Web Technologies and Privacy Policies for the Smart Grid

Sebastian Speiser, Andreas Wagner, Oliver Raabe, Andreas Harth
Karlsruhe Institute of Technology (KIT), Germany
{speiser,a.wagner,raabe,harth}@kit.edu

***Abstract*—The Smart Grid aims at making the current electricity grid more efficient, featuring an IT-layer that includes communication flows between a multitude of stakeholders. Its infrastructure is likely to be integrated with other IT-based systems. We argue that in order to enable the integration between already existing large-scale information networks (e.g., the Internet or the Web) and newly developed IT infrastructures such as Smart Cities, means for machine-understandable representation of data are required. Since these "smart" systems will exchange large volumes of highly sensitive data (billing and personal data), users should have the ability to specify their intent on how their data can be shared. Thus, systems require built-in mechanisms for protecting data. In this paper, we outline a Smart Grid architecture based on Semantic Web (Linked Data) technologies, and present mechanisms to allow for automated access control and enforcement of privacy rights on a technical level.**

***Index Terms*—Semantic Web, Privacy, Smart Grids, Information exchange.**

I. Introduction

The Smart Grid – a radical redesign of the energy grid – aims at profoundly changing the way how energy is created, distributed and consumed and promises to save considerable amounts of energy [1], [2]. An architecture for the Smart Grid has to be (1) flexible, i.e., fulfil customer requirements, but also allow for future extensions, (2) accessible, i.e., allow access to all market participants, (3) reliable, i.e., assure quality of supply, and (4) economic, i.e., provide best value and allow for innovation and competition [1]. In other words, the energy grid has to become more "like the Internet"[1] and allow for open data access.

As a result, more energy consumption data is being made available for billing, but also for improving the grid's quality and efficiency. In the brave new energy world not one stakeholder will have control over the grid (and the communication flows within), but many – energy providers, technology vendors, service companies, customers. These actors need to organise data exchange, e.g., power consumption data for billing and planning. In addition, car and appliance manufacturers have the opportunity to collect detailed data about the day-to-day usage of their products to improve their design. Further, as factories and urban infrastructures are being upgraded with information technology, all these systems have to be connected and their data has to be integrated with energy grid data.

Requirements. Given the Smart Grid vision, we can derive requirements for a data infrastructure [3], [4]:

- A flexible, open and light-weight data access is needed to enable seamless communication between market participants, leading to novel products and services. Standards only available under restrictive licenses to a selected number of market actors or over-specified regulations might stifle innovation. Thus, communication standards should be open to facilitate introduction of new products and methods, and to lower the barrier for new actors entering the market.
- New roles and processes within the Smart Grid require flexible data models, which enable a distinction between syntactic and semantic content. Thus, access methods and data formats should be high-level and support for data integration.
- Users should have the power to decide what data in which granularity to expose to whom.[2] In other words, a key requirement for the Smart Grid is to preserve data privacy. An architecture should allow for tight integration of legal aspects concerning data exchange and sharing, i.e., mechanisms allowing users to formulate their legal intentions in a machine-readable manner and enable an automated legal inference logic regulating the data usage.

Especially with regard to a technical enforcement of data privacy, a semantic data representation is key. Currently, however, implemented standards are either based on restricted technologies (such as EDIFACT) or complex service-oriented architectures, with no formal semantics associated. In particular, one proposed data format for the German Smart Grid is EDIFACT/MSCONS, which hinders a realisation of a privacy-aware system, as EDIFACT is not self-describing. Also, MSCONS lacks support for qualified signature procedures and privacy-related access authorisation.

Driven by Semantic Sensor Networks[3], e.g., [5], [6], and the Internet of Things, e.g., [7], [8], as well as recent works on Smart Grid infrastructure analyses, e.g., [3], [4], we suggest an architecture employing Web technologies. Such standards are widely used, well-known, and available under royalty-free licensing. In addition, Semantic Web technologies enable data publishing and integration in large, distributed environments. That is, their schema-less and self-describing nature facilitates flexible data integration, and allows for data privacy reinforcements via policies.

Contributions. Our contributions are: (1) We show how Semantic Web and *Linked Data technologies* may provide an infrastructure for data exchange and integration in the Smart Grid. (2) We specify a language for *access policies on data* so that users retain control over their private data, thereby providing the means for a privacy-aware framework.

Outline. First, we introduce Linked Data in Sect. II, followed by an example illustrating data access in Sect. III. Sect.

[1] *Building the energy Internet*, The Economist, 11th of May, 2004.

[2] "Informationelle Selbstbestimming" in German law.

[3] http://www.w3.org/2005/Incubator/ssn/XGR-ssn/

IV details the access process with policies and introduces our policy model. Policy matching is explained in Sect. V, and evaluated in Sect. VI. We cover related work in Sect. VII, and conclude with Sect. VIII.

II. URIs, HTTP, RDF(S) and Linked Data

Web technologies have proven suitable for large distributed IT-systems. Prominent examples include the Internet of Things [7], [8] or Semantic Sensor Networks, e.g., [5], [6].

URIs and HTTP. Resource addressing on the Web is based on Uniform Resource Identifiers (URIs) [9], which can be used to identify both real-world entities – such as a person, https://smartmeter.example.org/data#mary, or a car, http://car.example.org/data#uamp760e – and digital artifacts, e.g., a document. TCP/IP is applied as communication protocol stack, and HTTP is employed for data transfer. HTTP is stateless: client and server do not need to maintain a permanent connection. A client performs a lookup on a URI (request) and the respective server returns a piece of content back (response). In case state is required, e.g., to track a client's interaction history with a server, cookies can be used.

RDF. URIs and HTTP specify how to access content. To encode data in the Smart Grid, we propose the Resource Description Framework (RDF) [10] – resembling a directed data graph. A RDF graph comprises a set of triples: $\{\langle s, p, o\rangle\}$. Each triple associates an entity (subject) s with an object o via a predicate p, with s and p as URIs, and o as either URI or literal, e.g., a string or a numerical value.

There are different syntaxes for RDF. In this work, we will make use of the N3 notation, as it is human-friendly readable. For a full description of N3 see [11]. The N3 syntax basically is a list of triples separated by dots ("`.`"). To make this paper self-contained, we introduce the basic syntactic N3 primitives. Brackets (<>) denote URIs, quotes ("") denote literals (such as strings or integers), and blank node identifiers start with "_ :", denoting anonymous resources. There exists a number of syntactic shortcuts, for example ";" to introduce another predicate and object for the same subject. Namespaces can be introduced with the @prefix keyword.

RDF Schema (RDFS) adds additional expressivity in order to support the design of simple vocabularies encoded in RDF, cf. [12]. In particular, the predefined resources `rdfs:Class`, `rdfs:Resource` and `rdf:Property` may be used to model concepts (classes), roles (predicates) and resources. Furthermore RDFS introduces the following predefined properties: `rdf:type` as means for an instance-of relationship, `rdfs:subClassOf` for stating a subclass-of relation and `rdfs:subPropertyOf` for defining a subproperty-of dependency between two properties.

Vocabularies. Various RDF vocabularies have been defined in the Semantic Web.[4] For instance, in our scenario, Sect. III, we employ one vocabulary for temporal and another one for geographical data, cf. Table I. Further, multiple vocabularies for the Smart Grid have been proposed as RDF, e.g., Gridpedia[5] or [13]. Most notably, the Common Information Model (CIM) has been published in RDF [14], [15]. Note that RDF allows for easy integration of multiple vocabularies.

Linked Data. Linked Data (LD) principles dictate *how to publish* RDF(S) data, so that one may easily relate and integrate data from varying sources. The LD principles are [16]: (1) Identify entities via HTTP-URIs. (2) If someone looks up an HTTP-URI, useful RDF about that entity should be returned. (3) A resolved entity description should provide links to other entities. We argue that publishing Smart Grid RDF data as Linked Data is a good fit – as illustrated by the following scenario. Notice, while we restrict attention to read-only access, current LD efforts aim at read and write access.[6]

III. Linked Data in the Smart Grid

In this section, we outline a scenario illustrating how Linked Data enables innovative use of accumulated data. The scenario also describes how data sharing in a decoupled energy market is supported. Further, we show that such a scenario requires policies that allow the expression of data sharing restrictions specified by individuals or the law. We distinguish between data associated with *legal consequences*, in particular data required for billing, and all other data. The former data is referred to as *obligatory data*, while the latter is simply *non-obligatory data*. Non-obligatory data may be managed by the smart meter (for devices which do not have processing power) or the device itself. Obligatory data, however, must be stored by a trusted instance, e.g., the metering provider. The distinction is of importance as for obligatory data legal regulations specify publishing and data availability requirements. Thus, a trusted environment is necessary.

Scenario. Consider a person Mary, who lives at an apartment fitted with a smart meter; Mary owns a CoolWash washer and an UltraAmp 760e electric car, amongst other devices. The scenario is illustrated in Figure 1. For our scenario we require that these devices are accessible via TCP/IP and have the following hostnames: smartmeter.example.org, washer.example.org, and car.example.org. We assume that each device is accessible via HTTP, and URIs identify each resource (a Person, a Vehicle, or an Appliance). The URI denoting Mary, for example, is http://smartmeter.example.org/data#mary. Please note that we use the namespace definitions listed in Table I for brevity.

If we perform a lookup on Mary's URI, the server (the smart meter) returns an RDF data describing the resource:

```
sm:mary rdf:type foaf:Person ;
  foaf:name "Mary Doe" ;
  foaf:based_near sm:apt .
```

A request on `sm:apt` returns more data pertaining to the premise (such as latitude and longitude or address).

Requests on the URI of the washer `washer:w` provide data describing the appliance, including links to energy con-

[4] http://www.w3.org/wiki/TaskForces/CommunityProjects/LinkingOpenData/CommonVocabularies

[5] http://gridpedia.org

[6] http://www.w3.org/DesignIssues/ReadWriteLinkedData.html

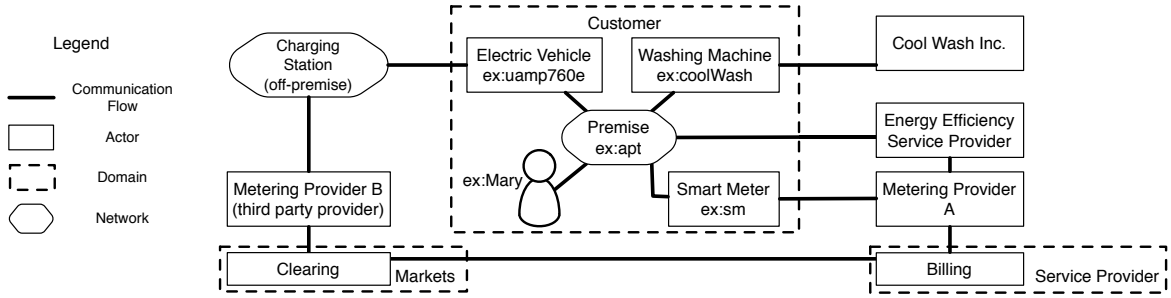

Fig. 1. Actors and their interactions in the Smart Grid (adapted from [2]).

Prefix	*URI*	*Description*
sm	https://smartmeter.example.org/data#	Smart Meter Data
washer	http://washer.example.org/data#	Washer Data
car	http://car.example.org/data#	Car Data
cw	http://coolwashinc.example.org/data#	CoolWash Inc Data
sg	http://smartgrid.example.org/vocab#	Smart Grid Vocab.
p	http://policy.example.org/vocab#	Policy Vocab.
xsd	http://www.w3.org/2001/XMLSchema#	XML Schema Vocab.
rdf	http://www.w3.org/1999/02/22-rdf-syntax-ns#	RDF Vocab.
rdfs	http://www.w3.org/2000/01/rdf-schema#	RDF Schema Vocab.
foaf	http://xmlns.com/foaf/0.1/	Person Vocab.
geo	http://www.w3.org/2003/01/geo/wgs84_pos#	Geo-location Vocab.
ical	http://www.w3.org/2002/12/cal/ical#	Temporal Vocab.

TABLE I
NAMESPACES AND VOCABULARIES USED IN OUR SCENARIO.

sumption data and data about the previously selected washing programs:

```
washer:w
  rdf:type sg:Appliance ;
  sg:manufacturer cw:company ;
  sg:owner sm:mary ;
  cw:washingData washer:program40 ;
  sg:consumption sm:data20100310 .
```

A lookup on `washer:program40` returns:

```
washer:program40
  rdf:type cw:WashingData ;
  foaf:name "Program 40 C" ;
  cw:totalCount "23"^^xsd:int .
```

The energy usage data resides at the metering system, so performing a lookup on `sm:usage2010031100` results in the following data snippet, indicating a consumption of 1.04 kWh during a late-night wash:

```
sm:data20100310
  rdf:type sg:Consumption ;
  rdf:value "1.04"^^sg:kWh ;
  ical:dtstart "2010-03-10T00:00:00" ;
  ical:dtend   "2010-03-10T01:00:00" .
```

In contrast to on-premise appliances, the *UltraAmp 760e* is mobile. We assume a TCP/IP connection to the car (e.g., via 3G), so requests may be performed as well. A HTTP lookup on the URI of the car `car:uamp760e` may provide the model description and current location:

```
  car:uamp760e rdf:type sg:Vehicle ;
      foaf:name "UltraAmp 760e" .
      geo:location _:loc20100331 .
  _:loc20100331 dc:date "2010-03-31T12:23:45";
      geo:lat "49.0047222" ;
      geo:lon "8.3858333" .
```

Note, a lookup targets a device directly, rather than requiring a centralised location which warehouses all data. We assume that access to the metering system is done via an encrypted channel (e.g., `https`), and recording of consumption data adheres to legal requirements.

We assume that the manufacturer of the *CoolWash* machine wants to request data about the washing machine (to, e.g., optimise future versions of the appliance based on real-world usage), a metering system provider wants to request power consumption data (for billing), and an energy optimisation consultancy wants to request all energy-related data (to help Mary optimise her energy consumption).

Data sharing is essential in a decoupled market scenario, where solely a metering system provider (MSP) has access to a customer's detailed usage data and provides an energy utility upon request with aggregated consumption data. Data sharing is also needed for roaming, as an electric vehicle may be charged at an off-premise charging station. The electric vehicle identifies itself at the charging station, and the power consumption data is sent to the customer's utility company for billing purposes. Summarising the scenario, there is a data flow which includes sensitive, fine-grained usage data and possibly additional personal data such as location data.

Policies. In the next sections, we specify policies that enable customers to articulate their legal intent in a machine-understandable manner. We distinguish two policy types:

- *Policies specified by a private party*: Consider customers, who wish to allow the manufacturer of their washing machine or their electric car access to usage data to help them improve their products. Such access is not mandated by law, however, a customer may want to allow it.
- *Policies specified by law*: Consider an energy provider, who requires access to data related to billing issues. Such access is enforced by law to ensure contract fulfillment.

IV. POLICY MODEL

Policy-aware data access begins with the consumers, who deploy private policies on their devices. As a special case the smart meter may serve as a gateway for incoming data requests to low-powered devices. The policies are also sent along with obligatory data, e.g., to metering providers.

A requestor sends along with a request a statement of identity and purpose of the data usage. Focusing on the

access chain, the appropriate handler first matches the request statement with the policies regulating the specific data and decides whether the access is allowed, and to which level the returned data is aggregated (e.g., time-wise). According to this decision the corresponding data, respectively an error code, is returned to the requestor. The data is accompanied by a policy specifying the terms of agreement.

Optionally signatures may be used to avoid attacks (e.g., "man-in-the-middle" attack): firstly, the data requestor could sign his statement regarding identity and purpose, such that in case of a violation, the data provider can prove it. Secondly, the data provider signs the combination of the returned policy and the hash value of the returned data. This way the data requestor can prove that his use of data was rightful, if he can present the appropriate signed policy and hash value. Note, RDF hashing techniques have been proposed in, e.g., [17].

Our model is similar to the traditional contractual model. A request as well as a policy may be seen as declaration of intent, therefore a matched policy represents a contract (if signatures are applied). This means that our model is an automated version of well-known and well-founded principles in the existing legal framework.

Our policy model gives users control over their data by enabling them to formulate their intents in machine-readable manner and thus allowing them to restrict or permit agents data access. However, allowing access to some data does not imply access to all data. We propose that users are allowed to define various data *perspectives*. A perspective is defined as a SPARQL[7] query. Using the CONSTRUCT operator a new graph can be defined depending on the original graph. This means that triples matching certain criteria, specified in the WHERE and FILTER clauses can be included or excluded from the graph of the perspective.

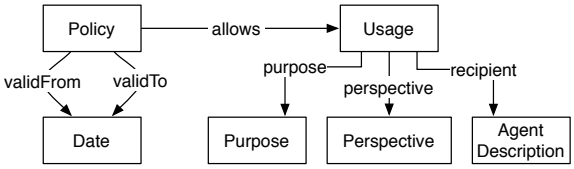

Fig. 2. Conceptual Model of Usage and Policies

Based on such perspectives we can describe a policy model as visualised in Figure 2. Here, by introducing `validFrom` and `validTo` a `Policy` models a timespan during which it is valid. Furthermore, it `allows` a number of `Usages`. An allowed usage applies to the data that is available by means of its `perspective`. An allowed usage is restricted to a specific `purpose` and to a `recipient`, characterised by an `AgentDescripton`, which can require that an actor (1) belongs to a specific class, e.g., a NGO, or (2) is a specific individual, e.g., Cool Wash Inc.

Our approach focuses on the expression of a user's intent (and matching the intent with an incoming request), but could be extended to enforce access control by verifying the requestor's authentication. This could, e.g., be realised by requiring that each requestor presents a logical proof of his identity, based on axioms from trusted sources, cf. [18].

[7]http://www.w3.org/TR/rdf-sparql-query/

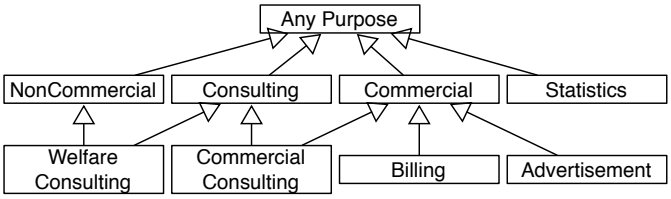

Fig. 3. Example Taxonomy of Purposes

Especially with respect to *purposes*, it makes sense to refer to URIs defined by a trusted third-party organisation, which provide useful and reliable definitions for different purposes. Note that this model is comparable to the Creative Commons approach, where, e.g., a non-commercial clause is defined that can be referenced by different licenses. An example taxonomy of purposes is illustrated in Figure 3.

The following is an example policy for the washer that provides full access to Mary and restricted access to the manufacturer (omitting consumption data).

```
washer:pol1
  rdf:type p:Policy ;
  ical:dtstart "2010-01-01" ;
  ical:dtend "2014-12-31" ;
  p:allows washer:fullaccess ;
  p:allows washer:access .

washer:fullaccess
  rdf:type p:Usage ;
  p:purpose purpose:Any ;
  p:recipient mary:i ;
  p:perspective washer:fullperspective .

washer:fullperspective
  p:definition """
    CONSTRUCT { ?s ?p ?o . }
    WHERE { ?s ?p ?o . }""" .

washer:access
  rdf:type p:Usage ;
  p:purpose purpose:Consulting ;
  p:recipient cw:company ;
  p:perspective washer:cwperspective .

washer:cwperspective
  p:definition """
   CONSTRUCT { ?s ?p ?o . }
   WHERE { ?s a sg:Appliance .
    ?s sg:manufacturer cw:company .
    ?s ?p ?o .
    FILTER (?p != sg:consumption) }""" .
```

V. Policy Matching

In this section, we describe how a data request can be matched with policies. A request is sent via HTTP POST including RDF data encoding identity of the requestor and the purpose for the data will be used. Consider CoolWash Inc requesting data from Mary's washer `washer:w`:

```
:req
   rdf:type sg:Request ;
   p:purpose p:Consulting ;
   p:recipient cw:company .
```

The rules matching requests with policies are implemented as SPARQL queries, using the CONSTRUCT operator to define new triples based on the conditions in the WHERE clause. Answering a request involves:

- The data provider of the corresponding URI approves the request, if there is a law or private policy that allows a matching usage. If there is no matching usage, the request is assumed to be forbidden, encoded in the following rule:
```
CONSTRUCT { ?r p:allowedBy ?u } WHERE
  { ?p rdf:type p:Policy .
    ?p p:allows ?u . ?u p:matches ?r . }
```
- A request matches an allowed usage, if the recipient of the request is the same as the allowed recipient and its purpose is the same as or a subclass of (i.e., more specialised) the purpose of the allowed usage. We encode the transitivity property of the subclass relation as a rule, and compute a fixpoint over the data to materialise the entire subclass hierarchy. The formal matching rule is as following:
```
CONSTRUCT { ?u p:matches ?r }
WHERE {
  ?u rdf:type p:Usage .
  ?r rdf:type p:Request .
  ?r p:purpose ?rp .
  ?u p:purpose ?up .
  ?rp rdfs:subClassOf ?up .
  ?r p:recipient ?rec .
  ?u p:recipient ?rec . }
```
- If a matching usage is found, then the perspective of the usage is applied to the data describing the requested URI. If all data is filtered out by the perspective or no matching allowed usage was found, an HTTP 404 (not found) error is returned to the requestor. Thus, an unauthorised requestor cannot infer whether a URI exists for which he was not authorised.
- Otherwise the data filtered by the perspective is returned to the requestor together with a policy including the allowed usage that matched the request.
- Optionally digital signatures can be applied as described in Section IV.

The dependencies of the different knowledge pieces used for the matching process are depicted in Figure 4.

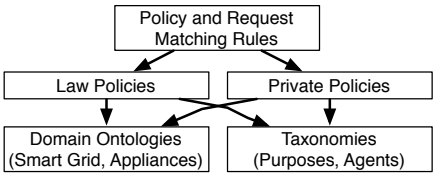

Fig. 4. Dependencies for Policy Matching

As requested purpose and recipient equal those defined by Mary's `washer:access`, the request is allowed with the corresponding perspective applied, resulting in:

```
washer:w
  rdf:type sg:Appliance ;
  sg:manufacturer cw:company ;
  sg:owner sm:mary ;
  cw:washingData washer:program40 ;
```

Note that the energy consumption data has been filtered out. The data is accompanied by a policy with a single allowed usage: `washer:access`.

VI. Evaluation

Setting. For evaluation we implemented a policy matcher based on Rasqal[8], which is a lightweight SPARQL engine programmed in C. We ran the policy matcher on two different hardware platforms: (1) a 2.4 GHz Core2Duo laptop with 4 GB RAM, and (2) a SheevaPlug device with an 1.2 GHz ARM processor and 512 MB RAM. While the first computer is representative for data servers, like they could be used by metering providers, the second platform represents a low power device, e.g., a washing machine. *The goal of our experiments is to show that matching of policies with varying size can be done efficiently.*

Results. For the experiment we created 100 different allowed usages by specifying combinations of recipients and purposes, modeled via taxonomies as shown in Figure 3. In the same way, we defined 10 different requests. We created policies, where the number of allowed usages varied between 1 and 75 in steps of 5. For each size, 10 policies with randomly drawn allowed usages were matched against every request[9]. We measured the average time for matching a request against a policy of a specific size. We distinguished between allowed and denied requests.

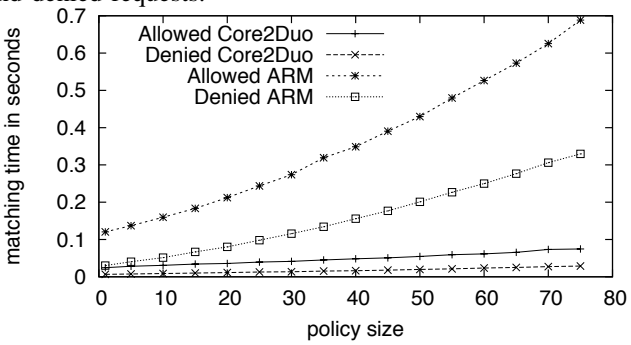

Fig. 5. Performance Measurements: Policy Matching

Figure 5 shows the matching times for both hardware platforms. We assume that a typical policy on an appliance is not larger than 10 allowed usages. Our measurements indicate that even for larger numbers of allowed usages, several requests per second can be matched by the ARM processor. Note that our implementation does not use any optimisation methods such as indexing, caching or multithreading. For realistic policies (i.e., less than 15 usages) the Core2Duo system can deny 100 unauthorised requests per second. The times for denying requests are significantly lower, which is useful as not much computing power has to be wasted for unauthorized and probably malicious requests.

[8]http://librdf.org/rasqal/

[9]Test data and source code are available at http://code.google.com/p/polen/.

VII. Related Work

There have been proposals for a Smart Grid communication infrastructure, e.g., [2], [4]. In fact, previous approaches also consider (Semantic) Web technologies to be applicable, e.g., [3], [19], [20]. Further, several RDF Smart Grid vocabularies have been proposed, e.g., [13], [15]. In these works, data privacy and protection is either not addressed or merely mentioned as future problem. On the other hand, there is work targeting these issues, e.g., [21], [22]. However, such approaches outline organisational/legislative measures, e.g., adjustment of privacy laws. In contrast, we aim at a technical solution, i.e., an automated enforcement of data privacy. Technical reinforcements have been proposed for specific aspects of the Smart Grid, e.g., metering (cf. [23], [24]) or vehicle-to-grid communication (cf. [25]). We argue, however, that data privacy should be enforced throughout the entire grid. In this work, we show how Linked Data principles are a good fit here.

In our previous works [26], [27], we address data privacy and illustrate legal use-cases as well as issues within a Smart Grid. However, while [26], [27] also emphasise the importance of a technical enforcement of privacy, the works primarily aim at legal questions. In contrast, this paper focuses more on a technical perspective.

Related to our policy model is work on access control for RDF, e.g., [28]. However, such approaches restrict initial data access and not its ongoing usage. [29] developed a data-purpose algebra that enables data usage modeling. Their work focuses on the verification of processes, whereas our work addresses policy expression and enforcement. Privacy policies can also be expressed using XML languages XACML [30] and EPAL [31]. EPAL's hierarchical types, which are used for purposes, define a schema for creating taxonomies, but it's lacking a formal foundation or support for linking different vocabularies. In fact, both languages rely on proprietary XML schema with semantics given by natural language documents. Our approach has the advantage that matching is defined in terms of simple rules based on formal logics. In [32], personal infospheres are introduced, which can be used to specify what data may be shared with whom. Our policy model can be seen as instantiation of this infosphere model.

Last, various works targeted a formal modeling of laws or legal aspects using Semantic Web technologies such as OWL, e.g., [33]. However, we aim a lightweight and simple model only capturing few concepts relevant for the user's data. This way, we can achieve an efficient and scalable policy matching.

VIII. Conclusion and Future Work

We outlined a Smart Grid communication based on Semantic Web technologies. In particular, we illustrated how Linked Data principles may be applied. Based on this architecture, we proposed and evaluated a lightweight policy approach allowing a technical reinforcement of data privacy.

We will extend the matching rules to allow more complex inferences, and evaluate the system in a larger setting. We plan to add cryptographic procedures to ensure the identity of participants and authenticity of policies. Further, we want to support active controlling of devices – actuators have to be taken into account to moderate energy demand.

References

[1] "European Technology Platform - SmartGrids Vision and Strategy for Europes Electricity Networks of the Future," European Comission, 2006.

[2] "NIST Framework and Roadmap for Smart Grid Interoperability Standards, Rel. 2.0," National Institute of Standards and Technology, 2012.

[3] A. Wagner, S. Speiser, and A. Harth, "Semantic web technologies for a smart energy grid: Requirements and challenges," in *ISWC*, 2010.

[4] S. Rohjans, C. Danekas, and M. Uslar, "Requirements for Smart Grid ICT-architectures," in *ISGT*, 2012.

[5] Compton et al., "The SSN Ontology of the W3C Semantic Sensor Network Incubator Group," *JWS*, vol. 17, no. 0, 2012.

[6] A. Sheth, C. Henson, and S. Sahoo, "Semantic sensor web," *Internet Computing, IEEE*, vol. 12, no. 4, pp. 78–83, 2008.

[7] L. Atzori, A. Iera, and G. Morabito, "The Internet of Things: A survey," *Computer Networks*, vol. 54, no. 15, pp. 2787 – 2805, 2010.

[8] Pfisterer, D. et al., "SPITFIRE: toward a semantic web of things," *Communications Magazine, IEEE*, vol. 49, no. 11, pp. 40–48, 2011.

[9] "URIs, URLs, and URNs: Clarifications and Recommendations 1.0," Joint W3C/IETF URI Planning Interest Group, W3C, Tech. Rep., 2001.

[10] *Resource Description Framework (RDF): Concepts and Abstract Syntax*. W3C Recommendation, 2004, http://www.w3.org/TR/rdf-concepts/.

[11] T. Berners-Lee, "Notation 3 – Ideas about Web architecture," Tech. Rep. [Online]. Available: http://www.w3.org/DesignIssues/Notation3

[12] *RDF Vocabulary Description Language 1.0: RDF Schema*. W3C Recommendation, 2004, http://www.w3.org/TR/rdf-schema/.

[13] Zhou Q. et al., "Semantic Information Modeling for Emerging Applications in Smart Grid," in *ITNG*, 2012.

[14] Uslar, M. et al., *The Common Information Model CIM: IEC 61968/61970 and 62325 - A practical introduction to the CIM*, ser. Power Systems, 2012.

[15] Y. Penya, A. Pena, and O. Esteban, "Semantic integration of IEC 60870 into CIM," in *SmartGridComm*, 2011.

[16] T. Berners-Lee, "Linked data design issues," W3C, Tech. Rep., 2006, http://www.w3.org/DesignIssues/LinkedData.html.

[17] G. Tummarello, C. Morbidoni, P. Puliti, and F. Piazza, "Signing individual fragments of an RDF graph," in *WWW*, 2005.

[18] D. J. Weitzner, J. Hendler, T. Berners-Lee, and D. Connolly, *Creating a Policy-Aware Web : Discretionary , Rule-based Access for the World Wide Web*. IRM Press, 2005, ch. I, pp. 1–31.

[19] D. Rech and A. Harth, "Towards a decentralised hierarchical architecture for Smart Grids," in *EDBT/ICDT Workshops*, 2012.

[20] S. Rohjans, "A standard-compliant ICT-architecture for semantic data service integration in Smart Grids," in *ISGT*, 2013.

[21] P. McDaniel and S. McLaughlin, "Security and Privacy Challenges in the Smart Grid," *IEEE Security and Privacy*, vol. 7, pp. 75–77, 2009.

[22] A. Cavoukian, J. Polonetsky, and C. Wolf, "SmartPrivacy for the Smart Grid: embedding privacy into the design of electricity conservation," *Identity in the Information Society*, vol. 3, no. 2, pp. 275–294, 2010.

[23] Rial, A. et al., "Privacy-preserving smart metering," in *WPES*, 2011.

[24] J. Strüker and F. Kerschbaum, "From a Barrier to a Bridge: Data-Privacy in Deregulated Smart Grids," in *ICIS*, 2012.

[25] M. Stegelmann and D. Kesdogan, "V2GPriv: vehicle-to-grid privacy in the Smart Grid," in *CSS*, 2012.

[26] O. Raabe, "Datenschutz im SmartGrid," *Datenschutz und Datensicherheit*, 2010.

[27] A. Wagner, S. Speiser, O. Raabe, and A. Harth, "Linked data for a privacy-aware smart grid," in *GI Jahrestagung*, 2010.

[28] Abel F. et al., "Enabling Advanced and Context-Dependent Access Control in RDF Stores," in *The Semantic Web*, ser. Lecture Notes in Computer Science, vol. 4825, 2007, pp. 1–14.

[29] Hanson C. et al., "Data-Purpose Algebra: Modeling Data Usage Policies," in *POLICY*, 2007.

[30] *eXtensible Access Control Markup Language (XACML) Version 2.0*. OASIS Standard, 2005, http://docs.oasis-open.org/xacml/2.0/.

[31] *Enterprise Privacy Authorization Language (EPAL 1.2)*. W3C Member Submission, 2003, http://www.w3.org/Submission/2003/07/.

[32] A. Harth and R. S. (eds.), "Dagstuhl Perspectives Workshop Report: Semantic Web Reflections and Future Directions," Tech. Rep., 2010.

[33] Rubino R. et al., "An OWL Ontology of Fundamental Legal Concepts," in *JURIX*, 2006.