



Matching Offerings and Queries on an Internet of Things Marketplace

Victor Charpenay¹(✉), Hoan Nguyen², Mohamad Ibrahim³, Achille Zappa²,
and Arne Bröring¹

¹ Siemens AG — Corporate Technology, Munich, Germany
victor.charpenay@siemens.com

² Insight Centre for Data Analytics, National University of Ireland, Galway, Ireland

³ Technical University of Clausthal, Clausthal-Zellerfeld, Germany

Abstract. A marketplace for the Internet of Things acts as the corner stone of an IoT ecosystem, by matching the offer (i.e., data or functionalities) with the demand coming from IoT applications (e.g. analytics). In this paper, we present the semantic matching implemented on the public BIG IoT marketplace, accessible at <https://market.big-iot.org/>.

Today, many Internet of Things (IoT) platforms have come up and provide data and functionalities of *things*, e.g., ThingWorx, Xively or Siemens MindSphere. In order to enable a vibrant and collaborative IoT ecosystem across these platforms, *marketplaces* are needed to enable *providers* to monetize the access to their platforms by *consumers* (e.g., applications or services). The BIG IoT project [1] offers such a marketplace and enables providers to register their IoT resources as *offerings* and consumers to formulate *queries* to discover these offerings. Once offerings and queries are registered on the BIG IoT marketplace, it is crucial to effectively support the matching between offerings and queries, so that consumers are reported which offerings suit their needs in near real-time.

1 Offering Model

The basis for our matching approach is a lightweight ontological model for IoT offerings and queries, which we illustrate in Fig. 1. Platform providers must register their offerings in the form of RDF documents we call *offering descriptions* (ODs), which describe in detail the offered platform resources as per our ontology. The OD is based on the *thing description* from the W3C's Web of Things working group [2].

At the core of our model are the two classes **Offering** and **OfferingQuery**. They respectively extend schema.org's classes **Offer** and **Demand**. A direct benefit of aligning with the well-known schema.org vocabulary is that BIG IoT offerings and queries could be crawled by search engines as an alternative to being registered on a marketplace. Every offering has a mandatory endpoint definition that specifies an IRI as well as the communication protocol and message type

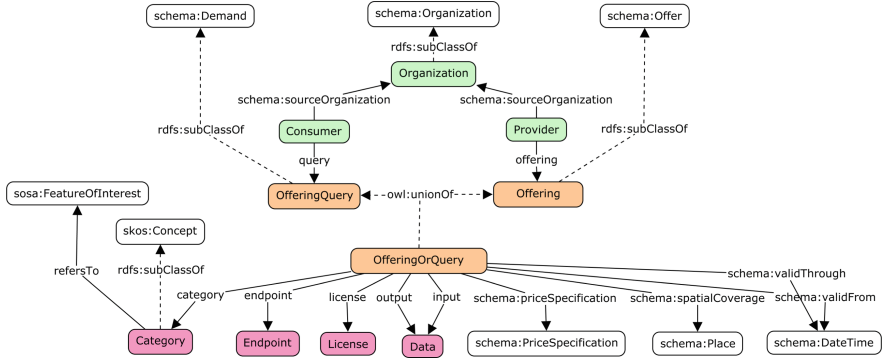


Fig. 1. Overview of the BIG IoT core ontology (<http://schema.big-iot.org/core/>).

used to access the resource (e.g. HTTP POST or CoAP GET). Queries may also define restrictions on the endpoint. To allow for large-scale data exchange, offerings and queries must provide input and output data definitions (inspired by JSON schema). These definitions are used both for matching and at access time to validate a consumer’s input and the provider’s output. Queries and offerings can also define a category, which is modeled as a SKOS concept in BIG IoT. Although categories are to be thought of as free-text tags defined by users (eventually becoming a folksonomy managed by the IoT community), we provide an initial category tree for the domains of mobility and environmental monitoring. Finally, queries and offerings can define non-functional properties such as spatial and temporal extent as well as price and license to refine the offering matching process [3].

RDF annotations in data definitions, as well as offering categories, are crucial for offering matching. In parallel to its core ontology, the BIG IoT marketplace also provides other cross-domain or domain-dependent ontologies, aligned both with schema.org and the widespread W3C SOSA ontology, designed to capture the semantics of IoT systems [4]. In particular, these ontologies extend the concept of `FeatureOfInterest`, e.g. by defining classes for concepts such as parking lot, traffic, or air pollution in a given area, which are used as a basis for the automatic derivation of a category tree for BIG IoT offerings.

2 Offering Matching

An important requirement that drives our approach is that results must remain intuitive to the various stakeholders of the BIG IoT marketplace. We therefore implemented a straightforward multi-modal boolean search algorithm, where for each aspect of the offering model presented above, the evaluation against an offering returns *true* or *false*. In practice, a SPARQL graph pattern is generated for each of these aspects from an offering query, the conjunction of which is evaluated against an RDF graph that includes all offerings. Figure 2 provides an

example of a SPARQL query generated for matching. We review all six aspects of matching in the following.

IoT data points are typically associated to a precise geographical location. The spatial extent of an offering represents either the actual location of a sensor or a bounding box for a set of data points. An offering matches as soon as the bounding boxes of an offering intersect with that of a query (1). Similarly, IoT data is located in time (temporal extent). An offering can provide historical data or live data. A query that includes a time range will match offerings with historical data that overlap with this range. If no temporal extent is given, offerings with live data only will match (2).

Offerings can also match against a category (3). Our model currently includes 38 categories, defined hierarchically¹. Subsumption reasoning is performed at registration time to add parent categories to an offering. For instance, `ParkingCategory` should also match offerings tagged with `ParkingSiteCategory` or `ParkingSpaceCategory`.

As mentioned in the previous section, offerings include data definitions annotated with RDF. For matching, we discard structural information and only keep the set of RDF terms (4). Matching occurs when the set of terms in an offering fully contains that of a query (for both input and output). As for categories, reasoning is performed at registration time to expand the set of terms with those found in `schema.org`, `SOSA` and our own ontologies (which include 64 classes and 74 properties so far). Super-classes and domain and range classes are added to the set of terms, as per RDFS semantics. Here, offerings annotated with

```

SELECT DISTINCT ?id
WHERE {
  ?offering a :Offering;
  :offeringId ?id;
  :isActive true^^xsd:boolean.

1  ?offering a :Offering;
   schema:spatialCoverage ?area.
   ?area wgs84:geometry ?geo.
   FILTER(bif:st_intersects(?geo, "...").

2  ?offering schema:validFrom ?validFrom;
   schema:validThrough ?validThrough.
   FILTER (?validFrom <= "1520812"^^xsd:long
    && ?validThrough >= "1518393"^^xsd:long)

3  ?offering schema:category
   <urn:big-iot:ParkingCategory>.

4  ?offering :input ?in;
   :output ?out.
   ?in :rdfAnnotation schema:GeoShape.
   ?out :rdfAnnotation mobility:ParkingSpace.

5  ?offering schema:priceSpecification ?spec.
   ?spec :pricingModel ?pModel;
   schema:priceCurrency "EUR"^^xsd:string;
   schema:price ?amount.
   FILTER(?amount <=0.02).
   VALUES ?pModel {
     :per_access_price
     :free_price
   }

6  ?offering :license ?license.
   ?license :licenseType :open_data_license.
}

```

Fig. 2. SPARQL query example

¹ See <http://big-iot.github.io/categories>.

`GeoCircle` (subclass of `GeoShape`) and `distanceFromParkingSpace` (whose domain includes `ParkingSpace`) will match.

Finally, offerings can also be matched against a price or a license. Our BIG IoT core model includes several pricing models: subscription, per access payment, etc. and prices given in a query are interpreted as *maximum* amounts (5). As for pricing models, our model includes several licenses in use on various Open Data platforms: Creative Commons, Open Government License, any commercial license and public domain. Offering licenses must be an exact match (6).

3 Related Work

In information science, matching supply with demand within a market has been extensively studied, especially in the context of Web services. In particular, many works have been conducted to allow for the matchmaking between user needs and services at a semantic level, as the development of ontological frameworks such as OWL-S [5] and WSMO [6] shows. The general idea behind semantic matchmaking is to find a subsumption relation between supply and demand (modeled as RDF classes) using automated reasoning [7].

One limitation to pure semantic matchmaking is the potentially high heterogeneity between the various semantic models used across service providers, as well as a potential mismatch in the semantics used by providers and consumers. To address this, ontology matching techniques exist (either based on syntactic, structural or semantic features) to align semantic models with each other [8]. However, most of these techniques require human input, which, in the case of an IoT marketplace, would hardly scale.

The offering model we present in Sect. 1 accounts for a lightweight alternative to OWL-S and WSMO suited for IoT data by integrating spatio-temporal extents, as well as license and price. In the three large-scale pilots in which BIG IoT has been deployed (Barcelona, Piedmont and northern Germany), we could observe a very high heterogeneity across platforms with respect to data modeling, which led us to implement a matching algorithm that does not rely on pure semantic matchmaking, while still leveraging semantic models.

4 Demonstration and Conclusion

Our demonstration consists of a walk through the public BIG IoT marketplace, available at <https://market.big-iot.org/> (see also Fig. 3). It currently contains 56 offering providers, providing 142 offerings, as well as 28 consumers. The number of offerings is still growing, as the project goes on. The data being offered on this marketplace is fairly diverse but offerings are mainly categorized in three domains: `EnvironmentalIndicatorCategory` (38), `MobilityFeatureCategory` (77) and `WeatherIndicatorCategory` (27). There are 12 offerings with the category `ParkingCategory` and charging less than 0.02€ per access, these two criteria being a subset of what is presented in Fig. 2. If we add the spatial and temporal extents, only 4 offerings remain.

Our demonstration should illustrate how our pragmatic approach towards semantic matchmaking already enables intelligent matching between offerings and queries of existing IoT platforms. This matching mostly exploits ontological models based on schema.org and SOSA that we developed for the domains of mobility and environmental monitoring.

In the coming months, the BIG IoT Marketplace and its APIs will be released as open source² and developers are encouraged to take part in this IoT ecosystem³. Further improvement of the matching algorithm is possible by leveraging structural information of data type definitions, which we plan in a future deployment of our marketplace. In parallel, we will further develop our domain models and plan to eventually contribute them to schema.org⁴.

Fig. 3. Parking query

Acknowledgements. This work is supported by the project ‘Bridging the Interoperability Gap of the Internet of Things’ BIG IoT funded by the European Commission’s Horizon 2020 program (Grant No. 688038).

References

1. Bröring, A., Schmid, S., Schindhelm, C.K., Khelil, A., Käbisch, S., Kramer, D., Phuoc, D.L., Mitic, J., Anicic, D., Teniente, E.: Enabling IoT ecosystems through platform interoperability. *IEEE Softw.* **34**(1), 54–61 (2017). <https://doi.org/10.1109/MS.2017.2>
2. Käbisch, S., Kamiya, T.: Web of Things (WoT) thing description. <https://www.w3.org/TR/wot-thing-description/>
3. Bröring, A., Ziller, A., Charpenay, V., Schmid, S., Thuluva, A., Anicic, D., Zappa, A., Linares, M., Mikkelsen, L., Seidel, C.: The BIG IoT API: Semantically Enabling Interoperability in IoT Ecosystems. *IEEE Pervasive Computing*, PCSI-2018-01-0008.R2 manuscript type: Special Issue - IoT Deployments (October/December 2018)
4. Haller, A., Janowicz, K., Cox, S., Le Phuoc, D., Taylor, K., Lefrançois, M.: Semantic sensor network ontology. <https://www.w3.org/TR/vocab-ssn/>
5. Martin, D., Burstein, M., Hobbs, J., Lassila, O., McDermott, D., McIlraith, S., Narayanan, S., Paolucci, M., Parsia, B., Payne, T., Sirin, E., Srinivasan, N., Sycara, K.: OWL-S: semantic markup for web services. <http://www.w3.org/Submission/OWL-S/>

² See <https://projects.eclipse.org/proposals/eclipse-bridge.iot>.

³ See our developer guide at <https://big-iot.github.io>.

⁴ See the <http://iot.schema.org> extension.

6. Roman, D., Keller, U., Lausen, H., De Bruijn, J., Lara, R., Stollberg, M., Polleres, A., Feier, C., Bussler, C., Fensel, D.: Web service modeling ontology. *Appl. Ontol.* **1**(1), 77–106 (2005)
7. Noia, T.D., Sciascio, E.D., Donini, F.M.: Semantic matchmaking as non-monotonic reasoning: a description logic approach. *J. Artif. Intell. Res.* **29**, 269–307 (2007)
8. Euzenat, J., Shvaiko, P.: *Ontology Matching*. Springer, Heidelberg (2007). <https://doi.org/10.1007/978-3-540-49612-0>