From the Physical Web to the Physical Semantic Web: knowledge discovery in the Internet of Things

Michele Ruta, Saverio Ieva, Giuseppe Loseto, Eugenio Di Sciascio
Politecnico di Bari
via E. Orabona 4, Bari (I-70125), Italy
email:{michele.ruta,saverio.ieva,giuseppe.loseto,eugenio.disciascio}@poliba.it

***Abstract*—The *Physical Semantic Web* is proposed as a paradigm enhancing the Google Physical Web approach for the Internet of Things. It allows semantic annotations to be associated to beacons instead of trivial identifiers, in order to enable more powerful expressiveness in human-things and things-things interactions. This paper presents a general framework for the Physical Semantic Web, based on machine-understandable annotations of physical resources and novel logic-guided resource discovery capabilities. Possible application scenarios are outlined to highlight the benefits of induced enhancements and the effectiveness of theoretical solutions.**

***Keywords*–Physical Web; Knowledge Representation; Ubiquitous computing**

I. INTRODUCTION AND MOTIVATION

The Physical Web is a paradigm in the Internet of Things framework devised by Google Inc. to enhance the interaction capabilities with real-world objects. A discovery service, running on a mobile device, retrieves Uniform Resource Locators (URLs) of nearby things through the *Eddystone* Bluetooth beacon protocol [1]. This requires neither centralized archives nor special-purpose applications. Finding URLs via Wi-Fi using mDNS (multicast DNS) and uPnP (Universal Plug and Play) is also supported. The Google Physical Web approach tends to preserve legacy applications: it empowers the human-thing interaction when native software is unfeasible or even impractical to be used: every smart object owns a web address, and this makes possible a simple and direct interaction bypassing dedicated apps and on-line backends. Although such an approach induces not negligible enhancements in the object networks manageability, several issues restrain a powerful adoption in even more complex Internet of Things (IoT) scenarios. Particularly, things-things interaction is not enabled yet and discovery mechanisms are too simplistic with respect to what needed in really autonomic IoT scenarios. Basically, interoperability problems have to be taken into account when coping with contexts evolving and modifying continuously.

This paper proposes to extend the Physical Web project exploiting the Semantic Web approach and theory, so enabling advanced resource advertisement and discovery features. Knowledge Representation (KR) promotes interoperability, being a possible means to overcome internal peculiarities of interacting entities. As of now, more and more studies indicate this could be also exploited in the Internet of Things. The so-called *Semantic Web of Things* (SWoT) [2] refers to scenarios where intelligence is embedded in the environment by deploying in the field a plethora of heterogeneous micro-devices, each acting as dynamic knowledge micro-repository.

In the proposed approach semantic annotations are encapsulated in beacons, to enhance representation capabilities of objects in the Physical Web. This adds the possibility of more complex interactions: things become resources exposing knowledge characterizing themselves without depending on any centralized actor and/or infrastructure. In addition, user agents running on mobile personal devices are able to dynamically discover the best available objects according to user's profile and preferences; not simply resources in the surroundings, but the ones better supporting users tasks and needs. The proposed approach still maintains the reference URL-based mechanism detecting all Eddystone-URL beacons in a given environment. Legacy applications are preserved, any off-the-shelf beacon and mobile device supporting the base Eddystone protocol can be adopted. Hence, each URL could target: (i) a basic web page where users access the document via browser; (ii) an annotated web page, where users may view the page and/or exploit features allowed by metadata semantics; (iii) a semantic annotation to be used by agents. Retrieved annotations are exploited in a semantic-based matchmaking [3] setting to compare a request (*e.g.*, user profile) with multiple beacon annotations (*i.e.*, object descriptions). Proper compression techniques are adopted to cope with verbosity of annotations and minimize data transfers. Any resource domain (shopping, transportation, gaming, points of interest, work, and so on) can be explored by simply selecting the conceptualization (*i.e.*, ontology) annotations are grounded on. For each <user profile, resource> pair, a score is the outcome of matchmaking: it assesses the affinity of the beacon with user preferences. Concept Abduction [4] non-standard inference provides also a full explanation about the score, evidencing compatible and missing features. Analogously, in case of incompatibility between preferences and beacon, the Concept Contraction inference [4] detects properties of the beacons causing the mismatch.

The remainder of the paper is organized as follows. The next section frames the background of the proposal, while the following Section III presents the envisioned Physical Semantic Web protocol and framework. Then Section IV introduces reference application scenarios to corroborate the comprehension of what proposed before Section V which closes the paper sketching future work.

II. RELATED WORK

In latest years, interesting approaches were developed to integrate knowledge-based frameworks in Wireless Sensor Networks and the Internet of Things. Resulting architectures

ISBN: 978-1-61208-505-0

largely vary in scope, but usually aim to: (i) exploit ontologies –*e.g.*, [5]– to annotate data, devices and services; (ii) share sensor data along the Linked Open Data (LOD) [6] guidelines by means of RESTful [7] or OGC's Sensor Web Enablement (SWE) [8] web service interfaces. *Sense2Web* [9] is a LOD-based platform to publish sensor data and link them to existing resources on the Semantic Web. Different ontologies were used to describe physical resources, query data and relations to deduce implicit knowledge and integrate sensor information coming from various sources. Likewise, the *Linked Stream Middleware* (LSM) platform [10] fuses data produced by sensors with other LOD sources, by enriching both sensor sources and data streams with semantic annotations. A processing engine is used to perform queries across both dataset types, mashup the data and compute results. Finally, [11] describes an application of knowledge representation to automatically create sensor compositions: user goals, functional and non-functional properties of sensors are described w.r.t. an OWL (Web Ontology Language) ontology so that the envisioned orchestration system is able to combine sensors and processes to satisfy a user request. In [12] ontology-based sensor descriptions allow the users to express requests in terms of device characteristics. Quantitative querying and semantic-based reasoning techniques are combined to improve the resource discovery and select appropriate sensors through exploratory search.

From a communication standpoint, the present paper is based on Bluetooth Low Energy (BLE) and Eddystone. Valid and supported alternatives include 6LoWPAN [13] at the network level and CoAP [14] at the application one. 6LoWPAN enables IPv6 packets to be carried on top of low-power wireless networks, while CoAP is an HTTP-like protocol for interconnected objects, designed for machine-to-machine interoperation of resource-constrained nodes. It follows the REST (REpresentational State Transfer) paradigm for making resources accessible, exploiting a binary data representation and a subset of HTTP methods. Each resource is a server-controlled abstraction, unambiguously identified by a URI (Uniform Resource Identifier). 6LoWPAN can be interfaced to IPv6 and CoAP/UDP to HTTP/TCP, so that sensor data can be accessed also from the classic Web. As an example, the *SPITFIRE* project [15] combines Semantic Web and networking technologies to build a service infrastructure aiming to develop advanced applications exploiting Internet-connected sensors and lightweight protocols, as CoAP. In that framework, sensors are described as RDF triples and service discovery is based on metadata (referred for example to device features or location).

A major issue of most proposals is the requirement for a stable Internet connection and/or a support infrastructure to enable discovery features. This makes them unsuitable connectionless scenarios and to mobile ad-hoc networks of resource-constrained objects. As an example, the work in [16] proposes "IoT gateways" to expose resources between CoAP and HTTP nodes; peer-to-peer (P2P) overlay network techniques are used to enable large-scale discovery among different networks. A key issue is that the approach is based only on a resource name resolution scheme, not allowing the use of articulate resource features for discovery, selection and ranking. Actually, all the above solutions except [12] only allow elementary queries on annotations, and then only basic discovery is possible.

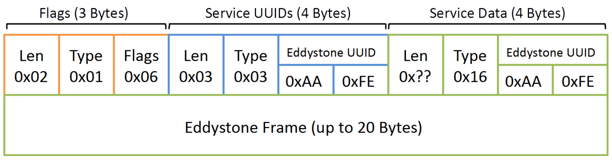

Figure 1: Eddystone Beacon message format

Ontology-based *complex event processing* [17] and *semantic matchmaking* [4] could be used to improve knowledge and object discovery in mobile and pervasive contexts. The latter exploits logic-based reasoning to support approximated matches, resource ranking and explanation of outcomes.

III. THE PHYSICAL SEMANTIC WEB: PROPOSED APPROACH

This section recalls the original Physical Web and outlines the proposed Physical Semantic Web evolution as a comprehensive framework for knowledge management in agent-based IoT environments.

A. *Protocol*

The Physical Web (PW) is an open approach, initially proposed and implemented by Google, aiming to enable on-demand interaction among objects and user devices. Every PW object should be able to expose information to surrounding devices, which can process it without requiring any specific application. The current PW is grounded on two basic elements: a wirelessly broadcast resource identifier (typically a URL) and a mobile device agent able to discover and show collected nearby URLs to the user. *Bluetooth Low Energy* (BLE)[1] beacons supporting the open *Eddystone* [1] application-level protocol are used to expose generic URLs. BLE was introduced in 2010 within the Bluetooth 4.0 Core Specification for Internet of Things scenarios. It uses the same 2.4 GHz Bluetooth radio with a simpler modulation scheme (strongly reducing power usage), ensuring multi-vendor interoperability and a long life-cycle for low-cost devices with standard coin-cell batteries. Over BLE, Eddystone protocol specification defines different formats for proximity beacon messages. All messages share a common PDU (Protocol Data Unit) format, reported in Figure 1. It is composed by: (i) 3 bytes for flags as defined in [18, Part A]; (ii) 4 bytes for service Universally Unique IDentifier (UUID) advertisement, containing the Eddystone Service UUID `0xFEAA`; (iii) resource data, comprising 4 bytes for data advertisement (also in this case the Eddystone Service UUID is included) and up to 20 bytes for the data message payload.

The protocol defines the following Eddystone message types, identified by means of the four most significant bits of the first octet in the data message.

- *Eddystone-UID* (`0x00`): broadcasts a unique 16-byte beacon ID as shown in Figure 2a. Namespace ID can be used to group a set of beacons, while the instance ID identifies individual devices in the group.

---

[1] https://www.bluetooth.com/what-is-bluetooth-technology/bluetooth-technology-basics/low-energy

This partition is useful to improve discovery and filter beacons according to one or more namespaces.

- *Eddystone-URL* (`0x10`): exposes an encoded schema prefix and a compressed and encoded URL (up to 17 bytes), fitting the message format reported in Figure 2b. The URL can be decoded and used by clients to manage the related resource (typically, open a Web page).
- *Eddystone-TLM* (`0x20`): transmits telemetry data useful for monitoring the health and operation of the beacon: battery voltage, device temperature and count of broadcast packets. TLM messages can be either *unencrypted* (version field `0x00`) or *encrypted* (`0x01`), following the format shown in Figure 2c and Figure 2d respectively. In the latter case, beacons must have been previously configured as Eddystone-EID and an identity key should be set during the configuration step.
- *Eddystone-EID* (`0x30`): includes an encrypted ephemeral identifier refreshing periodically during the beacon life-cycle (Figure 2e). This message type is used for security issues (*e.g.*, with encrypted TLM) and privacy-enhanced devices.

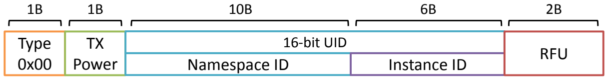

(a) Eddystone-UID message

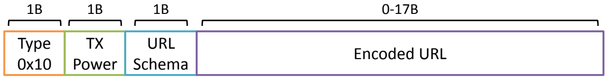

(b) Eddystone-URL message

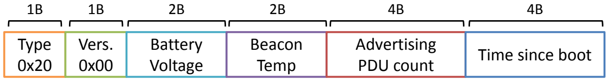

(c) Unencrypted Eddystone-TLM frame

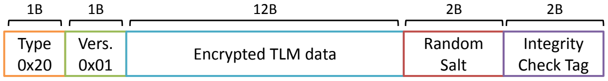

(d) Encrypted Eddystone-TLM frame

(e) Eddystone-EID frame

Figure 2: Format of Eddystone messages

In the current Physical Web proposal, any client supporting BLE can scan the surrounding environment and discover resources (exposed through Eddystone-URL messages) without requiring a directory service or dedicated apps, which are deemed impractical for simple interactions. Google provided reference client implementations for iOS, Android and Node.js. Eventually, scanners should be implemented as background services in operating systems, so as to require no software installation by users. Beacons broadcast URLs pointing to Web pages for informative purposes or to Web apps using advanced technologies to offer interactive user experiences, *e.g.*, push notifications, direct connection and remote control of smart devices via *Web Bluetooth* [19] specification, which enables Web pages to access Bluetooth 4 devices using the *Generic ATTribute (GATT)* Profile to read and/or write attribute values. When multiple beacons are detected in the same area, ranking is based on beacon proximity. Eddystone allows estimating distance by comparing the RSSI (Received Signal Strength Indicator) with the nominal transmitted power recorded in beacon messages (see Figure 2a and 2b). Previous user actions can be also taken into account by implementing the following optional features: (i) *history*, a cache of recently visited URLs; (ii) *favorites*, a list of bookmarked URLs; (iii) *spam*, a list of URLs marked as undesired.

Despite the benefits of a general-purpose and technologically open approach, the Physical Web has several limits: (i) explicit interaction with the user is always required; (ii) beacons can only broadcast a simple URL, not rich resource descriptions; (iii) beacon ranking is based on simplistic distance criteria, without considering common and/or conflicting characteristics of the advertised resources w.r.t. a discoverer's profile or request. This paper proposes an extension of the PW vision, the *Physical Semantic Web* (PSW), exploiting Semantic Web technologies to enable advanced resource discovery. Particularly, in the approach presented here, along with classic and standard web pages each Eddystone beacon could also target annotated web pages or logic-based resource annotations. Both could be used to perform a semantic-based matchmaking [4], which exploits standard and non-standard inferences to give a logic-based ranking of nearby resources w.r.t. a request based on the meaning of their descriptions.

From a communication standpoint, the usage of Eddystone-URL beacons presumes an Internet connection will be available to retrieve resources pointed by broadcast URLs. In several real-world scenarios MANETs (Mobile Ad-hoc NETworks) and point-to-point infrastructure-less connections provide a more flexible and effective solution for wireless low-power networking. They can be particularly useful in ubiquitous scenarios, where mobile objects must provide quick decision support and/or on-the-fly organization in such intrinsically unpredictable environments. As shown in Figure 3, the PSW exploits Eddystone-UID beacon messages to transmit: MAC address of the device exposing the resource; instance ID, adopted to identify a specific local resource provided by the object; the protocol to be used to retrieve the resource annotation (*e.g.*, Bluetooth, Wi-Fi Direct). Such an approach considerably increases the flexibility and autonomicity of the basic PW for what concerns resource management, dissemination and discovery.

B. *Framework*

The proposed framework enhancing the standard PW solution is based on four elements.

A. Machine-understandable standard language to express information with rich and unambiguous semantics. The devised PW extension supports both human-to-machine and machine-to-machine interactions. This grants flexibility for accommodating a wider range of scenarios, including agent-based systems with implicit interaction patterns or with no

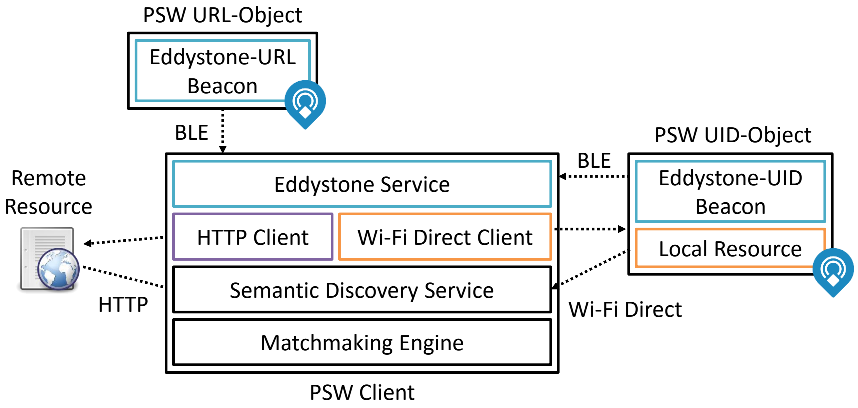

Figure 3: Physical Semantic Web architecture

user involvement at all. The proposal relies on Semantic Web languages, and particularly on Resource Description Framework (RDF) [20] and Web Ontology Language (OWL) [21]. RDF defines a general knowledge assertion model for characterizing and linking resources through statements structured as *subject-predicate-object* triples. OWL allows defining *ontologies* on top of RDF, *i.e.*, expressive vocabularies for modeling structured knowledge about particular domains. The semantics of OWL is formally grounded on the Description Logics (DLs) family of logical languages. An ontology and a set of individuals build the *Knowledge Base* (KB) used for automated *reasoning* supporting discovery in the reference domain.

B. Objects exposing semantic annotations. From now on, the term object will not refer to the RDF-specific meaning, but will denote material things equipped with storage, processing and communication facilities. Ontologies are assumed as stable background knowledge, available either through direct wireless exchange between smart objects or through the Web. Every object exposes a semantic annotation describing its state, capabilities and/or factual knowledge collected from the context it is dipped in. Each annotation corresponds to a KB individual, expressed w.r.t. a particular domain ontology. Objects materialize structured information with rigorous semantics and make it discoverable to nearby devices through BLE via Eddystone-UID frames, as explained in Section III-A. In this way, knowledge exchange can occur through point-to-point connectivity, even without Internet connection. Whenever Internet access is available, they also publish the same knowledge fragments on the Web and advertise them via Eddystone-URL frames, in order to make them retrievable through the standard PW mechanism. No customized application-level protocols are needed to mediate interactions and knowledge sharing: starting from a logical core information grounded on a reference ontology, an object is able to update and enrich its annotation during its lifecycle, in order to reflect evolution in its perceptions, goals and functionalities.

C. Knowledge discovery and sharing. When an object exposes a semantic-based annotated information, agents running on nearby devices and objects can discover it via PSW, as outlined in Section III. The PSW push policy allows agents to be notified of nearby annotation instances. Nevertheless, discovery is driven by application requirements, expressed as a semantic-based request in a matchmaking problem. The discoverer agent collects UIDs and URLs from neighboring devices and preselects only the ones corresponding to semantic annotations having the same reference ontology as the request. This preliminary filter sets the general knowledge domain for the current discovery session and excludes irrelevant knowledge fragments, so reducing the communication and computational load of the subsequent matchmaking step. The matchmaking outcome is a list of annotations ranked by semantic similarity w.r.t. the request. In classical PW scenarios, the user is in control of the discovery process: she selects one of the returned results pointing to a Web page with human-readable information and possible actions. In more advanced scenarios, the discovery process is performed autonomously by an agent device, equipped with knowledge representation tools to select the best result(s) and guide automatic interactions between objects.

D. Semantic matchmaking. Knowledge discovery is supported by a rigorous semantic matchmaking framework to rank a set of *resources* according to relevance with respect to a *request*, where the resources and the request must be satisfiable concept expressions with respect to a common ontology. Standard reasoning services for matchmaking include *Subsumption* and *Satisfiability*. Given a request R and a resource S, Subsumption verifies whether all features in R are included in S: its outcome is either *full match* or not. Satisfiability checks whether any constraint in R contradicts some specification in S, hence it divides resources in *compatible* (a.k.a. *potential matches*) and *incompatible* (a.k.a. *partial matches*) w.r.t. the request. This approach is inadequate for fully autonomic scenarios, because full matches seldom occur and incompatibility is frequent when dealing with complex expressions from independent heterogeneous sources. One would like to determine *what* constraints caused incompatibility or missed full match. In order to produce a finer resource ranking and a logic-based explanation of outcomes, the framework exploits *Concept Abduction* and *Concept Contraction* non-standard inference services [4]. Given a request R incompatible with an available resource S, Contraction detects what part G (for *Give up*) of R is conflicting with S. If one retracts G from R, K (for *Keep*) is obtained, which represents a contracted version of the original request, such that it is compatible with S. On the other hand, if R and S are compatible but S does not match R completely, Abduction identifies what additional feature set H (for *Hypothesis*) should be assumed in S in order to reach a full match. *Penalty functions* are associated to Abduction and Contraction, in order to compute a semantic distance metric for ranking a set of resources w.r.t. a given request. The overall matchmaking process is summarized in Figure 4. Efficient implementations of the above inferences exist for mobile and embedded computing architectures on moderately expressive DLs [3]. The final ranking score integrates semantic distance with context-aware data-oriented attributes:

$$f(R,S) = 100[1 - \frac{penalty(R,S)}{penalty(R,\top)}(1 + dist(R,S))]$$

where $penalty(R,S)$ is the penalty induced by Abduction and Contraction from request R and resource S; this value is normalized dividing by the penalty between R and the universal concept (a.k.a. *Top* or *Thing*), which depends exclusively on axioms in the reference ontology. The $dist(R,S)$ term is the physical distance between the discoverer agent and the discovered resource's owner. Nearer resources are usually preferred, because (i) knowledge locality is often important in pervasive applications and (ii) shorter hops in wireless communications are more reliable and less energy-consuming. The formula for f translates the semantic distance measure into a *relevance* 0-100% ascending scale.

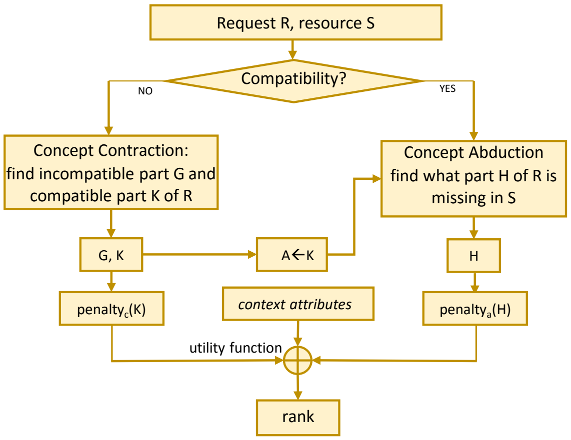

Figure 4: Semantic context-aware matchmaking

The above elements implement a *ubiquitous Knowledge Base* (u-KB) model [2] where the classical components of a KB (ontologies and individuals) are attached to ubiquitous smart objects and no centralized storage or processing infrastructure is required. At the *field layer*, mobile hosts extract information from the environment through embedded sensing and identification technologies. At the *discovery layer*, hosts communicate by exposing and searching knowledge fragments through the Physical Semantic Web. Since objects can expose multiple resources referring to different ontologies, several u-KBs can co-exist in the same physical space: the PSW protocol allows to materialize a proper subset of the u-KB of interest when needed for reasoning.

IV. APPLICATION SCENARIOS

Several reference scenarios are now outlined to exemplify practical applications of the Physical Semantic Web approach. Basically, proposed scenarios exploit two different architectures:

- Internet-based scenarios: a wireless Internet connection is available on the user's mobile device (*e.g.*, via 3G, LTE, Wi-Fi communication). Moreover, the user installed the PSW mobile app to discover and interact with nearby beacons;
- Connection-less scenarios: an Internet connection is not available and PSW objects can communicate via peer-to-peer connections in infrastructure-less networks.

It is important to notice that basic PW framework only considers Internet-based scenarios whereas PSW also enables to exploit BLE beacons in many real-world IoT environments.

1. Distributed Sea Surveillance. Monitoring large marine areas is essential to prevent environmental emergencies and rescue victims of nautical disasters. Innovative surveillance strategies leverage multimodal sensor platforms mounted on water drones or on buoys placed in strategic locations. In such scenarios, Internet connectivity is generally unavailable: *nomadic networks* are a more effective paradigm to minimize data transfers. This can save both bandwidth and power of stationary nodes (*e.g.*, buoys), which must run unattended for long time spans. Each sensor node must be able to perform a mining and summarization process on the potentially large streams of raw data, in order to extract only interesting events and patterns. Such summaries can be expressed with short, high-level formal annotations in Semantic Web languages. Exploiting the PSW proposal, the mobile nodes of the nomadic network (*e.g.*, drones) patrolling an area can discover and collect annotations of relevant events and conditions from stationary nodes. If intervention is required, mobile nodes can provide immediate assistance for minor issues or alert the command center in case of more complex operations.

2. Precision Agriculture. Let us consider a simple case study where an agricultural land is divided into several fields, farmed with different type of products each characterized by a set of features (*e.g.*, age, growth stage). Each field is managed autonomously by a team of robots (sensors and actuators), acting as smart objects and able to process data to produce shareable useful knowledge. *Monitor* robots collect data in the field, create a local semantic-based annotation and expose an *Eddystone-UID* beacons indicating how to retrieve the description. In particular the beacon message contains the MAC address of a Bluetooth device (not necessarily the same BLE device exposing the beacon) and the ID of the file to be retrieved by means of the Bluetooth File Transfer profile. These descriptions are discovered, downloaded through a P2P connection and then processed by *actuators*, equipped with a mobile matchmaker, to identify the most suitable areas where perform required actions (*e.g.*, irrigation, fertilization).

3. Self-driving vehicles. An application of the Physical Semantic Web will be proposed to support long-term performance evaluation tests of vehicles. Currently this kind of tests deals with several issues, as continuous long-term road test sessions are a very intensive physical and mental task for human drivers often spanning several weeks. Self-driving vehicles can overcome these limitations, also extending the maximum duration of tests. However novel challenges must be taken in consideration. Autonomous driving requires many complex interactions with the surrounding environment. The vehicle has to check its status, understand the external context and decide how to respond to the detected conditions. By embedding BLE beacons in the road surface, self-driving cars can be not only guided to remain on track, but also informed about context conditions in an articulated way. Some test centers, like *Porsche Nardò Technical Center* (http://www.porscheengineering.com/nardo/), are already starting investigations in this direction. Further developments could be then applied to real-world self-driving scenarios. By embedding a BLE beacon in a car, the vehicle becomes a *moving beacon* able to annotate and share both context and vehicles properties in order to dynamically adapt its driving style in presence of modifications of the test environment (*e.g.*, change of weather conditions and wear of vehicle components).

V. CONCLUSION AND FUTURE WORK

The paper proposed a theoretical framework enabling the Physical Semantic Web, a novel paradigm enhancing the Physical Web program by Google. It basically applies the Semantic Web of Things vision to the real world. Models for knowledge sharing and discovery have been extended to fully autonomic environments populated by smart objects. Application scenarios have been also presented to make evident benefits the approach could drive to reach. Future work will be oriented to: (i) validate the approach through an extensive experimentation (supported by Google Inc.) with native Physical Web devices; (ii) align the early Physical Semantic Web proposal to the latest PW features (*e.g.*, *FatBeacon* specification); (iii) further investigate objects interaction schemes in order to enable more effective data management; (iv) extend the current PW Android client (http://github.com/google/physical-web/tree/master/android) integrating the proposed semantic-based discovery of BLE beacons; (v) implement and test the multi-robot scenario, described in IV, exploiting both simulation software and off-the-shelf robots. Robot Operating System (ROR) will be used as reference platform.

More information about the Physical Semantic Web project can be found on the reference web page (http://sisinflab.poliba.it/swottools/physicalweb) whereas all software updates will be uploaded on the GitHub repository (http://github.com/sisinflab-swot/physical-semantic-web), created as a fork of the official Physical Web project.

ACKNOWLEDGMENT

The authors acknowledge partial support of Google Inc. for valuable endorsements within the frame of IoT Research Award.

REFERENCES

[1] "Eddystone protocol specification," http://github.com/google/eddystone, [Online; accessed 20-Sep-2016].

[2] M. Ruta, F. Scioscia, and E. Di Sciascio, "Enabling the Semantic Web of Things: framework and architecture," in 2012 IEEE Sixth International Conference on Semantic Computing. IEEE, 2012, pp. 345–347.

[3] F. Scioscia, M. Ruta, G. Loseto, F. Gramegna, S. Ieva, A. Pinto, and E. Di Sciascio, "A mobile matchmaker for the Ubiquitous Semantic Web," International Journal on Semantic Web and Information Systems, vol. 10, no. 4, dec 2014, pp. 77–100.

[4] M. Ruta, E. Di Sciascio, and F. Scioscia, "Concept abduction and contraction in semantic-based P2P environments," Web Intelligence and Agent Systems, vol. 9, no. 3, 2011, pp. 179–207.

[5] M. Compton, P. Barnaghi, L. Bermudez, R. Garcia-Castro, O. Corcho, S. Cox, J. Graybeal, M. Hauswirth, C. Henson, A. Herzog 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, 2012.

[6] 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, 2011, pp. 1–136.

[7] K. Janowicz, A. Bröring, C. Stasch, S. Schade, T. Everding, and A. Llaves, "A restful proxy and data model for linked sensor data," International Journal of Digital Earth, vol. 6, no. 3, 2013, pp. 233–254.

[8] A. Bröring, P. Maué, K. Janowicz, D. Nüst, and C. Malewski, "Semantically-enabled sensor plug & play for the sensor web," Sensors, vol. 11, no. 8, 2011, pp. 7568–7605.

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

[10] 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, Nov. 2012, pp. 42–51.

[11] 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.

[12] 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," Sensors Journal, IEEE, vol. 14, no. 2, 2014, pp. 406–420.

[13] N. Kushalnagar, G. Montenegro, and C. P. Schumacher, "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals," Internet proposed standard RFC, vol. 4919, August 2007.

[14] C. Bormann, A. Castellani, and Z. Shelby, "CoAP: An Application Protocol for Billions of Tiny Internet Nodes," IEEE Internet Computing, vol. 16, no. 2, 2012, pp. 62–67.

[15] D. Pfisterer, K. Romer, D. Bimschas, O. Kleine, R. Mietz, C. Truong, H. Hasemann, A. Kroller, M. Pagel, M. Hauswirth et al., "SPITFIRE: Toward a Semantic Web of Things," Communications Magazine, IEEE, vol. 49, no. 11, 2011, pp. 40–48.

[16] S. Cirani, L. Davoli, G. Ferrari, R. Léone, P. Medagliani, M. Picone, and L. Veltri, "A scalable and self-configuring architecture for service discovery in the internet of things," Internet of Things Journal, IEEE, vol. 1, no. 5, 2014, pp. 508–521.

[17] K. Taylor and L. Leidinger, "Ontology-driven complex event processing in heterogeneous sensor networks," The Semanic Web: Research and Applications, 2011, pp. 285–299.

[18] Bluetooth SIG, "Supplement to the Bluetooth Core Specification, Version 5," 2014. [Online]. Available: https://www.bluetooth.org/DocMan/handlers/DownloadDoc.ashx?doc_id=291904

[19] W3C Draft Community Group, "Web Bluetooth," https://webbluetoothcg.github.io/web-bluetooth/, Tech. Rep., Sep. 2016.

[20] G. Schreiber and Y. Raimond, "RDF 1.1 Primer," W3C, W3C Note, Jun. 2014, http://www.w3.org/TR/rdf11-primer.

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