Managing Things and Services with Semantics: A Survey

Matthias Thoma*‡ Torsten Braun‡ Carsten Magerkurth* Alexandru-Florian Antonescu*‡
*SAP (Switzerland) Inc., Althardstrasse 80, 8105 Regensdorf, Switzerland
‡Communication and Distributed Systems, University of Bern, Neubrückstrasse 10, 3012 Bern, Switzerland
matthias.thoma@sap.com, braun@iam.unibe.ch, carsten.magerkurth@sap.com, alexandru-florian.antonescu@sap.com

***Abstract*—This paper presents a survey on the usage, opportunities and pitfalls of semantic technologies in the Internet of Things. The survey was conducted in the context of a semantic enterprise integration platform. In total we surveyed sixty-one individuals from industry and academia on their views and current usage of IoT technologies in general, and semantic technologies in particular. The semantic enterprise integration platform aims for interoperability at a service level, as well as at a protocol level. Therefore, also questions regarding the use of application layer protocols, network layer protocols and management protocols were integrated into the survey. The survey suggests that there is still a lot of heterogeneity in IoT technologies, but first indications of the use of standardized protocols exist. Semantic technologies are being recognized as of potential use, mainly in the management of things and services. Nonetheless, the participants still see many obstacles which hinder the widespread use of semantic technologies: Firstly, a lack of training as traditional embedded programmers are not well aware of semantic technologies. Secondly, a lack of standardization in ontologies, which would enable interoperability and thirdly, a lack of good tooling support.**

I. Introduction

Research on semantics and semantic management of Internet of Things systems has attracted a lot of interest in the last ten years, but so far, has failed to gain widespread use in industrial applications. From an European research point of view (but not limited to that) a lot of resources, both in manpower as well as in financial support, has gone into semantic research. Most projects that have been ramped up in past five years in the context of Future Internet [1], Internet of Things [2] or Industry 4.0 [3] use semantic technologies in one way or the other. Just to name a few, the SENSEI project [4] for example, was funded with €14.9 million. The more recent Internet-of-Things Architecture project (IOT-A) [5], which is considered as an EU flagship project, received around €11.9 million. More semantic IoT-related projects will start as part of HORIZON 2020. So there is definitely a lot of research being conducted, but when looking into commercialized products it is obvious that semantic technologies in IoT so far failed to deliver on its promise. Semantic technologies still are not of wide-spread use in real-world applications.

In this work, we first present our vision of a semantic management of things, services and devices. We then continue with presenting a recent study on semantics in Internet-of-Things applications, that was mainly conducted to gain insight into potential further usage cases of this platform and further development options. Our work on linked services motivated surveying transport and application layer protocols, which otherwise is not much connected with semantic technologies. While the main focus of the survey was to get a feeling about the view of the community on semantics, it also revealed some interesting insights about application level protocols, transport level protocols and network management which is of interest for a broader audience.

II. Managing Entities and Services

A. Introduction

In the following we briefly present our vision of a semantic enterprise integration platform [6][7]. This section provides an overview of the main building blocks of a semantic IoT platform. It introduces several interaction points between semantics on one hand and management, technologies and protocols on the other hand. First, we introduce the overall architecture and then continue with explaining the management of entities (Section II-C), services (Section II-D) and devices (Section II-E).

B. Architecture

The overall architecture of our system is shown in Figure 1. It consists of several semantic repositories storing information about business entities and a service repository storing information about available services. Furthermore, it utilizes a special SQL-like construction language, called SPBEQL [6], to construct one entity out of the business repositories.

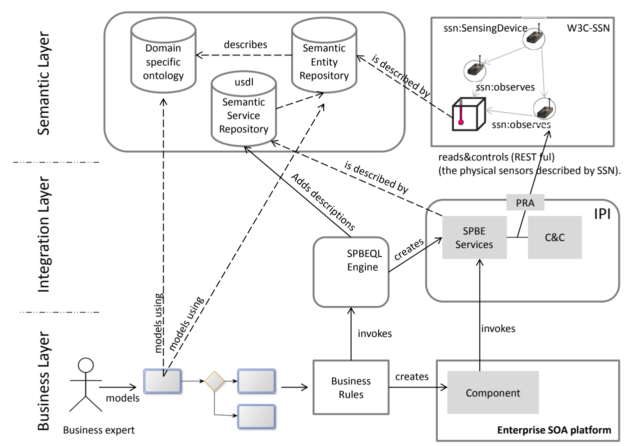

Fig. 1. Integration platform architecture [6]

A domain expert is modelling IoT processes from which business rules can be deducted. Based upon that, there is an automatic generation and configuration of code running on the sensor devices.

C. Entities

We are using the concept of semantic physical business entities (SPBE) [6] to abstract things. The entities do not need to be explicitly defined upfront, but can be generated by the user of the system by combining ontologies. SPBEs are defined as follows:

> A *Semantic Physical Business Entity* is a conceptual representation of a real-world object processed by one or more enterprise IT systems. Information about it is discoverable. It is described through well defined vocabularies, that make internal and external relationships explicit. [6]

Generally speaking, this means that entities can be constructed on the fly by a query language. SPBEs are composed of discoverable information described via ontological links.

D. Services

We are applying the same semantic web principles to services as we do to entities. All services are described in a service description language called Linked USDL4Iot [8]. Most approaches for enabling interoperability between systems use high-level protocols like HTTP[9] or SOAP/XML[10] for integration. Some work has been done on transferring that concept to the Internet of Things [11]. Using such high-level protocols interoperability comes for a price in terms of computing power, energy consumption and delay. Some attempts to overcome the interoperability challenge include

978-1-4799-0913-1/14/$31.00 © 2014 IEEE

the use of ”more compact” application layer protocols like CoAP [12] or stripped-down versions of enterprise application protocols [13]. Nonetheless, as also the survey shows, in IoT often custom low level protocols are used.

Instead of allowing only one, or only a predefined set of endpoint technologies (like CoAP and HTTP) we are using the concept of semantic service descriptions, that separate the service description from the actual service endpoint. We define the term service description as follows:

> A service description is a description of all essential properties of a given service, as well as the means to access it. A service description is independent of an actual implemented callable service. [7]

This means, that a service and its endpoint are semantically described, which allows a seamless integration into enterprise systems and enables high-level as well as low-level service (endpoint) interoperability.

E. Devices

Device management is currently only done with the Semantic Sensor network (SSN) ontology[14], as the system was originally tailored towards sensor/actor networks. Nonetheless, the system is not limited to SSN. Further ontologies, such as SensorML[15] and the Device Description Language (DDL) [16], that is widely used in industrial automation, could be integrated as well.

III. SURVEY

A. Introduction

We conducted a survey on semantics within the Internet of Things domain with an emphasis on integration of enterprise IT systems. As already outlined in Section I semantics are currently under intense research from both industry and academia but so far failed to achieve a breakthrough in actual industrial usage. Our objective was to identify the actual needs of IoT with regard to semantic support and to identify current shortcomings.

B. Methodology

The survey was distributed among internal and external experts, from both industry (among others: SAP, IBM, NEC, Orange, Telefonica) and academia. While some experts were recruited directly, the majority of the respondents were self-selected. They survey was conducted online and anonymity was guaranteed and technically enforced by the system. As IoT is a very broad field, we explicitly excluded all kinds of mobile phone development (for example sensing with mobile phones) and limited protocol related questions specifically to systems where an ISO/OSI-like stack is being used. Nonetheless, we also briefly surveyed the usage of other technologies (Bluetooth, RFID. NFC and Mobile).

C. Threats to internal or external validity

As the study was conducted anonymously it is not possible to validate that the claims made are valid. Nonetheless, we added some sanity checks that allowed to filter non valid responses. No incentives were given for participating. Most industrial participants worked with IoT-systems in industrial automation, retail or logistics. The responses by participants from academia were (if a sector was chosen) mainly from automation and logistics, and the broad areas of smart city. Other sub-fields of IoT may have different requirements regarding protocols, but we expect the tendencies discovered in our survey to be generalizable.

D. Results

In the following we present and discuss the results of the survey. We categorize the results into four groups: (i) General statistical information about the participants and their skillsets, (ii) Protocols, (iii) Semantics and (iv) Enterprise Integration.

1) General

The total number of participants who participated the survey was sixty-one. Their experience levels as well as their origin and skillsets are detailed out in Table I. There were nearly as much participants from industry as from academia. The majority of participants had at least three years of professional experience and a more than basic understanding of IoT and semantics. Most people from academia had experience (skills) on the advanced or expert level. Naturally, the expertness in enterprise software and systems was higher for the industry participants.

Participants		Career level	
Total	61	Entry	8
Industry	32	Advanced	17
Academia	29	Professional	36
Experience (in years)		**Skills (IoT)**	
1-2	6	No experience	0
3-5	31	Beginner	11
5-9	15	Some experience	22
10-14	5	Advanced	21
>15	4	Expert	7
Skills (Semantics)		**Skills (Enterprise systems)**	
Beginner	26	Beginner	28
Some experience	12	Some experience	11
Advanced	17	Advanced	14
Expert	6	Expert	8
Sector			
Industrial automation	14	Home automation	4
Retail	12	Transportation and logistics	7
Smart City	5	Healthcare	2
Vehicular communications	5	Other / none	12

TABLE I
PARTICIPANT GROUP: EXPERIENCE AND SKILLS

Most of the projects were either in the area of (wireless) sensor network or other connected (embedded) constrained devices. While the actual devices were almost always constrained, the network was not. There was a large group using 802.15.4 based wireless, but also some with 802.11 networks, a regular (Ethernet) wired connection or combinations thereof.

2) Protocols

Application layer protocols, as shown in Table II, seem still to be dominated by custom written protocols. Nonetheless, standardized protocols like HTTP, CoAP[12] or MQTT[17] are used by nearly half of the participants. SOAP, which is otherwise widely used in enterprises [18], does not seem to play a role at all. While not adopted widespread yet, in future, most people seem to anticipate CoAP as one of the major players, winning shares from all other standardized protocols and the custom ones. Nonetheless, when comparing the protocols planned to be used in future own developments (i. e. what the participants really plan) and the expected future usage of the industry as a whole (i. e. what they think the industry will move towards), than the expectations towards CoAP are even higher. The number of people planning with and expecting custom protocols is still quite high. In terms of network/transport layer protocols there seems to be an expected shift towards IPv6/6LoWPAN and UDP/CoAP based protocols. It is surprising, given the size of the ZigBee Alliance, that the ZigBee protocol suite is not used more often. ZigBee was almost always selected together with 6LoWPAN, so that most likely even within the ZigBee universe its IPv6 enhancements (Zigbee IP) are used to ensure interoperability.

The views on the community with regards to the current IoT-protocols has been surveyed with a 4-point Likert-style questionnaire. The Likert items as well as the responses are illustrated in Figure 2. There seems to be consensus that most future IoT-applications will be IP-based to some degree, and a bias towards ReST-based architectures. This matches the results in Table II.

Dedicated management protocols, like the Simple Network Management Protocol (SNMP)[19] or CMIP/CMIS[20], seem to be not that widely used at the moment as one would expect. Most participants, used (if at all), a custom application-specific layer on top of the already used application layer protocol. Most participants (>70%) consider nowadays management protocols as not sufficient enough (Figure 2). Interestingly, many respondents also do not plan to use a standardized protocol in future. Generally speaking, the need for management seems to be recognized though, as a more widespread

Application layer	Now	Future (planned)	Future (expected)
CoAP	8%	28%	44%
HTTP	14%	12%	21%
SOAP	2%	2%	6%
CAN	2%	4%	4%
MQTT	5%	6%	7%
KNX	6%	6%	7%
MODBUS	4%	6%	7%
Other/Custom	47%	38%	10%
Zigbee	8%	10%	8%
Transport layer	**Now**	**Planned**	**Expected**
UDP	19%	21%	14%
Reliable UDP (non CoAP)	24%	23%	12%
TCP	13%	11%	9%
UDP + CoAP	8%	24%	39%
Custom/other (TCP-like)	13%	7%	14%
Custom/other	23%	14%	12%
Network layer	**Now**	**Planned**	**Expected**
IPv4	5%	3%	2%
IPv6	21%	23%	35%
6LoWPAN	25%	45%	32%
Custom 802.15.4 protocol	21%	14%	10%
Custom (other)	22%	9%	12%
Zigbee	6%	6%	9%
(Dedicated) Network Management	**Now**	**Planned**	**Expected**
CMIP/CMIS	4%	3%	5%
SNMP	15%	24%	49%
IEC104	0%	2%	8%
Custom protocol based on application layer protocol	44%	37%	20%
Other custom protocol	37%	34%	18%
None / not at all	24%	18%	4%

TABLE II
USED PROTOCOLS (IN PERCENT FROM TOTAL NUMBER OF PROTOCOLS USED, NONE VALUES IN PERCENT FROM PARTICIPANTS), MULTIPLE ANSWERS WERE POSSIBLE

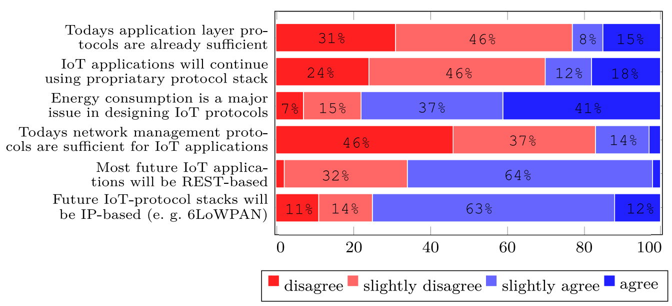

Fig. 2. Usage and potential of IoT-protocols (on a 4-point Likert scale)

use is expected in future. When asked why not using an existing management protocol, the vast majority answered that they expected the overhead of a standardized protocol as too high, or that they fear negative consequences with regard to performance and power consumption. The usage of a custom protocol, tailored towards their specific needs, seems to give them more confidence in the qualitative and quantitative properties of the system, even for the price of a lack of interoperability. As these concerns are not new, there are efforts to run (subsets) of, for example, SNMP [21][22] and NETCONF [23] on resource constrained nodes. CoAP-based protocols (e.g. [24]) have also been investigated, nonetheless, these are also non standardized custom protocols on top of the application layer protocol.

Considering that the Internet of Things originated from RFID and to some degree was driven by the Auto-ID Labs [25], we surveyed other technologies than those based on ISO/OSI (Internet)-like stacks. As can be seen in Table III quite some people use RFID, Bluetooth and QR codes. NFC still seems not that much used at the moment by the participants. Nonetheless, this might be due to the participants as none of them was from financial sector or mobile ticketing, where NFC has gained some usage. Those who are using NFC

Technology	Now	Future (planned)	Future (expected)
RFID	25%	27%	24%
QR codes	5%	8%	11%
Barcodes	12 %	10%	7%
NFC	10%	15%	39%
Bluetooth	7%	6%	19%
None	41%	34%	– %

TABLE III
TECHNOLOGIES USED (IN PERCENT OF PARTICIPANTS), MULTIPLE ANSWERS WERE POSSIBLE

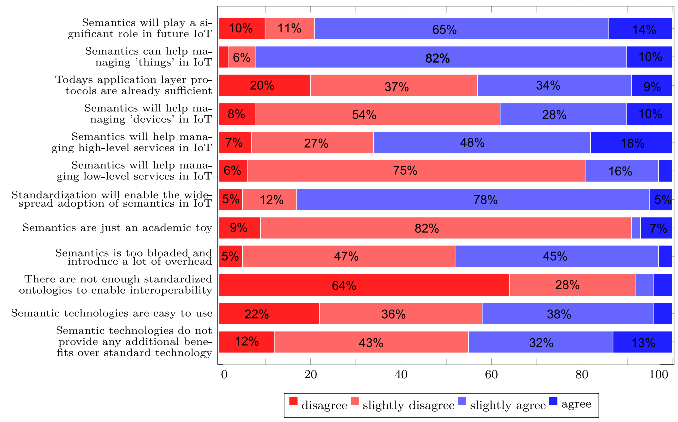

Fig. 3. Usage and potential of semantics (on a 4-point Likert scale)

are solely from retail and industrial automation.

3) Semantics

The general attitude of the community towards semantics was surveyed with a four point Likert-style questionnaire. The individual Likert items and the distribution of the answers are shown in Figure 3. Most participants agree that semantics will play a role in future IoT systems. Some nonetheless, think that it is too bloated/an academic toy, or as one of the participants wrote ”a hype from bored academics that noone will remember in a few years” When asked what is needed for a widespread adoption of semantics in the IoT (see Figure 4) the by-far most often mentioned issues were knowledge / awareness of development staff and standardization, followed by development tool support. Infrastructure and tool support for domain experts did not seem to be an issue. As illustrated in Figure 5 the main advantages of semantics is seen in high-level interoperability and the management of things, followed by reasoning and the management of devices. Interoperability at an endpoint level, as suggested in our integration platform, is not yet seen as an area where semantics can contribute.

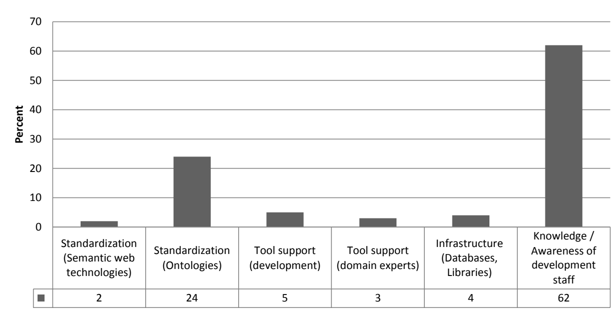

Fig. 4. Main obstacle for not using semantic technologies (in percent)

Given the fact that most people think that there is potential in semantics, it is interesting to see its actual usage: As shown in Table IV, out of all participants 41% do not use semantics at all and 34% do not plan to use it in a further project. From those who use semantics the vast majority is using it for the description of things, devices or services. Reasoning on top of the semantic data seems to be a topic that many people have on their radar. When asked which (domain specific) ontologies are used, most answers centered around custom/problem specific ontologies. The only ontology that was mentioned more often was SSN. The lack of standardization or at least of de-facto standards has been mentioned several times. This corresponds to the results in Figure 2.

Technology	Now	Future (planned)
Description of endpoint level services	7%	22%
Description of Things	67%	78%
Description of Devices	32%	48%
Description of high-level services	38%	45%
Reasoning	22%	35%
Configuration	5%	12%
None / Not at all	41%	34%

TABLE IV
SEMANTIC TECHNOLOGY USAGE (PER PARTICPANTS NOT SELECTING NONE), MULTIPLE ANSWERS WERE POSSIBLE

The time horizon of supporting semantics in a product or product prototype (industry), or do research is shown in Table V. Industry participants generally expected to not do or use semantics in the next 2 years and more ($> 70\%$). Participants from academia do plan to work with semantics often and within a timespan of less than two years (also around 70%).

	Industry	Research
not at all	5%	12%
next 6 months	8%	24%
1 year	11%	35%
2 years	24%	17%
more than 2 years	52%	12%

TABLE V
USE OF SEMANTIC TECHNOLOGIES (TIME HORIZON)

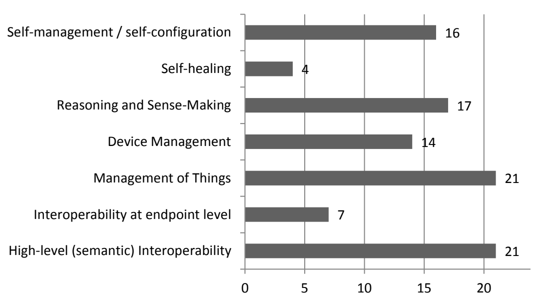

Fig. 5. Attitude towards opportunities/usages of semantics in IoT

4) Enterprise integration

Enterprise Integration is still mainly done through middleware or via HTTP proxies, in cases where appropriate. Direct communication, without an intermediary, which is one of the key elements in the Internet of Things vision, has not yet been widely adopted in enterprise integration frameworks. As shown in Figure 6, only 17% of the enterprise integration projects were done via a direct access (e. g. IPv6), while most solutions seem still to use a gateway solution. This is most likely due to historical reasons, as IoT used to use custom protocols which required a gateway solution. Most IoT-systems run by our participants have either a fully automated configuration scheme or some kind of technical administrator responsible. End-users are only in 14% of all cases solely responsible for the configuration of their device. Most monitoring and management (Figure 7) activities in such environments are also done on this intermediary. Data gathering and aggregation is mostly done either solely on the device (37%) or on the gateway (46%). This, to some degree, might also explain why management protocols did not gain widespread use yet. Most monitoring platforms, on the other hand, are centralized, and often do not yet provide a real-time view on the system. A device-level only monitoring (e. g. by the user only) is not very common.

IV. CONCLUSION

While the survey is only a still picture of a subset of the current situation in IoT, some conclusions can be drawn. The

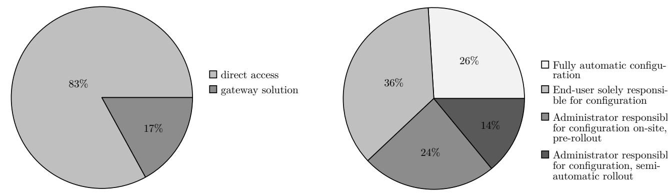

(a) Direct access to motes on network/transport layer vs gateway solutions

(b) Configuration

Fig. 6. Enterprise integration

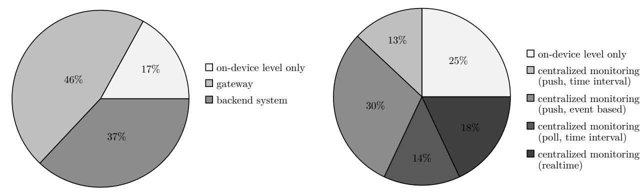

(a) Data gathering/aggregation

(b) Data monitoring platform

Fig. 7. Monitoring

Internet of Things domain remains to be highly heterogeneous. While semantics are expected to play a role in future IoT systems, there is still a way to go. Most participants see some benefit in the semantic management of things, devices and services. Nonetheless, when looking into the actual situation and the planned usage of semantics in IoT, these benefits seem not to be strong enough to stimulate large scale usage in the near future. One possible reason here could be a lack of training in semantics and the more "bit and byte"-oriented skillset current embedded developers have.

It is obvious that semantic management of services needs to take the different protocols into account, even if there seems to be the expectation that IPv6/6LoWPAN will play a crucial role and finally make the Internet of Things vision a reality. CoAP, while currently not used at a large scale, is expected to be for IoT what was HTTP for the WWW. Nonetheless, the number of people using and still expecting the use of custom protocols in the future is quite high and gateways or proxies will still be widely used. It seems as if a convergence towards an internet standard might not happen as soon as expected. Semantic service descriptions could fill that gap and allow integration of different protocols. SSN seems to have emerged to a de-facto standard for research in sensor networks. It was also one of the very few mature enough choices available when we sketched our own integration platform, a view that seems to be shared by others. The management of things and corresponding functionality (like discovery) seemed most promising to the participants of our study, followed by devices and services. Here, in our opinion, the community has still a way ahead before it stands on a common ground.

Our survey is only one small piece towards a qualitative and quantitative understanding of real-world usage in the emerging field of the Internet of Things. Not only in terms of the use of semantics, but also in the use of protocols and the management of things, services and devices. The authors suggest further empiric work to broaden the databasis and deepen the understanding of used protocols, needed management functionality and problems arising in real IoT deployments. Especially, the first industrial deployments of semantic platforms will lead to further insights into the real problems arising when using semantic technologies and if semantics can really hold its promises.

ACKNOWLEDGMENT

Would like to thank all participants in the study for their valuable feedback. Furthermore, the authors would like to thank Martin Zabel for his help in finalizing the paper.

REFERENCES

[1] "European Future Internet Portal." [Online]. Available: http://www.future-internet.eu/

[2] N. Gershenfeld, R. Krikorian, and D. Cohen, "The Internet of things." *Scientific American*, vol. 291, no. 4, p. 76, 2004.

[3] W. Dangelmaier, C. Laroque, and A. Klaas, "Industrie 4.0– Handlungsfelder der Digitalen Fabrik zur Optimierung der Ressourceneffizienz in der Produktion."

[4] "Sensei - Integrating the Physical with the Digital World of the Network of the Future." [Online]. Available: http://www.sensei-project.eu/

[5] "Internet of Things Architecture (IOT-A)." [Online]. Available: http://www.iot-a.eu/public

[6] M. Thoma, K. Sperner, T. Braun, and C. Magerkurth, "Integration of WSNs into enterprise systems based on semantic physical business entities," in *Wireless Days (WD), 2013 IFIP*, 2013.

[7] M. Thoma, T. Braun, and C. Magerkurth, "Enterprise Integration of Smart Objects using Semantic Service Descriptions," in *Proc. of IEEE Wireless Communication and Networking Conference (WCNC 2014), Istanbul, Turkey*, 2014.

[8] M. Thoma, T. Braun, K. Sperner, and C. Magerkurth, "Linked USDL4IoT," Tech. Rep., 2014. [Online]. Available: http://www.iot4bpm.de

[9] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee, "Hypertext transfer protocol–HTTP/1.1," 1999.

[10] D. Box, D. Ehnebuske, G. Kakivaya, A. Layman, N. Mendelsohn, H. F. Nielsen, S. Thatte, and D. Winer, "Simple object access protocol (soap) 1.1," 2000.

[11] N. Glombitza, D. Pfisterer, and S. Fischer, "Integrating wireless sensor networks into web service-based business processes," in *Proceedings of the 4th International Workshop on Middleware Tools, Services and Run-Time Support for Sensor Networks*, ser. MidSens '09. New York, NY, USA: ACM, 2009, pp. 25–30.

[12] Z. Shelby, K. Hartke, C. Bormann, and B. Frank, "Constrained Application Protocol (CoAP), Internet-Draft," 2011.

[13] M. Thoma, T. Kakantousis, and T. Braun, "REST-based sensor networks with OData," in *Wireless On-demand Network Systems and Services (WONS), 2014 11th Annual Conference on*. IEEE, 2014.

[14] M. Compton, P. Barnaghi, L. Bermudez, R. García-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, pp. 25–32, 2012.

[15] M. Botts and A. Robin, "Opengis sensor model language (sensorml) implementation specification," *OpenGIS Implementation Specification OGC*, pp. 07–000, 2007.

[16] M. Thron, *EDDL Elektronic Device Description Language*. Oldenbourg Industrieverla, 2011.

[17] U. Hunkeler, H. L. Truong, and A. Stanford-Clark, "MQTT-SA publish/subscribe protocol for Wireless Sensor Networks," in *Communication Systems Software and Middleware and Workshops, 2008. COMSWARE 2008. 3rd International Conference on*. IEEE, 2008, pp. 791–798.

[18] P. Fremantle, S. Weerawarana, and R. Khalaf, "Enterprise services," *Communications of the ACM*, vol. 45, no. 10, pp. 77–82, 2002.

[19] J. Case, M. Fedor, M. Schoffstall, and C. Davin, *A simple network management protocol (SNMP)*. Network Information Center, SRI International, 1989.

[20] U. Warrier, L. Besaw, L. LaBarre, and B. Handspicker, "Common management information services and protocols for the internet (CMOT and CMIP)," *Request for Comments*, vol. 1189, 1990.

[21] S. Kuryla and J. Schönwälder, "Evaluation of the resource requirements of SNMP agents on constrained devices," in *Managing the Dynamics of Networks and Services*. Springer, 2011, pp. 100–111.

[22] J. Schönwälder, H. Mukhtar, S. Joo, and K. Kim, "SNMP optimizations for constrained devices," *IETF draft draft-hamid-6lowpan-snmpoptimizations-03. txt, IETF*, 2010.

[23] A. Sehgal, V. Perelman, S. Kuryla, and J. Schonwalder, "Management of resource constrained devices in the internet of things," *Communications Magazine, IEEE*, vol. 50, no. 12, pp. 144–149, 2012.

[24] O. Mazhelis, M. Waldburger, G. S. Machado, B. Stiller, and P. Tyrväinen, "Retrieving monitoring and accounting information from constrained devices in internet-of-things applications," in *Emerging Management Mechanisms for the Future Internet*. Springer, 2013, pp. 136–147.

[25] "Auto id labs." [Online]. Available: http://www.autoidlabs.org/