pervasive computing proposals. On the other hand, the full power of semantic-based annotation and discovery is exploited in the second step.

In the search and rescue case study, the mission goal is the request for the resource discovery algorithm, which has to find the most suitable composition of services/resources provided by team units. Each robot is characterized by a set of operational *preconditions* that must be satisfied and a set of produced *effects*, which may comply preconditions for further team units. A preliminary procedure is performed to select devices managing the same "resource classes" using OUUIDs. Figure 3 shows the chunks of the OWL ontology that has been modeled to express relevant concepts and properties for the domain, grounding semantics and inference tasks on the *Attributive Language with unqualified Number restrictions* ($\mathcal{ALN}$) logical language of the *Description Logics* (DL) [2] family.

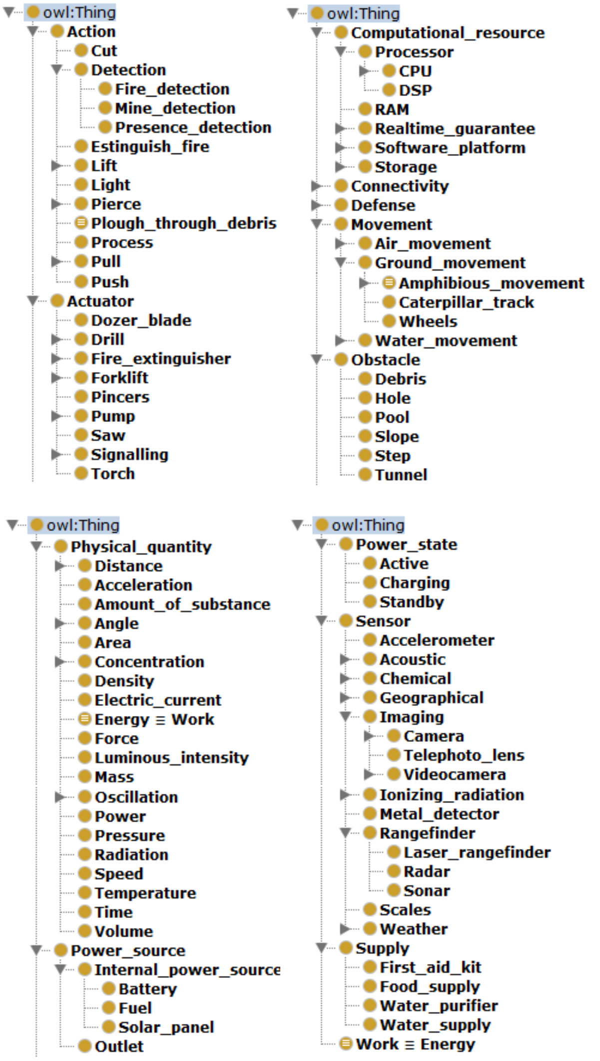

Figure 3: Disaster recovery KB class hierarchy created for the case study

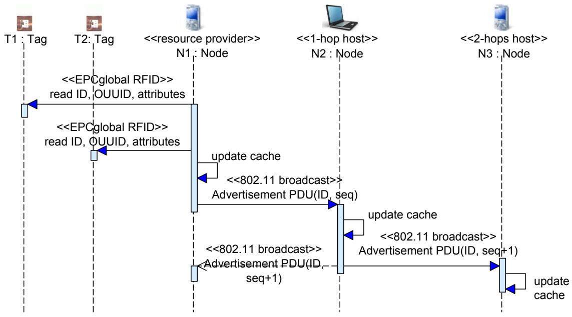

Figure 4: Information dissemination

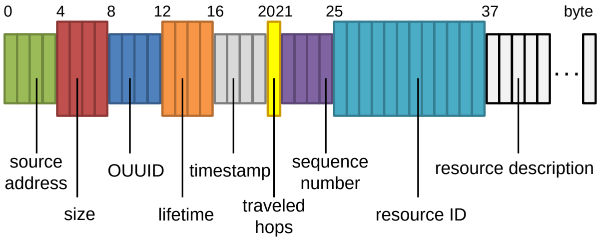

Figure 5: Structure of a cache record

4.2 Information Dissemination

An efficient dissemination protocol is fundamental to balance network resource usage and ease of resource retrieval. Resource providers periodically send ***Advertisement PDUs*** (Protocol Data Units) also specifying the maximum number of hops that the advertisement must travel (`MAX_ADV_DIAMETER`). During such lifetime, advertisements are forwarded using data-link-layer broadcasts and can be stored in the cache memories of intermediate nodes. Figure 4 shows the typical sequence of the information dissemination phase and involved actors. In the depicted example, embedded micro-devices consist in EPCglobal RFID tags storing semantically annotated object descriptions, while network nodes are PDAs (personal digital assistants) equipped with RFID reader and IEEE 802.11 transceiver.

All resources, detected at field layer, are advertised by means of a single advertisement PDU. Hence, the size of the frame increases proportionally with the number of resources in the network host range. PDU fields are explained in what follows.

– `TYPE`: the kind of PDU (see Table 2).

– `FLAGS`: it contains one status flag to distinguish the kind of transmission (unicast or broadcast); the remaining flags are reserved for future purposes.

– `TRAVELED HOPS`: the number of hops already traversed by the frame. A node sets this value to 1 and it is increased every time a node forwards the frame.

Table 2: PDU types

Type	Bit set	PDU
A	0	Advertisement
B	1	Cache entry
C	2	Solicit
D	3	Request
E ... L	4 ... 7	Reserved

– `NUMBER OF RESOURCES`: how many micro-devices are in the Cluster Head (CH) radio range.

– `ADVERTISEMENT ID`: the CH's sequence number.

– `NODE SEQUENCE NUMBER`: the sequence number of the node forwarding the frame. If the frame has been sent by a node, this value coincides with the previous one.

– `SOURCE ADDRESS`: the IP address of the CH.

– `RESOURCE PARAMETERS`: a composite, variable-length field depending on the number of advertised resources. For each resource it contains: OUUID value, remaining life time, maximum hops number for the advertisement travel, and resource ID.

A network node, which has detected one or more micro-devices in its direct range, broadcasts an advertisement every `DEFAULT_RTIME` milliseconds (values of exploited constants are reported in Table 3). Nearby nodes forward the frame by broadcasting it to their neighbors; hence, the reader listens to the echo of the advertisement PDU it originally transmitted. Thus, it can obtain a confirmation of the presence of other nodes in its neighborhood and update its routing table. If the reader does not receive any echo within `POLLING_TIME` milliseconds (less than `DEFAULT_RTIME`), it will retransmit the advertisement, presuming that a collision or a transmission error has occurred. After `MAX_RETRIES` retries, it can be presumed there are no neighbors, so the transmission of the advertisement can be scheduled after a longer timeout in order to reduce power consumption.

When a node receives an advertisement, it extracts information about the resources and, in case of "new" elements, it adds cache entries; otherwise, before updating stored data, the node verifies if the received information is more recent or has ran across a shorter path than the existing one. In particular, the entry is updated if its `ADVERTISEMENT_ID` is smaller or its `TRAVELED_HOPS` field is greater or equal w.r.t. the received frame. This means that the arrived PDU is more recent or has ran across a shorter path.

If the cache is updated and the maximum advertisement diameter has not been reached, then the advertisement is forwarded; otherwise the whole frame is silently discarded. This simple mechanism ensures each node in the network sends the same advertisement at most once. Furthermore, each host

Table 3: Protocol constant parameter values

Name	Meaning	Value
`DEFAULT_RTIME`	Time interval between two subsequent advertisement frame transmissions	30000 ms
`POLLING_TIME`	Time a network node waits for the echo of the advertisements	7500 ms
`MAX_ADV_JITTER`	Maximum value for random time waited when advertisement frames are forwarded	600 ms
`ONE_HOP_WAIT`	Timer set by a requester node after sending a solicit PDU, waiting for cache contents reception	2000 ms
`HOP_TIME`	Time a node needs to process and forward a solicit PDU sent by a neighbor	50 ms
`ACK_RTT`	Timer set by a requester node waiting for acknowledgment after a solicit has been sent	50 ms
`DISCOVERY_DIAMETER`	Current search diameter (in hops) during discovery phase	4
`MAX_RETRIES`	Maximum number of retransmissions before a reader presumes there are no neighbors	5

Figure 6: Example of *Wheels*, *Laser_rangefinder* and *Fire_extinguisher* resource descriptions

waits a random time $t \in [0,$ `MAX_ADV_JITTER`$]$ before transmitting. This aims to reduce collision probability when using data link layer protocols that do not provide acknowledgment for broadcast transmissions, such as IEEE 802.11 MAC protocol. After this check, the node verifies if the maximum advertisement diameter has been reached and builds the new advertisement PDU to be forwarded to its neighbors, according to the following procedure:

1. The sequence number of the node is written into the `NODE_SEQUENCE_NUMBER` field of the new frame.
2. For each advertised service, the `TRAVELED_HOPS` value is compared against the `MAX_HOPS` found in the received PDU.
3. If `TRAVELED_HOPS` equals `MAX_HOPS` the information is discarded, because the maximum advertisement diameter for that service has been reached.
4. Otherwise the `TRAVELED_HOPS` value is increased by 1 and the service information is included into the new frame. In this case, both `TRAVELED_HOPS` and `NODE_SEQUENCE_NUMBER` are increased.

Each reader manages a cache table where it stores characteristics of both resources in its radio range and resources it has "seen" in the network. Figure 5 shows the structure of a typical entry. Content of each field is as follows.

– `Source address`: the address of the resource provider.
– `Size`: the description size (in bytes).
– `OUUID`: numeric identifier for the specific ontology.
– `Lifetime`: remaining time of a resource/tag.
– `Timestamp`: last reference to the entry (read/write).
– `Traveled hops`: distance (hops number) between provider and cache holder.
– `Sequence number`: it is referred to the last resource provider.
– `ID`: the unique code of a resource.
– `Resource description`: the annotation of a resource.

An entry is added to the cache whenever the node receives an advertisement or a cache content frame.

In the search and rescue case study, each robot advertises its own description annotated in the $\mathcal{ALN}$ DL. Descriptions focus on provided and required capabilities rather than technically-oriented specifications of robot components derived from datasheets, since a service-oriented approach is required for mission-driven orchestration. Modeled capabilities include on-board sensing and acting devices, supported wireless communication protocols and the kinds of terrain and obstacles that can be passed. OWL *object properties* (*a.k.a. roles*) are exploited to relate pairs of classes in service instance descriptions. For example, as shown in Figure 6, the *overcomes* property relates a `Movement` type with the types of `Obstacles` it can surmount, while *measures* property links a type of `Sensor` with the `Physical_quantity` it evaluates and *performs* links an `Actuator` type with the `Actions` it can do.

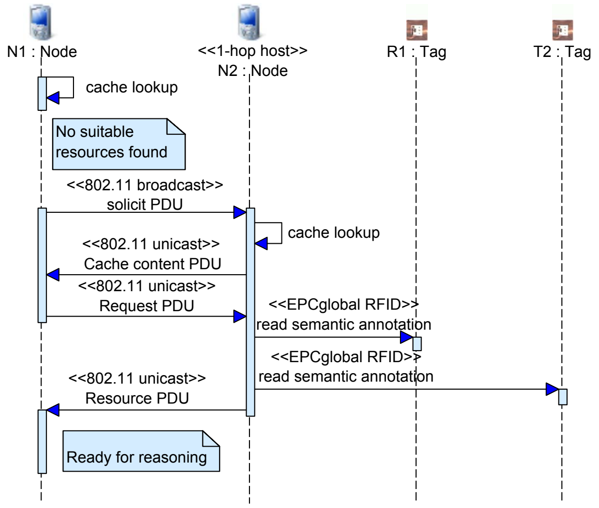

Figure 7: Semantic-based resource discovery

4.3 Service/Resource discovery

When starting a resource retrieval process, a node generally attempts to cover the request by using resource descriptions stored within its own cache memory. If some description is missing, a ***Request PDU*** is sent in unicast. On the contrary, if a requester has no resource descriptions in its cache or if managed resources are not enough to satisfy the request, the node broadcasts a ***Solicit PDU*** including the maximum travel diameter (`TOTAL_HOPS`). When receiving a solicit, a node replies (in unicast) providing cache table entries matching parameters contained within the solicit frame. If it does not manage any information satisfying the solicit, it will reply with a “no matches” message. During their travel, replies to the request and solicit PDUs are used to update the cache memory of forwarding nodes. Figure 7 depicts discovery procedure in case of RFID tag/reader interaction. A node must firstly look within its local cache table for entries compatible with the request. In case of a cache hit, it will require the corresponding annotated descriptions from their respective owners, by means of unicast ***Request PDUs***. In what follows the meaning of introduced PDU fields is summarized.

– `TYPE`: it is set to 3.
– `FLAGS`: analogous to the corresponding field of the advertisement PDU.
– `TRAVELED HOPS`: how many hops the frame has already gone across.
– `LAST HOP SEQUENCE NUMBER`: the sequence number of the last node processing the request.
– `DESTINATION SEQUENCE NUMBER`: the sequence number of the destination node.
– `DESTINATION ADDRESS`: the address of the last node processing the request.
– `PROVIDER ADDRESS`: the destination node.
– `OUUID`: the ontology unique identifier.
– `NUMBER OF REQUESTS`: the number of requested resource descriptions.
– `DATA`: the size of this field depends on the number of requests; it contains the IDs of the tags of required descriptions. Each ID is 96 bit long, like in the advertisement PDU.

If a requester has no resource descriptions in its cache or if managed resources do not satisfy the request, it requires further descriptions. So it will transmit a ***Solicit PDU*** to nearby nodes. Basically, the soliciting mechanism is analogous to the advertising one, exploiting controlled broadcast of request in an *expanding ring* fashion. A node generating a Solicit waits for an acknowledgment from each neighbor for `ACK_RTT` seconds. In this way the requester elicits information about nearby nodes: it can exactly know the number of neighbors. Each node located within `DISCOVERY_DIAMETER` hops from the requester, after receiving a solicit PDU, replies with a ***Cache content PDU*** in unicast toward the node the solicit comes from. Nodes receiving the PDU update their own cache and recursively send back the PDU, till the original requester node receives the information. A Cache content PDU has a variable length according to the number of contained resource handles. Hence the cache update could involve more records. Other PDU fields are:

– `N`: the number of resources handles (and then cache entries) the packet transports.
– `REQUEST ID`: the identifier of the original request.
– `LAST HOP SEQUENCE NUMBER`: the sequence number of the node sending the packet.
– `DESTINATION ADDRESS`: requester IP address.

The information dissemination/discovery framework and the inference procedures are both completely general and suitable for any scenario where semantic-based service/resource discovery and composition could be required. All specific aspects of a particular case study are modeled within the knowledge base. This is clarified by means of the running example in the search and rescue scenario.

Algorithm *resourceComposer* is based on the Concept Covering inference [42]. By limiting the expressiveness of resource descriptions to $\mathcal{ALN}$, reasoning complexity is polynomial for standard and nonstandard inference services [42], thus making the approach suitable to mobile and embedded computing. The algorithm takes as inputs: a set of services (resources) S, a request D, a (possibly empty) set of initial preconditions and a reference ontology. It returns a composed service flow CS (Composite Service), possibly with the part of the request $D_{uncovered}$ which could not be covered (*i.e.*, satisfied) by available services or resources. For each

Algorithm 1: *resourceComposer* algorithm

Input: Request *D*, Set *S* of cached service descriptions, Minimum Threshold Covering Level (*MTCL*), maximum discovery range (*Total_hops_MAX*), a set of initial preconditions (P_0), and the Knowledge Base *B*.

Output: A triple $\langle CS, D_{uncovered}, Covering\_Level \rangle$; where *CS* is the composed service flow, $D_{uncovered}$ is the part of the request not covered, and *Covering_Level* is the achieved request covering level.

```
$D_{uncovered} := D$; $Total\_hops := 0$; $continue := TRUE$;
repeat
    if $Total\_hops \neq 0$ then
        Update the cache and the set S with all retrieved service descriptions;
    end
    $\langle CS, D_{uncovered}, Covering\_Level \rangle :=$ resourceComposer($D_{uncovered}$, S, $P_0$, B) ;
    if ($Covering\_Level \geq MTCL$) or ($Total\_hops = Total\_hops\_MAX$) then
        $continue := FALSE$;
    end
    $Total\_hops := Total\_hops + 1$ ;
    Broadcast a Solicit PDU with the new Total_hops value;
until $continue = TRUE$;
```

element in *S*, if preconditions are satisfied, it is added to the set *EX(CS)* of executable services. Maximal covering of *D* is then performed by testing components of *EX(CS)* one at a time; if an element contributes to cover *D*, it is added to *CS*. In the first iteration, the algorithm is run by composing resources in the cache of the requester device itself. It outputs a temporary *CS* as well as the uncovered part of the request $D_{uncovered}$. If the covering level is under a given threshold, both $D_{uncovered}$ and *CS* are stored and the requester broadcasts a Solicit PDU in an expanding ring fashion to require more resources.

For the sake of readability, *Manchester syntax* [17] –as stylized by the *Protégé* KB editor [27]– is used to represent OWL concept expressions annotating the request and resources in the following example in the search and rescue scenario. Let us suppose the mission goal in the explosion area described in Section 2 is as follows: *Move through the debris, overcoming slopes and tunnels; detect, locate and extinguish fires; detect, locate and provide first aid to people.* This request could be expressed in OWL 2 as:

Request: *D* = { measures **some** (owl:Thing) **and** measures **only** (Distance) **and** Air_analysis **and** Imaging_sensor **and** performs **some** (owl:Thing) **and** performs **only** (Removing_debris **and** Detecting_fire **and** Extinguishing_fire **and** Detecting_presence) **and** overcomes **some** (owl:Thing) **and** overcomes **only** (Slope **and** Tunnel) **and** First_aid_kit }

Initial preconditions: P_0 = { WiFi **and** Fuel **and** Battery **and** DSP }

Preconditions and effects are referred to each available robot (mobile unit, m.u.). They are composed by each m.u. to characterize its own state and capabilities by annotating information derived from embedded devices and proximity sensors at the lowest layer in Figure 1. These field level annotations are exposed through the semantic support micro-layer in Figure 1 to the semantic information-centric internetworking layer, as explained in Section 4.2. Let us assume the following robots are available:

mu₁: *Environment unit* = ⟨ P_1, E_1 ⟩ = ⟨ (WiFi **and** Internal_power_source), (GPS **and** Altimeter **and** Anemometer **and** Barometer **and** Hygrometer **and** Thermometer **and** Air_analysis **and** Toxic_gas_analysis) ⟩

mu₂: *Weather unit* = ⟨ P_2, E_2 ⟩ = ⟨ (WiFi **and** Internal_power_source), (GPS **and** Altimeter **and** Anemometer **and** Barometer **and** Hygrometer **and** Thermometer) ⟩

mu₃: *Mine unit* = ⟨ P_3, E_3 ⟩ = ⟨ (Fuel), (Beacon **and** Loudspeaker **and** Metal_detector) ⟩

mu₄: *Fire victim detection unit* = ⟨ P_4, E_4 ⟩ = ⟨ (Battery **and** GPS **and** Imaging_sensor **and** Microphone **and** Thermometer **and** Air_Analysis **and** DSP), (Presence_detector **and** Fire_detector) ⟩

mu₅: *Rescue unit* = ⟨ P_5, E_5 ⟩ = ⟨ (Presence_detector **and** Fire_detector), (Pincers **and** Fire_extinguisher **and** First_aid_kit) ⟩

mu₆: *Scout unit* = ⟨ P_6, E_6 ⟩ = ⟨ (GPS **and** Altimeter **and** Anemometer **and** WiFi **and** Fuel), (Wheels **and** Dozer_blade **and** Forklift **and** Videocamera **and** Laser_rangefinder **and** Microphone) ⟩

In case of a supposed covering threshold of 90%, the following covering steps are accomplished by means of Algorithm 1:

a. $CS = \emptyset$
$D_{uncovered} = D$
$Covering\_Level = 0\%$

b. $EX(CS) = \{ mu_1, mu_2, mu_3 \}$
$CS = (mu_1)$
$D_{uncovered}$ = (measures **some** (owl:Thing) **and** measures **only** (Distance) **and** Imaging_sensor **and** performs **some** (owl:Thing) **and** performs **only** (Removing_debris **and** Detecting_fire **and** Extinguishing_fire **and** Detecting_presence) **and** overcomes **some** (owl:Thing) **and** overcomes **only** (Slope

`and Tunnel) and First_aid_kit)`
Covering_Level = 9.1%
Among executable resources, only mu_1 contributes to cover the request since it provides air analysis.

c. $EX(CS) = \{ mu_2, mu_3, mu_6 \}$
$CS = (mu_1, mu_6)$
$D_{uncovered}$ = `(performs some (owl:Thing) and performs only (Detecting_fire and Extinguishing_fire and Detecting_presence) and First_aid_kit)`
Covering_Level = 63.6%
mu_6 is triggered, since its required preconditions P_6 are satisfied (by the conjunction of P_0 and E_1). Also notice that mu_2 would cause an effect duplication with mu_1.

d. $EX(CS) = \{ mu_2, mu_3, mu_4 \}$
$CS = (mu_1, mu_6, mu_4)$
$D_{uncovered}$ = { `(performs some (owl:Thing) and performs only (Extinguishing_fire) and First_aid_kit)`
Covering_Level = 81.8%
mu_4 provides detection of fire sources and human presence.

e. $EX(CS) = \{ mu_2, mu_3, mu_5 \}$
$CS = (mu_1, mu_6, mu_4, mu_5)$
$D_{uncovered} = owl : Thing$
Covering_Level = 100%

5 EXPERIMENTAL RESULTS

The proposed approach has been tested using *ns-2* network simulator (`http://nsnam.sourceforge.net/wiki/index.php/User_Information`), to evaluate the effectiveness of the framework and find possible performance issues. A simulation campaign of the full protocol stack has been conducted for semantic-enhanced information-centric networking in complex MANET environments. The following performance metrics have been considered.
1. *Network load*, assessed by means of the total packets generated at discovery layer.
2. *Hit ratio*, *i.e.*, percentage of successful resource retrieval, where a hit has to be intended as the delivery of at least three complete descriptions referred to the same ontology. This reference value has been elicited from preliminary tests, using different domain ontologies. Obviously, the greater the number of available resources per node, the greater the probability of success.
3. *Duration* of resource discovery sessions.
Furthermore, the sensitivity of such properties has been assessed w.r.t. changes in network topology due to node mobility. Several scenarios have been created for simulation, each comprising 50 nodes moving in a plain 1000 m x 1000 m area. Each simulated host is equipped with an IEEE 802.11 transceiver with omnidirectional antenna, 2 Mb/s nominal bandwidth and 250 m range. Two-way ground signal propagation model [3] was adopted. Scenarios were set up with 3, 5 and 7 hosts detecting annotated resources in their direct range (hereafter *providers*) while 15, 30 and 45 hosts act as *requesters*. Resources are available at the beginning of each simulation, whereas requests are generated at randomly chosen instants, uniformly distributed within the simulation time. Moreover, for each combination of reference parameters, 8 simulations were run using different values for the seed of the ns-2 random number generator. Obtained results have been averaged in order to filter out the bias deriving from conditions of single scenarios (*e.g.*, high link breakage ratio or network partitions). The host motion follows the random waypoint model [31], which is characterized by two parameters: speed S and pause time P. The simulation starts with hosts remaining stationary for P seconds, then each host selects a random destination and moves toward it with a fixed speed, randomly chosen in the range [0, S]. After reaching the destination, the host pauses again for P seconds, then selects another destination and repeats the previous steps till the end of the simulation, which lasts 900 s. For each scenario, the value of S was set to 1 m/s while P varied from 0 (non-stop motion) to 50, 250, 500 and 900 s (no motion).

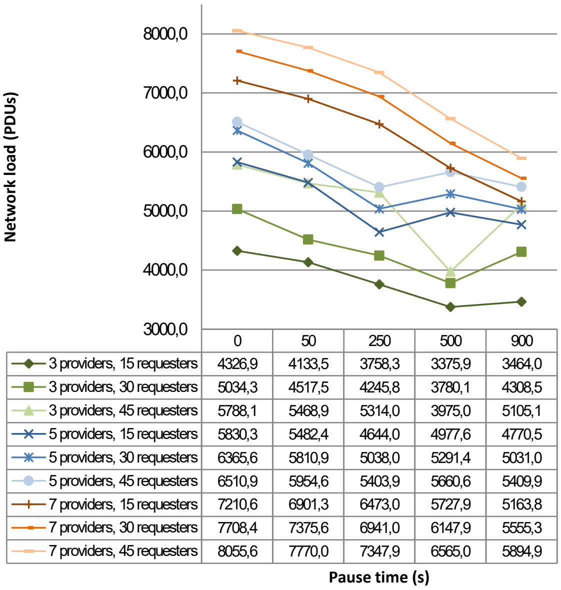

	0	50	250	500	900
3 providers, 15 requesters	4326,9	4133,5	3758,3	3375,9	3464,0
3 providers, 30 requesters	5034,3	4517,5	4245,8	3780,1	4308,5
3 providers, 45 requesters	5788,1	5468,9	5314,0	3975,0	5105,1
5 providers, 15 requesters	5830,3	5482,4	4644,0	4977,6	4770,5
5 providers, 30 requesters	6365,6	5810,9	5038,0	5291,4	5031,0
5 providers, 45 requesters	6510,9	5954,6	5403,9	5660,6	5409,9
7 providers, 15 requesters	7210,6	6901,3	6473,0	5727,9	5163,8
7 providers, 30 requesters	7708,4	7375,6	6941,0	6147,9	5555,3
7 providers, 45 requesters	8055,6	7770,0	7347,9	6565,0	5894,9

Figure 8: Network load of dissemination and discovery

In what follows, outcomes of performance evaluation are presented. Figure 8 reports on the overall network packets generated by dissemination and discovery protocols. Results show that traffic has higher correlation with the number of providers rather than requesters.

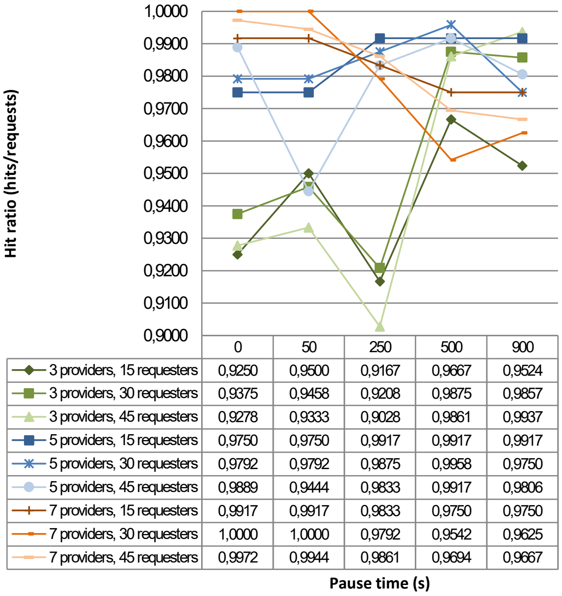

	0	50	250	500	900
3 providers, 15 requesters	0,9250	0,9500	0,9167	0,9667	0,9524
3 providers, 30 requesters	0,9375	0,9458	0,9208	0,9875	0,9857
3 providers, 45 requesters	0,9278	0,9333	0,9028	0,9861	0,9937
5 providers, 15 requesters	0,9750	0,9750	0,9917	0,9917	0,9917
5 providers, 30 requesters	0,9792	0,9792	0,9875	0,9958	0,9750
5 providers, 45 requesters	0,9889	0,9444	0,9833	0,9917	0,9806
7 providers, 15 requesters	0,9917	0,9917	0,9833	0,9750	0,9750
7 providers, 30 requesters	1,0000	1,0000	0,9792	0,9542	0,9625
7 providers, 45 requesters	0,9972	0,9944	0,9861	0,9694	0,9667

Figure 9: Hit ratio

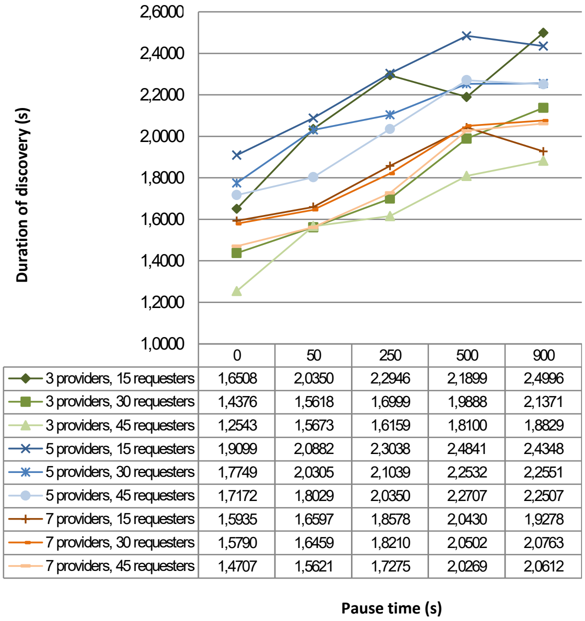

	0	50	250	500	900
3 providers, 15 requesters	1,6508	2,0350	2,2946	2,1899	2,4996
3 providers, 30 requesters	1,4376	1,5618	1,6999	1,9888	2,1371
3 providers, 45 requesters	1,2543	1,5673	1,6159	1,8100	1,8829
5 providers, 15 requesters	1,9099	2,0882	2,3038	2,4841	2,4348
5 providers, 30 requesters	1,7749	2,0305	2,1039	2,2532	2,2551
5 providers, 45 requesters	1,7172	1,8029	2,0350	2,2707	2,2507
7 providers, 15 requesters	1,5935	1,6597	1,8578	2,0430	1,9278
7 providers, 30 requesters	1,5790	1,6459	1,8210	2,0502	2,0763
7 providers, 45 requesters	1,4707	1,5621	1,7275	2,0269	2,0612

Figure 10: Average duration of discovery

This happens because advertisements are regularly sent in a proactive way by providers, even if there are no requests. On the other hand, solicit, cache content and request PDUs are produced on-demand by requesters and their neighbors. Furthermore, the number of generated packets decreases when pause time increases. As nodes remain stationary for more time, probability decreases that radio links are lost and advertisements are not echoed: hence advertisers do not schedule PDUs for retransmission, so decreasing the overall traffic. When normalizing the number of packets w.r.t. the number of nodes and the simulation duration, values range from a minimum of 0.092 to a maximum of 0.364 packets per node per second. Such data are comparable to results in Figure 14 of [7], which reported about [750, 850] packets per node in 4500 s simulation runs with the most efficient protocol variant, corresponding to [0.167, 0.189] packets per node per second. The hit ratio is depicted in Figure 9. It is very high in general, with values above 90% in all tests and above 95% in more than half the tests. Results are very close to the ones reported in Figure 15 of [7] and even to the perfect recall reported by [44], clearly indicating the relevance of the proposed approach. More precisely, hit ratio increases with the number of resource providers. This is motivated because: (i) with more providers the capillarity of dissemination is increased; (ii) the proposed discovery mechanism induces by itself a load balancing of requests, since the closest hosts with available resources will be contacted first by a requester. Finally, Figure 10 reports on the time values for a successful discovery. Analysis evidences that, for a given number of providers, the service time decreases as the number of clients increase. This is due because when a solicit is answered by cache content PDUs, intermediate nodes cache the resource records. This reduces latency in the response to later requests. Results seem significantly worse than the [0.3, 0.7] s range reported in Figure 16 of [7], but actually they are not, because in that work a hit was defined as the successful retrieval of one resource while here it means three resources. Furthermore, [47] reported that in their approach "semantic matchmaking increases exponentially as more data and ontologies are provided" and "semantic reasoning is only carried out at concrete moments of time when information or ontologies change. Once the knowledge base has been augmented via semantic reasoning, the queries can be resolved against the data repository with [linear time] lexical-like performance as long as the information remains the same". The proposed approach, on the contrary, does not incur penalties in discovery performance when the size of KB or information volatility increase. Nevertheless, overall values are still slightly high w.r.t. the requirements of pervasive computing. This may be influenced by the fact that current protocol implementation is not optimized for execution speed. Further techniques should be devised, however, in order to reduce duration of discovery sessions. Based on the above results, the proposed approach can be deemed as a step toward solving the core issues of information-centric networking in the Semantic Web of Things. Effectiveness and relevance seem evident, as well as the opportunity of further optimization.

6 CONCLUSION

This paper presented a content-centric networking approach for the Semantic Web of Things. A peer-to-peer, collaborative and dynamic framework supports information dissemination and semantic-based resource discovery. Information gathered in the field through different identification and sensing technologies is exploited at discovery layer. A thorough evaluation of the approach using *ns-2* simulator test correctness and effectiveness. An implementation of the system is ongoing with real computing devices to provide a more significant assessment of real-world performance. Future work includes: (i) a wider validation of the approach; (ii) an improvement of resource retrieval, through more advanced query and matchmaking schemes; (iii) an enhancement of ontology management, by means of semantic-aware fragmentation and on-the-fly rebuilding of TBoxes in the KB.

ACKNOWLEDGEMENTS

The authors acknowledge partial support of Italian PON project *RPASInAir - Remotely Piloted Aircraft Systems Integration in non-segregated Air space for services* and Apulia Region POR project *Sì-Ca.Re. - Integrated Monitoring and Care System for Patients with Cardiorenal Syndrome*.

REFERENCES

[1] M. Amadeo, C. Campolo, J. Quevedo, D. Corujo, A. Molinaro, A. Iera, R. L. Aguiar, and A. V. Vasilakos, "Information-centric networking for the Internet of Things: Challenges and opportunities," *IEEE Network*, vol. 30, no. 2, pp. 92–100, 2016.

[2] F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. Patel-Schneider, *The Description Logic Handbook*. Cambridge University Press, 2002.

[3] A. Bakre and B. Badrinath, "I-TCP: Indirect TCP for mobile hosts," in *International Conference on Distributed Computing Systems*, vol. 15. IEEE Computer Society Press, 1995, pp. 136–136.

[4] P. Barnaghi, M. Presser, and K. Moessner, "Publishing linked sensor data," in *Proceedings of the 3rd International Workshop on Semantic Sensor Networks (SSN), International Semantic Web Conference*, ser. CEUR Workshop Proceedings, vol. 668, 2010.

[5] L. M. Borges, F. J. Velez, and A. S. Lebres, "Survey on the characterization and classification of wireless sensor network applications," *IEEE Communications Surveys & Tutorials*, vol. 16, no. 4, pp. 1860–1890, 2014.

[6] D. Brickley and R. Guha, "RDF Schema 1.1," W3C, W3C Recommendation, Feb. 2014, https://www.w3.org/TR/rdf-schema/.

[7] D. Chakraborty, A. Joshi, Y. Yesha, and T. Finin, "Toward distributed service discovery in pervasive computing environments," *IEEE Transactions on Mobile Computing*, vol. 5, no. 2, pp. 97–112, 2006.

[8] W. Colitti, K. Steenhaut, and N. De Caro, "Integrating Wireless Sensor Networks with the Web," *Extending the Internet to Low powerand Lossy Networks (IP+ SN 2011)*, 2011.

[9] M. Compton, P. Barnaghi, L. Bermudez, R. García-Castro, O. Corcho *et al.*, "The SSN ontology of the W3C semantic sensor network incubator group," *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 17, pp. 25–32, Dec. 2012.

[10] C. Dannewitz, D. Kutscher, B. Ohlman, S. Farrell, B. Ahlgren, and H. Karl, "Network of Information (NetInf) – An information-centric networking architecture," *Computer Communications*, vol. 36, no. 7, pp. 721–735, 2013.

[11] R. De Virgilio, E. Di Sciascio, M. Ruta, F. Scioscia, and R. Torlone, "Semantic-based RFID Data Management," in *Unique Radio Innovation for the 21st Century: Building Scalable and Global RFID Networks*. Springer, 2011, pp. 111–142.

[12] M. Erdelj, M. Król, and E. Natalizio, "Wireless sensor networks and multi-UAV systems for natural disaster management," *Computer Networks*, vol. 124, pp. 72–86, 2017.

[13] J. A. Fisteus, N. F. Garcia, L. S. Fernandez, and D. Fuentes-Lorenzo, "Ztreamy: A middleware for publishing semantic streams on the Web," *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 25, pp. 16–23, 2014.

[14] D. Guinard, V. Trifa, S. Karnouskos, P. Spiess, and D. Savio, "Interacting with the SOA-based Internet of Things: Discovery, Query, Selection, and On-Demand Provisioning of Web Services," *IEEE Transactions on Services Computing*, vol. 3, no. 3, pp. 223–235, 2010.

[15] A. Gyrard, A. Zimmermann, and A. Sheth, "Building IoT-based applications for smart cities: How can ontology catalogs help?" *IEEE Internet of Things Journal*, vol. 5, no. 5, pp. 3978–3990, 2018.

[16] T. Heath and C. Bizer, "Linked data: Evolving the web into a global data space," *Synthesis Lectures on the Semantic Web: Theory and Technology*, vol. 1, no. 1, pp. 1–136, 2011.

[17] M. Horridge and P. Patel-Schneider, "OWL 2 Web Ontology Language Manchester Syntax (Second Edition)," W3C, W3C Note, Dec. 2012, http://www.w3.org/TR/owl2-manchester-syntax/.

[18] C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed diffusion: A scalable and robust communication paradigm for sensor networks," in *Proceedings of the 6th Annual International Conference on Mobile Computing and Networking*. ACM, 2000, pp. 56–67.

[19] V. Jacobson, D. K. Smetters, N. H. Briggs, M. F. Plass, P. Stewart, J. D. Thornton, and R. L. Braynard, "VoCCN: Voice-over Content-centric Networks," in *Proceedings of the 2009 Workshop on Re-architecting the Internet*, ser. ReArch '09. ACM, 2009, pp. 1–6.

[20] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking Named Content," in *Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies*. ACM, 2009, pp. 1–12.

[21] M. Jelasity, A. Montresor, and O. Babaoglu, "Gossip-based aggregation in large dynamic networks," *ACM Transactions on Computer Systems (TOCS)*, vol. 23, no. 3, pp. 219–252, 2005.

[22] D. Kempe, J. Kleinberg, and A. Demers, "Spatial Gossip and Resource Location Protocols," *Journal of ACM*, vol. 51, no. 6, pp. 943–967, Nov. 2004.

[23] I.-Y. Ko, H.-G. Ko, A. J. Molina, and J.-H. Kwon, "SoIoT: Toward a user-centric IoT-based service framework," *ACM Transactions on Internet Technology (TOIT)*, vol. 16, no. 2, p. 8, 2016.

[24] D. Le-Phuoc, H. Q. Nguyen-Mau, J. X. Parreira, and M. Hauswirth, "A middleware framework for scalable management of linked streams," *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 16, pp. 42–51, Nov. 2012.

[25] U. Lee, I. Rimac, and V. Hilt, "Greening the Internet with Content-centric Networking," in *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*. ACM, 2010, pp. 179–182.

[26] A. Makhoul, C. Guyeux, M. Hakem, and J. M. Bahi, "Using an epidemiological approach to maximize data survival in the Internet of Things," *ACM Transactions on Internet Technology (TOIT)*, vol. 16, no. 1, p. 5, 2016.

[27] M. A. Musen, "The Protégé project: A look back and a look forward," *AI matters*, vol. 1, no. 4, pp. 4–12, 2015.

[28] G. Pallapa and S. Das, "Resource Discovery in Ubiquitous Health Care," in *21st International Conference on Advanced Information Networking and Applications Workshops*, vol. 2, 2007, pp. 1–6.

[29] B. Parsia, S. Rudolph, M. Krötzsch, P. Patel-Schneider, and P. Hitzler, "OWL 2 Web Ontology Language Primer (Second Edition)," W3C, W3C Recommendation, Dec. 2012, http://www.w3.org/TR/owl2-primer.

[30] C. Perera, A. Zaslavsky, C. Liu, M. Compton, P. Christen, and D. Georgakopoulos, "Sensor Search Techniques for Sensing as a Service Architecture for the Internet of Things," *IEEE Sensors Journal*, vol. 14, no. 2, pp. 406–420, 2014.

[31] G. Resta and P. Santi, "An analysis of the node spatial distribution of the random waypoint mobility model for ad hoc networks," in *Proceedings of the second ACM international workshop on Principles of mobile computing*. ACM New York, NY, USA, 2002, pp. 44–50.

[32] W. Ruan, Q. Z. Sheng, L. Yao, T. Gu, M. Ruta, and L. Shangguan, "Device-free indoor localization and tracking through human-object interactions," in *17th International Symposium on a World of Wireless, Mobile and Multimedia Networks*. IEEE, 2016, pp. 1–9.

[33] M. Ruta, T. Di Noia, E. Di Sciascio, and F. Donini, "Semantic-Enhanced Bluetooth Discovery Protocol for M-Commerce Applications," *International Journal of Web and Grid Services*, vol. 2, no. 4, pp. 424–452, 2006.

[34] M. Ruta, F. Scioscia, E. D. Sciascio, and G. Loseto, "Semantic-Based Enhancement of ISO/IEC 14543-3 EIB/KNX Standard for Building Automation," *IEEE Transactions on Industrial Informatics*, vol. 7, no. 4, pp. 731–739, Nov 2011.

[35] M. Ruta, T. Di Noia, E. Di Sciascio, F. Scioscia, and G. Piscitelli, "If objects could talk: A novel resource discovery approach for pervasive environments," *International Journal of Internet Protocol Technology (IJIPT)*, vol. 2, no. 3/4, pp. 199–217, 2007.

[36] M. Ruta, F. Scioscia, T. Di Noia, and E. Di Sciascio, "A hybrid ZigBee/Bluetooth approach to mobile semantic grids," *Computer Systems Science and Engineering*, vol. 25, no. 3, pp. 235–249, 2010.

[37] M. Ruta, F. Scioscia, E. Di Sciascio, F. Gramegna, and G. Loseto, "Mini-ME: the Mini Matchmaking Engine," in *OWL Reasoner Evaluation Workshop (ORE 2012)*, ser. CEUR Workshop Proceedings, I. Horrocks, M. Yatskevich, and E. Jimenez-Ruiz, Eds., vol. 858. CEUR-WS, 2012, pp. 52–63.

[38] M. Ruta, F. Scioscia, S. Ieva, G. Loseto, F. Gramegna, and A. Pinto, "Knowledge discovery and sharing in the IoT: the Physical Semantic Web vision," in *Proceedings of the Symposium on Applied Computing*, 2017, pp. 492–498.

[39] M. Ruta, F. Scioscia, A. Pinto, E. Di Sciascio, F. Gramegna, S. Ieva, and G. Loseto, "Resource annotation, dissemination and discovery in the Semantic Web of Things: a CoAP-based framework," in *Green Computing and Communications (GreenCom), 2013 IEEE and Internet of Things (iThings/CPSCom), IEEE International Conference on and IEEE Cyber, Physical and Social Computing*. IEEE, 2013, pp. 527–534.

[40] G. Schreiber and F. Gandon, "RDF 1.1 XML syntax," W3C, W3C Recommendation, Feb. 2014, http://www.w3.org/TR/rdf-syntax-grammar/.

[41] F. Scioscia and M. Ruta, "Building a Semantic Web of Things: issues and perspectives in information compression," in *Semantic Web Information Management (SWIM'09). In Proceedings of the 3rd IEEE International Conference on Semantic Computing (ICSC 2009)*. IEEE Computer Society, 2009, pp. 589–594.

[42] F. Scioscia, M. Ruta, G. Loseto, F. Gramegna, S. Ieva, A. Pinto, and E. Di Sciascio, "Mini-ME matchmaker and reasoner for the Semantic Web of Things," in *Innovations, Developments, and Applications of Semantic Web and Information Systems*. IGI Global, 2018, pp. 262–294.

[43] The W3C SPARQL Working Group, "SPARQL 1.1 Overview," W3C, W3C Recommendation, Mar. 2013, https://www.w3.org/TR/sparql11-overview/.

[44] A. Toninelli, A. Corradi, and R. Montanari, "Semantic-based discovery to support mobile context-aware service access," *Computer Communications*, vol. 31, no. 5, pp. 935 – 949, 2008.

[45] K.-N. Tran, M. Compton, and R. G. Jemma Wu, "Semantic Sensor Composition," in *3rd International Workshop on Semantic Sensor Networks. Proceedings of the 9th International Semantic Web Conference (ISWC 2010)*, ser. CEUR Workshop Proceedings, D. R. D. Taylor K., Ayyagari A., Ed., vol. 668. CEUR-WS, nov 2010, pp. 33–48.

[46] D. Trossen, M. Sarela, and K. Sollins, "Arguments for an Information-centric Internetworking Architecture," *SIGCOMM Comput. Commun. Rev.*, vol. 40, no. 2, pp. 26–33, 2010.

[47] J. I. Vazquez and D. L. de Ipiña, "mRDP: An HTTP-based lightweight semantic discovery protocol," *Computer Networks*, vol. 51, no. 16, pp. 4529 – 4542, 2007.

[48] G. Venkataramani and P. Iyer, "Semantics aware rfid middleware for personalized web services to retail consumers," 2008.

[49] S. Voulgaris, M. Jelasity, and M. van Steen, "A Robust and Scalable Peer-to-peer Gossiping Protocol," in *Second International Conference on Agents and Peer-to-Peer Computing*. Springer-Verlag, 2004, pp. 47–58.

[50] G. Xylomenos, C. N. Ververidis, V. A. Siris, N. Fotiou, C. Tsilopoulos, X. Vasilakos, K. V. Katsaros, and G. C. Polyzos, "A survey of information-centric networking research," *IEEE Communications Surveys & Tutorials*, vol. 16, no. 2, pp. 1024–1049, 2014.

[51] L. Yao, Q. Z. Sheng, A. H. Ngu, and X. Li, "Things of interest recommendation by leveraging heterogeneous relations in the Internet of Things," *ACM Transactions on Internet Technology (TOIT)*, vol. 16, no. 2, p. 9, 2016.

[52] M. Zhang, H. Luo, and H. Zhang, "A survey of caching mechanisms in information-centric networking," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 3, pp. 1473–1499, 2015.

[53] R. Y. Zhong, G. Q. Huang, S. Lan, Q. Dai, X. Chen, and T. Zhang, "A big data approach for logistics trajectory discovery from RFID-enabled production data," *International Journal of Production Economics*, vol. 165, pp. 260–272, 2015.

AUTHOR BIOGRAPHIES

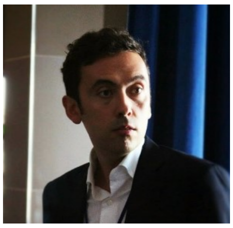

Michele Ruta received the master's degree in Electronics Engineering from Polytechnic University of Bari in 2002 and the Ph.D. in Computer Science in 2007. He is currently full professor. His research interests include pervasive computing and ubiquitous web, knowledge representation systems and applications for wireless ad-hoc contexts. On these topics, he has co-authored more than 110 papers in international journals, edited books and conferences.

Floriano Scioscia received the master's degree in Information Technology Engineering with honours from Polytechnic University of Bari in 2006, and the Ph.D. in Information Engineering in 2010 from the same institution. He is currently an assistant professor at the Information Systems Laboratory (SisInfLab) in the same university. His research interests include knowledge representation systems and applications for pervasive computing, cyber-physical systems and the Internet of Things. He co-authored over 80 papers in international journals, edited books and conferences.