Semantic-based Knowledge Dissemination and Extraction in Smart Environments

Michele Ruta[1], Floriano Scioscia[1], Eugenio Di Sciascio[1], Domenico Rotondi[2], Salvatore Piccione[2]
[1]DEE - Politecnico di Bari, Via Orabona 4, I-70125, Bari, BA, Italy
[2]TXT e-solutions S.p.A., Via Frigia 27, I-20126 Milano, MI, Italy
{m.ruta, f.scioscia, disciascio}@poliba.it, {domenico.rotondi, salvatore.piccione}@txtgroup.com

***Abstract*—The merging of Semantic Web technologies and the Internet of Things is originating the so-called *Semantic Web of Things* (SWoT), which promises to seamlessly integrate the real and digital worlds, to further enhance the available services in smart spaces and to increase the number of communicating users and objects. This paper presents a semantic content-centric framework enabling cooperative environments where resources can be discovered, queried and inventoried by autonomous objects in a peer-to-peer, collaborative way, without requiring a central control and coordination. Logic-based standard and non-standard inference services can be exploited in matchmaking and negotiation for service/resource discovery and on-the-fly reasoning, as well as for large-scale analyses. Features of the proposed approach are detailed and discussed w.r.t. the key requirements and research challenges of next-generation smart cities and factory automation.**

Keywords - Semantic Web of Things, Pervasive Computing, Smart Environments, Resource Discovery, Stream Reasoning

I. INTRODUCTION

The *Semantic Web of Things* (SWoT) is an emerging vision in Information and Communication Technology (ICT), joining the Semantic Web and the *Internet of Things* (IoT). While the Semantic Web initiative aims to allow applications to share, reuse and combine, even dynamically, information and services available in the World Wide Web, the IoT vision promotes the pervasive computing paradigm, envisaging embedded intelligence into ordinary physical objects, contexts and spaces. The goal of the SWoT is to integrate semantically rich and easily accessible information into the physical world, so connecting smart objects and digital entities. The SWoT vision has significant impact on human-computer (more generally, human-device) interaction models. It reduces the user's effort in extracting/providing contextual information and the amount of data to be provided to the system. More decentralized and autonomous interaction patterns are possible, where devices extract data from objects deployed into the environment and automatically process them in a context-aware way, in order to better support the current activity of a user and satisfy his/her needs.

Such a vision complies with content-centric (a.k.a. data-centric) network approaches [8], where information, not network nodes, grounds the protocol design. The SWoT paradigm, nevertheless, requires autonomic capabilities in information storage, retrieval and management and particularly resource discovery becomes a pivotal feature. Current mobile information-centric discovery paradigms employ elementary "string matching" of encoded resource attributes, which results unsuitable for advanced applications and complex processes where data streams produced by smart objects in multiple contexts have to be aggregated and managed.

In this paper a semantic-based content-centric framework is proposed, enabling cooperative environments where autonomous objects can be discovered, queried and inventoried without requiring a central control and coordination. The proposal includes a general architectural model bridging Mobile Ad-hoc Networks (MANETs) and a peer-to-peer distributed application-layer protocol for knowledge dissemination and discovery. Information is gathered through different identification and sensing technologies, and is exploited by inference engines and semantic-aware applications, in either pervasive or Web contexts, through a uniform set of operations. Ideas and technologies adapted from the Semantic Web vision allow to overcome the limitations of classical resource discovery and to provide robust knowledge-based stream analysis. Formal Knowledge Representation (KR) foundations and the Linked Data guidelines for information sharing [3] provide common vocabularies to describe resources and express requests. In particular, RDF (Resource Description Framework) is adopted as language for resource annotation w.r.t. RDFS (RDF Schema) domain vocabularies (ontologies). That allows semantic-based applications to leverage querying and reasoning tools, grounded on formal logics, originally conceived for the Semantic Web. The underlying logic-based framework allows standard and non-standard inference services to be exploited both (i) in matchmaking and negotiation for pervasive on-the-fly service/resource discovery and (ii) in semantic stream reasoning for large-scale information analyses.

The remainder of the paper is organized as follows. The next section surveys relevant related work, while Section III outlines motivation and possible application scenarios. Section IV describes the proposed framework focusing on its architecture and on information dissemination and resource discovery mechanisms; early experiments are also presented. Before conclusion, Section V discusses the preliminary application of the proposed theoretical model: particularly to address main technological challenges outlined in [10] for smart cities, and exploitation in a next-generation factory automation scenario.

II. RELATED WORK

The SWoT vision requires decentralized and collaborative middleware solutions, not only to support resource/service discovery and utilization, but also to characterize information with contextual meta-information. Several hierarchical DNS-like service infrastructures have been proposed to resolve object IDs. Nevertheless, they are basically designed to provide access to simple object data, lacking semantically rich and accurate information. Other solutions in literature [4] [15], although more oriented to semantic information management, rely on centralized brokers for registration and discovery of information resources, and consequently suffer from scaling issues. In [5], Christophe proposed an approach based on the Web of Things, where ontological models describe resources and their relationships, and a given framework was able to process them. Additionally, in [6] the usage of those ontological models to improve retrieval of resources in Web of Things environments was outlined. Similarly, in [17] a proposal is made for linking physical entities to online semantic information sources. Such approaches, even if distributed, still lack autonomic and local reasoning capabilities. The framework in [18] for semantic-based resource discovery is closely related with the present paper, as it also proposes a decentralized collaborative paradigm, even if based on a direct reuse of traditional Semantic Web technologies not optimized for pervasive computing environments.

All the above approaches aim to enable on-the-fly object/information discovery in pervasive computing. This is not enough for next-generation applications of pervasive computing (including supply chain management, smart cities and grids, factory automation and so on), which require also a continuous processing and mining of large-scale streaming data collected from metering and sensing devices in the field, in order to infer context or detect events and trends. *Data Stream Management Systems* (DSMSs) were deeply investigated in the last decade in order to support continuous queries over *data flows*. The majority of current DSMS solutions are based on extensions of the relational data model. They usually characterize each data record via trivial taxonomies of property categories defined on a per-application basis. This makes current DSMS unsuitable to systems that must integrate and analyze data coming from many heterogeneous sources. Hence, *Semantic stream processing* approaches have been proposed recently to infer information from RDF data streams, augmented with temporal properties. They are based either on: (i) timestamping RDF statements and extending SPARQL query language and engines to support continuous query processing [1] [2] [12], or (ii) using a SQL-like language, reformulating RDF entailment rules as SQL queries and using a tuple-based DSMS stream processor to execute them [19]. All the above solutions are limited to the simple entailment regime of RDF, therefore cannot support more advanced inference services. In [13] a framework was introduced to provide a compact representation of large concept streams and to find informative commonalities in them via non-standard inference services. It was applied to automated pattern analysis and trend discovery in supply chains, where semantic-enhanced RFID allowed storing a semantic product annotation directly on good's tags.

For what concerns infrastructure and protocols, the present paper shares core ideas with other information-centric approaches for general-purpose computer networks. In [16] current challenges of node-centric networks are discussed: at present, issues are addressed by applications on a case-by-case basis, although they could be solved in a more effective and efficient way by holistic information-centric network models. Moreover benefits of content-centric frameworks are analyzed in [11] from the energy efficiency standpoint, a relevant aspect in mobile networks. Most information-centric paradigms use route-by-name for resource discovery and routing [9]. Among them, CCN (Content-Centric Networking) [8] provides a practical approach deeply integrated within IP routers. Further notable features are end-to-end security and the exploitation of the broadcast property of wireless channels, which makes the protocol suitable to MANETs. Nevertheless some open issues remain. Primarily, the unique name of each CCN packet is a binary string with a hierarchical structure [8], so that resource discovery based on meaningful properties related to actual resource content is not possible.

III. MOTIVATION

The overall goal of the SWoT is to allow *resources* (e.g., objects, places, events, phenomena) to be easily and exhaustively described by means of semantic data stored within an associated tag or sensor. Each resource annotation is expressed using Semantic Web languages and reference vocabularies (ontologies), these latter being based on Knowledge Representation (KR) models and formalisms. Such an approach provides tangible benefits ranging from several fields and applications, like the management of supply chains or of the life cycle of industrial products, and including: accurate description of raw materials, components and processes; improved item tracking; support of u-commerce (ubiquitous commerce) and integration of knowledge discovery and reasoning into home and office appliances without expensive investments in infrastructure. Furthermore, the underlying logic-based framework could support management and analyses on large concept streams through inference procedures. A compact but semantically rich description of relevant entities and processes could be exploited for fully automated mining, pattern analysis and trend discovery, based on informative semantic commonalities [13], which is essential in complex system monitoring scenarios [10]. Anyway, the SWoT paradigm implies the following technological challenges:

1. Knowledge-based systems conceived for fixed networks are hardly adaptable to mobile and pervasive computing environments, which are characterized by dependency on context, severe resource limitations and mobility of users and devices. Approaches based on centralized control and data storage are impractical in such scenarios, where each mobile host can access information only within its wireless communication range. Pervasive computing needs a decentralized and collaborative coordination among autonomous mobile hosts.

2. A significant requirement for novel SWoT frameworks is to preserve backward compatibility with respect to

standard identification and sensing technologies, thus allowing legacy applications to co-exist with new services. Similarly, a key requirement for the practical adoption of novel information-centric networking approaches is interoperability with standard protocols at the network and transport layers [8].

3. The adoption of XML-based Semantic Web languages to describe domain vocabularies and to annotate resources requires information compression for efficient data storage and management in mobility. Extensively applied compression techniques provide overall benefits because they decrease data size, improve bandwidth utilization, and reduced energy consumption [7].

IV. Framework

Figure 1 sketches the architecture of the proposed framework. The ***ubiquitous Knowledge Base*** (***u-KB***) layer provides common access to information embedded in semantic-enhanced devices and sensors in the environment. Internet Protocol (IP) is exploited for basic addressing and routing in local networks (typically wireless and ad-hoc) and internetworking (i.e., wide area networks and the Internet).

In order to support both semantic annotation exchange and backward compatibility, each mobile identification and sensing technology requires a semantic support micro-layer for adapting it to the framework whose technical details depend on the specific protocol(s) to be adapted (see [7] for a description of the integration of Bluetooth Service Discovery Protocol and EPCglobal RFID).

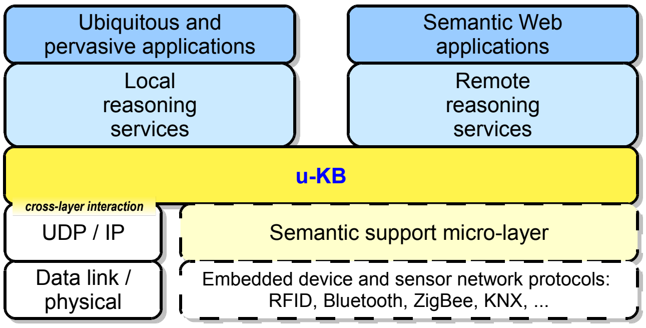

Figure 1. Ubiquitous Knowledge Base (*u-KB*) framework.

Applications use the information-centric network to disseminate and discover knowledge resources, upon which they can execute logic-based queries and reasoning services. Mobile hosts with embedded reasoning engines [14] enable ubiquitous semantic-aware applications, while gateway nodes can use the same protocol primitives to expose annotated resources towards remote hosts and to forward remote requests inside the local network, so enabling integration with the Semantic Web and Linked Data [3]. Ubiquitous and Web applications have substantial differences in terms of functional, performance, and architectural requirements. The *u-KB* layer provides common primitives for resource discovery in both cases. Indeed, local reasoning services can be exploited in pervasive computing where ad-hoc networks provide the communication infrastructure and mobile handheld devices with moderate computational and storage capabilities act as network nodes. On the other hand, remote reasoning services can be used for stream analysis of large amount of data coming from field devices dipped into an environment.

The peculiarity of the proposed *u-KB* framework is to unify the information representation and access mechanism so that, by simply selecting proper inferences, multiple goals can be pursued: different implicit information can be derived starting from available explicit object descriptions.

A. Architecture

A classic Knowledge Base is composed of an *ontology* - containing general *conceptual* knowledge about the problem domain - and a set of asserted *facts*, from which further entailed knowledge can be derived. In the proposed approach, the KB becomes *ubiquitous*: ontology segments can be managed by one or more hosts, while individual resources are scattered within a smart environment physically tied to devices deployed in the field. Since several object classes, described w.r.t. different ontologies, can co-exist in a physical space, they share the system infrastructure. Nevertheless, each individual resource annotation refers to an ontology providing the conceptual knowledge for a particular domain. Ontology Universally Unique Identifier (OUUID) codes [13] are adopted to mark ontologies unambiguously and to associate each individual to the describing vocabulary. OUUIDs are preferred to URLs, suggested in Linked Data best practices [3], because URLs are generally much longer than OUUIDs introducing overhead in bandwidth-constrained mobile ad-hoc networks targeted by our framework. Moreover OUUID is easily mapped to data types for resource class identifiers used by most standard mobile discovery protocols. However ontology access is still granted by means of OUUID-to-URL mapping mechanisms [7]. In detail, each resource is characterized by:

- 96-bit ID, globally unique item identifier (e.g., the 96-bit EPC code for RFID tags or the 64-bit MAC address for ZigBee sensors);
- 64-bit OUUID;
- a set of attributes that allow to integrate and extend logic-based reasoning services with application-specific and context-aware information processing;
- a semantic annotation, stored as a compressed RDF document fragment.

In order to compress RDF annotations, a *homomorphic* scheme for XML documents is adopted, enabling query processing directly on compressed annotations, without full preliminary decompression.

The overall architecture is basically a two-level infrastructure (see Figure 2). Pervasive identification and sensing technologies are exploited at the *field layer*, while the *discovery layer* is related to the inter-host communication. Each network host acts as a cluster head for field devices in its direct range, using available communication interfaces (e.g., RFID, ZigBee). It is worth noting that the proposed approach is fully decentralized.

In short, the information-centric framework is based on four interaction stages:

1. extraction of object parameters (carrying object characteristics from the field to the discovery layer);
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;
4. extraction of selected resource annotations (to carry semantic-based descriptions from the field to the discovery layer) for queries and reasoning.

In stages 1 and 4, the semantic support micro-layer of each field technology (e.g., RFID, ZigBee) allows to embed in, and retrieve from, pervasive objects machine-understandable information. It translates data structures and primitives between the uniform resource characterization outlined above and the technology-specific formats. This allows heterogeneous micro-devices at the field layer to be represented homogeneously at the discovery layer.

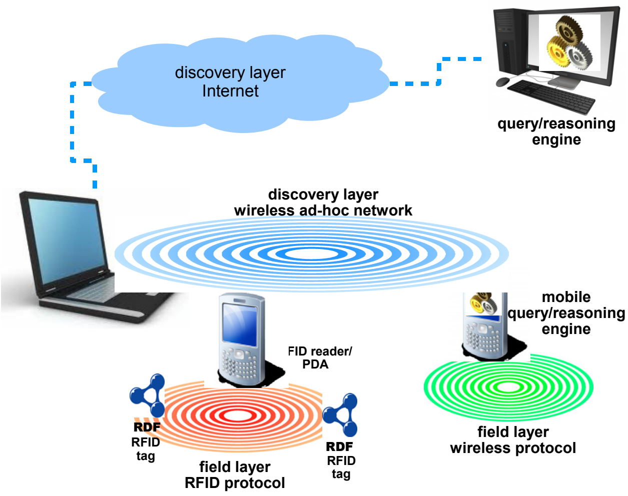

Figure 2. Field and Discovery layers in the *u-KB* framework architecture.

B. *Information dissemination*

An efficient dissemination protocol is a crucial element to balance network usage and effectiveness of resource retrieval. Address and main characteristics of each object are advertised by the related cluster head, using small-size messages. To this aim, resource providers periodically send advertisement messages to inform, and detect the presence of, other nodes in their direct range; each advertisement specifies the maximum number of hops it can travel. During a given lifetime, advertisements are forwarded in broadcasts and can be stored in the cache memories of intermediate nodes. A node that receives an advertisement extracts the resources' information and, in case of "new" elements, stores them; otherwise verifies if the received information is more recent, or has run across a shorter path, and updates its cache. If the memory is updated and the maximum advertisement hops has not been reached, the advertisement is forwarded; differently it is silently discarded. This simple mechanism ensures that each node in the network sends the same advertisement at most once.

C. *Resource discovery*

The discovery procedure occurs in two steps to avoid an uncontrolled flooding and wasting of bandwidth and power. The first step is syntax-based and aims to select resource descriptions potentially relevant for the requester via OUUID matching and contextual parameters evaluation. When starting a resource retrieval process, a node generally attempts to cover the request by using its own cached information. If some semantic description is missing, it can be retrieved in unicast using specific *Request* messages. On the contrary, if a node has no (or not enough) cached information, it can send a *Solicit* message with a given maximum travel diameter to get new resource locators. When receiving a *Solicit*, a node replies (in unicast) providing cache table entries matching the *Solicit* frame. If it does not have any information satisfying the *Solicit*, it will reply with a "no matches" message. During their travel, replies and *Solicit* packets are used to update the cache memory of forwarding nodes. Basically, the soliciting mechanism is analogous to the advertising one, exploiting controlled broadcast of request in an *expanding ring* fashion.

The second discovery step is semantic-based and aims to select the best available resources. The requester downloads semantic annotations directly from the provider, so preparing the further reasoning and semantic query processing. This hybrid, on-demand approach has been chosen because semantic descriptions are needed only in the last discovery phase, whereas a preliminary ontology-based context-aware selection procedure is mandatory. In this way, the induced traffic is significantly reduced.

D. *Early experiments*

The proposed approach has been implemented and preliminarily tested using network simulators. Early results (full analysis not reported) showed that: (i) overall network load is acceptable; (ii) network traffic has higher correlation with the number of providers rather than requesters. This happens because advertisements are regularly sent in a proactive way by providers, even if there are no requests. Furthermore, the data dissemination protocol adapts well to increasing node mobility, when radio links are lost more often; (iii) the hit ratio is very high in general, with values above 90% in all tests; (iv) discovery time decreases as the number of clients increase. This is due to the fact that intermediate nodes cache the resource records, so reducing latency of later requests. Overall values are still slightly high w.r.t. the requirements of pervasive computing scenarios.

V. APPLICATIONS

Hereafter some considerations are presented about the potential adoption of the *u-KB* framework in two widespread smart environments.

A. *Smart city*

A brief assessment of benefits and limits of the proposed framework is provided here w.r.t. the key research challenges outlined in the position paper [10] for pervasive computing approaches aimed at smart city scenarios.

1. Reasoning on stream data: the *u-KB* approach allows exploiting expressive background knowledge (ontologies) and to support non-standard inference tasks for stream reasoning [13], enabling more robust event/trend monitoring, identification and explanation as compared to the approaches mentioned in Section II.

2. Knowledge representation and expressivity: the proposed framework aims at a careful trade-off between knowledge expressiveness and performance of inference services. Previous research and case studies on semantic matchmaking for on-the-fly resource discovery [7] and on inference services for semantic stream reasoning [13] showed that even moderately expressive logic languages such as AΛN (Attributive Languages with unqualified Number restrictions) can adequately capture the semantics of relevant real-world scenarios.

3. Data dynamicity: as explained in Section IV, the *u-KB* primary goal is to closely match the dynamicity of resources. On-the-fly materialization of a subset of the KB allows extracting only the information that is actually available and relevant to the context and user/application needs.

4. Performance and scalability: performance was a key driver in the framework design: annotation language selection, information compression, protocol design and optimization, semantic matchmaking and reasoning exploitation. Preliminary results obtained in a medium-sized (about 50 nodes) *ns-2* simulated test bed evidenced that: the overall network load induced by the *u-KB* dissemination is acceptable (about 0.4 kB/s per node); effectiveness of resource discovery is high (multiple resources were retrieved in over 90% of tests); discovery duration is slightly high (1.5-2.5 s) w.r.t. the requirements of pervasive computing scenarios; compact representation of concept collections allows scalable stream storage and reasoning.

5. Data uncertainty: the *u-KB* model copes with the uncertainty caused by data noisiness and incompleteness in a straightforward way, by naturally inducing a graceful degradation of the QoS in discovery and inference, as Section IV reports. Further techniques to detect incorrect data and/or fill knowledge gaps can be added on top of the general framework, according to application-specific needs.

6. Data privacy and security: the inherent locality of the *u-KB* model mitigates the risk of unwanted exposition of data outside the nearby physical environment. While no specific security feature is embedded at the moment, the framework leverages transparently the security features of the underlying identification/sensing technologies, if present. Finally, semantic-based summarization of concept streams can be fine-tuned to provide the required level of anonymity.

B. Factory automation

Even if a concrete practical application of the *u-KB* model is an ongoing work, early analyses in a specific environment can highlight perspectives, benefits and issues of an extensive adoption of the proposed framework. Particularly, the EU FP7 IoT@Work project (http://www.iot-at-work.eu) aims at designing an IoT architecture that takes into account the needs of the industry with reference to factory automation systems. Specifically, a SWoT approach is devised, namely *Plug&Work IoT* focusing on most common communication issues, and attempting to improve both flexibility and reliability. Specific features being explored and developed are related to system's auto-configuration and security improvements. In what follows some project's related details will be provided in order to make clear as it can be a preliminary proof of concept for the theoretical framework outlined above.

An IoT@Work enabled factory shop floor should reduce operational and capital expenditure in automation issues. By transforming automation devices into *Internet-enabled things,* the plant engineers and maintainers will not take care of configuring each component/element/line/plant during design and commissioning phases. Particularly, the IoT@Work Directory Service (IoT@DiS), a RESTful service, acts as an access point for semantically enriched information about production units and devices deployed in the shop floor. The system is designed in such a way to also support an easy to use access mode simply *targeting the device* about which you need information, using a handheld (smartphone, PDA, tablet, etc.) and exploiting semantic-based deductions.

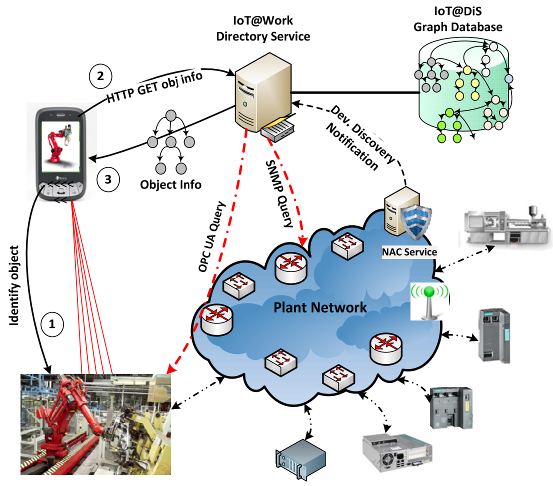

Figure 3. IoT@Work Directory Service

The IoT@DiS manages RDF compliant *semi-permanent* information (as usual for a directory service) related to entities in the production environment. Some of them are stored within its graph oriented database, while others are dynamically acquired querying shop floor devices via *SNMP* (Simple Network Management Protocol) or *OPC UA* (OPC Unified Architecture) services. Ad-hoc services in the shop floor network are in charge of discovering new devices and supporting their configuration. Such services also notify the IoT@DiS about the detection of new devices (see *Dev. Discovery Notification* in Figure 3). The IoT@DiS component still retains its role as a common access point to a set of information relevant for end-users (e.g., workers, supervisors, maintainers, etc.), as well as for production related services that need to acquire semantic annotations about entities in the production plant. It can therefore be considered as an intermediary/gateway between the shop floor and the application layer. The deployment of *u-KB* compliant solutions in the IoT@Work scenario adds flexibility to the basic *Plug&Work IoT* approach. Basically, the directory service acts as a knowledge layer to enable both resource discovery and their proper usage via a semantic-based matchmaking. On the other hand, devices and plant components simply exploit ubicomp features and wireless or wired communication in order to provide and/or obtain information useful for the plant working.

VI. Conclusion

An information-centric networking approach for the Semantic Web of Things was presented. This approach envisages a peer-to-peer, collaborative and dynamic framework supporting semantic information dissemination and resource discovery in pervasive environments. Information gathered in the field through different identification and sensing technologies can be exploited at the discovery layer for local on-the-fly queries/inferences or large-scale stream reasoning. Some applications are presented as proof of concept and to witness the effectiveness of the proposed approach. Experiments and evaluations can be carried out only on practical instances of the *u-KB* theoretical model: early results of experiments performed through network simulators are presented.

Acknowledgments

The authors acknowledge partial support of EU ECTP 2012-14 project "GAIA: Generalized Automatic exchange of port Information Area", and EU FP7 Project IoT@Work (Grant Agreement N. 257367).

References

[1] D. Barbieri, D. Braga, S. Ceri, E. Della Valle, M. Grossniklaus, "Continuous queries and real-time analysis of social semantic data with c-sparql", Proc. of the Second Social Data on the Web Workshop (SDoW 2009), Washington DC (USA), Oct. 2009, URL: http://ceur-ws.org/Vol-520/paper02.pdf, accessed on October 12, 2012

[2] A. Bolles, M. Grawunder, J. Jacobi, "Streaming sparql-extending sparql to process data streams", The SemanticWeb: Research and Applications, Lecture Notes in Computer Science, vol. 5021 (2008), pp.448–462, doi: 10.1007/978-3-540-68234-9_34

[3] C. Bizer, T. Heath, T. Berners-Lee, "Linked data-the story so far". Int. Jour. on Semantic Web and Information Systems, vol. 5(3), 2009, pp. 1–22, doi:10.4018/jswis.2009081901

[4] D. Chakraborty, A. Joshi, Y. Yesha, T. Finin, "Toward distributed service discovery in pervasive computing environments", IEEE Trans. on Mobile computing, vol. 5(2), 2006, pp. 97–112, doi:10.1109/TMC.2006.26

[5] B. Christophe, "Semantic Profiles to Model the 'Web of Things", in 7th Int. Conf. on Semantics Knowledge and Grid (SKG'11), pp. 51-58, Oct. 2011. doi: 10.1109/SKG.2011.29

[6] B. Christophe, V. Verdot, V. Toubiana, "Searching the 'Web of Things", in 5th IEEE Int. Conf. on Semantic Computing (ICSC'11), pp. 308-315, Sept. 2011. doi: 10.1109/ICSC.2011.69

[7] T. Di Noia, E. Di Sciascio, F.M. Donini, M. Ruta, F. Scioscia, E. Tinelli, "Semantic-based Bluetooth-RFID interaction for advanced resource discovery in pervasive contexts", Int. Jour. on Semantic Web and Information Systems, vol. 4(1), 2008, pp. 50–74, doi:10.4018/jswis.2008010104

[8] V. Jacobson, D. Smetters, J. Thornton, M. Plass, N. Briggs, R. Braynard, "Networking named content", Proc. 5th Int. Conf. on Emerging networking experiments and technologies (CoNEXT '09). New York, USA: ACM, 2009, pp. 1–12, doi:10.1145/1658939.1658941

[9] T. Koponen, M. Chawla, B. Chun, A. Ermolinskiy, K. Kim, S. Shenker, I. Stoica, "A data-oriented (and beyond) network architecture", ACM SIGCOMM Computer Communication Review, vol. 37(4), 2007, pp. 181–192, doi:10.1145/1282380.1282402

[10] F. Lecue, S. Kotoulas, P. Mac Aonghusa, "Capturing the Pulse of Cities: A Robust Stream Data Reasoning Approach", position paper for PlanetData Project, EU 7th FP 2007-2013. IBM Research - Smarter Cities Technology Centre, 2011, URL: http://planet-data-wiki.sti2.at/uploads/d/d3/Capturing_the_Pulse_of_Cities%2C_Freddy_Lecue.pdf, accessed on October 12, 2012

[11] U. Lee, I. Rimac, V. Hilt, "Greening the internet with content-centric networking", Proc. 1st Int. Conf. on Energy-Efficient Computing and Networking (e-Energy '10). New York, USA: ACM, 2010, pp. 179–182, doi:10.1145/1791314.1791342

[12] A. Rodriguez, R. McGrath, Y. Liu, J. Myers, "Semantic Management of Streaming Data", Workshop on Semantic Sensor Nets at International Semantic Web Conference (ISWC 2009), CEUR Workshop Proceedings, Vol. 522, URL: http://ceur-ws.org/Vol-522/p12.pdf accessed on October 12, 2012

[13] M. Ruta, S. Colucci, F. Scioscia, E. Di Sciascio, F.M. Donini, "Finding Commonalities in RFID Semantic Streams", The Fifth Int. Workshop on RFID Technology - Concepts, Applications, Challenges (IWRT 2011), Procedia Computer Science, vol. 5, 2011, pp. 857-864, doi:10.1016/j.procs.2011.07.118

[14] M. Ruta, F. Scioscia, E. Di Sciascio, "Mobile Semantic-based Matchmaking: a fuzzy DL approach", Proc. 7th Extended Semantic Web Conference (ESWC 2010), Lecture Notes in Computer Science, vol. 6088. Berlin/Heidelberg, Germany: Springer, 2010, pp. 16–30, doi:10.1007/978-3-642-13486-9_2

[15] A. Toninelli, A. Corradi, R. Montanari, "Semantic-based discovery to support mobile context-aware service access", Computer Communications, vol. 31(5), 2008, pp. 935–949, doi:10.1016/j.comcom.2007.12.026

[16] D. Trossen, M. Särelä, K. Sollins, "Arguments for an information-centric internetworking architecture", ACM SIGCOMM Computer Communication Review, vol. 40(2), 2010, pp. 26–33, doi:10.1145/1764873.1764878

[17] W. Van Woensel, S. Casteleyn, E. Paret, O. De Troyer, "Mobile Querying of Online Semantic Web Data for Context-Aware Applications", IEEE Internet Computing, vol. 15(6), 2011, pp. 32-39, doi: 10.1109/MIC.2011.108

[18] J. Vazquez, D. Lopez-de Ipiña, "mRDP: An HTTP-based lightweight semantic discovery protocol", Computer Networks, vol. 51(16), 2007, pp. 4529-4542, doi:10.1016/j.comnet.2007.06.017

[19] O. Walavalkar, A. Joshi, T. Finin, Y. Yesha, "Streaming knowledge bases", Proc. of the Fourth Int. Workshop on Scalable Semantic Web knowledge Base Systems (SSWS), Karlsruhe, 2008