Linked Timelines: Temporal Representation and Management in Linked Data

Gianluca Correndo, Manuel Salvadores, Ian Millard, and Nigel Shadbolt

Electronics and Computer Science, University of Southampton,
Southampton, UK.
{gc3,ms8,icm,nrs}@ecs.soton.ac.uk
http://ecs.soton.ac.uk/

Abstract. This paper addresses the issue of representing time entities (i.e. instants and intervals) as Linked Data, and how to exploit topological temporal relationships in order to increase the connectivity degree within Linked Data sets. Describing and efficiently managing temporal information in knowledge management systems is important. Information is volatile, dependant on a number of contexts for its interpretation, among them "time". Many data sets contain information that is valid only within a given time frame (e.g. roles fulfilled by different people at different times), whereas others describe temporal events. In this paper we present an approach to describe temporal entities as reusable URIs that can be adopted by data publishers as a temporal context for their information resources. The approach identifies a set of discrete temporal entities as relevant for a certain domain (e.g. financial years for the public sector) while a RESTful API is provided to users to dynamically create their own temporal entities. Once a dynamic temporal URI is resolved, information is provided to situate such URI in reference to the domain relevant entities. The URI resolution employs simple topological temporal reasoning in order to exploit the qualitative relationships between entities. We also provide a usage scenario of our approach based on a backlinking service and using Public Sector Information published in Linked Data format within the EnAKTing project.

Keywords: Linked Data, time, reasoning

1 Introduction

The Linked Data initiative represents the first collaborative effort to create a *Web of Data* (WoD henceforth) at scale, providing a few, simple guidelines for publishing content using well established standards [5]. Such guidelines and standards are leading the way to a new paradigm of interaction between government and citizens in the UK and around the world. In order to pursue better access for citizens to information held by local as well as national public organisations, the UK government has launched a public initiative for publishing Public Sector Information (PSI), adopting Linked Data as recommended future best practice. Data sets recently delivered to the public include: government expenses, NHS trusts' performances, public transportation, and a whole set of statistics about crime, mortality, census, environment, school and social indicators. Some of the data sets mentioned have been published already in a Linked Data format, others have been translated within the EnAKTing project[1], and many others are waiting to be made available in the Linked Data cloud.

The nature and validity of the information is often related to a time frame and is therefore not universal. For example, the definition of a constituency is temporary (e.g. Southampton Test constituency was established in 1950, the date where the previous Southampton constituency ceased to exist). The representation of temporal entities within the Linked Data cloud is therefore an essential step in order to provide temporal context to other entities.

A common trait of PSI, composed largely of national statistics, is its temporal validity (i.e. the time frame when the data was collected). Having a continuous collection of data helps policy analysts study trends, but it requires a coherent representation of time entities in the information ecosystem.

In this paper, we present a pattern and a tool for representing time lines that support the temporal contextualization of linked data entities. Such a pattern allows us to describe qualitative relations among temporal entities as well as absolute time points and intervals that can then be queried and exploited in order to retrieve relevant resources. Allen's temporal intervals algebra [2] is used to describe qualitative temporal relationships between entities.

2 Background

Time representation, querying and reasoning is a well known topic in computer science where many proposals have been produced; from formal systems, to ontologies and query languages. Within the database community the representation and management of temporal information can be found for example in the temporal extension to SQL named TSQL2 [12]. The semantic web community took inspiration from the philosophical roots of TSQL2, based on Allen's time intervals algebra [2], and encoded its semantics into a Time ontology [10]. The use of a Time ontology in OWL allows users to create temporal information in RDF although its semantics exceeds the capabilities of normal RDF query languages like SPARQL.

Allen provided a model of time based on intervals [2] whose semantics proved to be extended to represents time points too [3]. Allen's original time intervals' relationships depicted in Figure 1 are: **before**, **equal**, **meets**, **overlaps**, **during**, **starts**, and **finishes**. These relationships, along with their inverses, fully express all the possible temporal relationships that can hold between two intervals.

Such a formalization of time has been encoded, as reported above, in an OWL ontology [10] that introduces the concept of *Time Entity* that is composed of two disjointed subclasses, *Time Instant* and *Time Interval*. The Allen relationships are used in this time ontology for providing semantics to the properties

[1] `http://enakting.org`

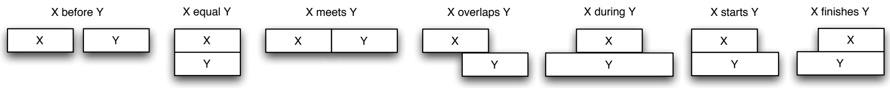

Fig. 1. Time interval relationships

between *Time Entity* instances. *Time Instants*, although not present in the original Allen model, are treated as point form intervals whose starting and ending point coincide.

Temporal validity of RDF statements was originally investigated by Gutierrez et al. [9] where a system of temporal labelling for single RDF statements was devised. A *temporal label* is a natural number $[t]$ that is used for labelling the temporal validity of a triple $\langle s, p, o\rangle$. An RDF triple annotated with a temporal label is then called a *temporal triple* and is represented as $\langle s, p, o\rangle[t]$. Temporal triples can then be extended to intervals $\langle s, p, o\rangle[t_1, t_2] = \{\langle s, p, o\rangle[t]|t_1 \leq t \leq t_2\}$. *Temporal graphs* are defined as a set of temporal triples so that the validity of the included statements can be effectively queried. The concept of *temporal graphs* has been recently implemented by exploiting RDF named graphs. The approach has then been extended introducing additional support for the SPARQL language, called τ-SPARQL [14], designed to handle statements temporal validity. A recent approach to RDF annotation [13] included the capability to reason, among other domains, over temporal information.

Contextualization in the WoD relies on authoritative sources of URIs for naming entities. A typical example of this usage pattern can be seen by looking at *DBpedia* [4], used as a common target reference when aligning data sets. More recently, ontologies like the Ordnance Survey (OS henceforth) Administrative Geography ontology [8], has been used to geographically contextualize Public Sector Information data sets [6]. The process of information contextualization implies the reuse of authoritative URIs and their successive exploitation in order to discover relevant resources. Such processes are facilitated by the adoption of URIs for naming things, but it is hampered by the very architecture of the Web. In fact, the Web does not handle back links that are the target relationships when exploiting hub data sets for information contextualization. Instead they need to be scraped and indexed separately in order to allow the discovery of relevant resources.

Creation of URIs for describing reference time intervals has been proposed firstly by Ian Davis with his `http://placetime.com` site where he uses ISO 8601 standard [1] to format instants and time intervals into dynamically resolvable URIs. This approach has been incorporated in some recent proposals for creating Linked Data for PSI[2].

In this paper we present a pattern for describing temporal information within Linked Data. This approach allows users to refer to reference time intervals sup-

[2] `http://www.epimorphics.com/web/wiki/using-interval-set-uris-statistical-data` accessed at 17/05/2010

plied by data providers (e.g. quarter of financial years) as well as unanticipated temporal entities within a time line. The resolution of such dynamic URIs gives not only useful information about the time entity itself (like that provided by `http://placetime.com`) but also temporal topological relationships in reference to the managed discrete time entities.

Temporal reasoning adopted here is a proper subset of Allen's classical time algebra, namely only topological temporal relationships are taken into account in order to retrieve relevant URIs. The use of backlinks, jointly with a lightweight temporal query support, allows us then to retrieve the contextualised resource.

3 Motivation

The Linked Data principles [5] promote a WoD whose architecture is inherently decentralised, relying on the reuse of available data, by means of linking, in order to give semantics and context to new data. In a Linked Data perspective therefore, the process of data contextualization is tightly connected with the process of data linkage and data cloud connectivity.

The RDF data model inherits XML schema (`xsd` henceforth) support for data types (i.e. `xsd:date` and `xsd:dateTime`). However, the sole use of `xsd` without the adoption of a linkable URI representation for relevant time entities, diminishes the level of connectivity of the overall data cloud, making the information contextualization process more difficult. `xsd` time related data types provide in fact a uniform representation of the data semantics whose interpretation is flattened into a shared time line, regardless of the context of use of the time information. Moreover, thinking of the WoD as a hypermedia system browsable by means of software agents, the node connectivity is as important as the schema that provides nodes with semantics.

For instance, let us consider the concept of a **financial year**. Although it is present in many countries, its actual extension differs from country to country, but even if the extension of the financial year in the UK and in India is the same (both of them start in April), it would be quite unsettling asserting that they are the same thing. Considering the semantics of *financial year* merely as a product of its extensional features (i.e. starting and ending date), this sentence would lead to no semantic clash at all. However, from an information retrieval point of view, binding UK relevant data to the Indian financial year, still seems quite inaccurate and misleading. In order to refer to different temporal entities we need to lift up from a flat representation of time and create distinct name spaces that restore contextual differences and identities. In the previous example, this would lead to the creation of different URIs for the UK and Indian financial year and then to the correct link between the UK relevant information and the UK temporal context. The concomitant creation of contextually equivalent name spaces for temporal entities only defers the creation of authoritative sources of URIs. Co-reference systems [7] in fact allow us to aggregate equivalent URIs enabling a later integration without inhibiting the publication of information.

Another issue that presents itself when dealing with temporal dimensions in publishing and dealing in general with statistical data is the heterogeneous temporal definition used when collecting observations. Data observations are in fact provided (explicitly or not) for different timespans and with different granularity (by solar or financial year, quarter etc.). It is therefore necessary to provide reasoning services for reconciling when possible these differences facilitating a semantically enabled retrieval of information.

As an example, consider two data sets from `http://data.gov.uk` about NHS performance[3] and mortality[4]. These data sets have been translated into Linked Data entities but unfortunately, the first data was described by quarter of financial year while the latter by whole financial years. Either at retrieval and at aggregation phase, the containment information between temporal entities must be explicitly represented and exploited in order to produce meaningful results. As an example, the knowledge that (i) financial years are composed by four quarters, that (ii) each of them is temporally contained by it, and that (iii) statistical observation can be summed up when aggregating to supersets of dimensions can be exploited for producing aggregate statistics that were not originally given.

Furthermore, considering more broadly general Linked Data sets, if we are to represent time validity of entities linking them with temporal entities, we would be able to exploit temporal entity semantics.

For an effective description of time that could help the above mentioned issues, the following high level requirements must be satisfied: explicit URI representation of temporal entities of general interest, and support for some form of temporal reasoning.

In the rest of the paper we will describe a proposal for the description and management of temporal URIs for Linked Data entities.

4 Time Representation in Linked Data

In order to allow users to explicitly express temporal facets of their data reusing reference URIs we developed the concept of *Linked Timelines*, knowledge bases about general instants and intervals that expose resolvable URIs. *Linked Timelines* adopt the OWL time ontology as the standard vocabulary for describing temporal entities. Temporal entities in a timeline contain RDF statements not only defining their starting and final instant, but also temporal relationships between other entities managed in the timeline creating a lattice of temporal entities. Such information will be used in order to infer temporal topological information as will be described in this section.

The OWL Time ontology[5] has been extended with three *datatype* properties: `hasXSDStart`, `hasXSDEnd`, and `hasXSDDuration` for describing with XML Schema datatype `xsd:dateTime` and `xsd:duration` the starting and ending instant of a time interval and its duration respectively. These extensions act as a short-cut for defining temporal intervals without creating first the instances for the starting and ending instants. Intervals and not Instants themselves are in fact

[3] `http://nhs.psi.enakting.org`
[4] `http://mortality.psi.enakting.org`
[5] `http://www.w3.org/2006/time#`

the main objects represented within timelines. Therefore the start-end instants would either be anonymous nodes, which are to be avoided whenever possible when publishing linked data, or verbose second class instances that would clutter the knowledge base. Such short-cut has been provided in the OWL Time ontology for the Time Instants with the creation of the `inXSDDateTime` datatype property, but not for time intervals. We introduced the possibility of defining intervals reusing entirely XML schema datatypes in order to limit the creation of entities to the ones of actual interest. An example of a temporal entity described within a *Linked Timeline* can be seen in Figure 2 where the RDF description for the solar year 2007 is reported in Turtle syntax. Note that in Figure 2 there is no explicit usage of the property `hasXSDEnd`, but its value can be inferred by exploiting the semantics of `time:intervalMeets` property and knowing the starting instant of `<http://time.psi.enakting.org/id/2008>` (i.e. the solar year 2008).

```
<http://time.psi.enakting.org/id/2007> a timepsi:Year ;
    rdfs:label "Solar year 2007" ;
    timepsi:hasXSDStart "2007-01-01T00:00:00Z"^^xsd:dateTime ;
    timepsi:hasXSDDuration "P1Y"^^xsd:duration ;
    time:intervalContains  <http://time.psi.enakting.org/id/2007/08/Q1> ,
     <http://time.psi.enakting.org/id/2007/08/Q2> ,
     <http://time.psi.enakting.org/id/2007/08/Q3> ;
    time:intervalStartedBy  <http://time.psi.enakting.org/id/2006/07/Q4> ;
    time:intervalMeets <http://time.psi.enakting.org/id/2008> ;
    time:intervalMetBy <http://time.psi.enakting.org/id/2006> .
```

Fig. 2. Turtle representation of the solar year 2007

The `http://time.psi.enakting.org` timeline contains relevant URIs for the UK financial year (other domains could describe different kinds of temporal entities), therefore alongside solar years we find financial years and quarters of financial year. Other timelines can be set up in order to represent time entities relevant to other domains and applications. The RDF representation of the solar year 2007 in Figure 2 for example, additionally to the starting date and duration, contains the relations between the year 2007 and other *relevant* intervals such as the the solar year 2006 (that meets 2007) and 2008 (met by the 2007). The relevance criteria is determined by the data publisher and, in this example, the relevant entities are the quarters of financial years contained in the year 2007 (i.e. the fourth quarter of the financial year 2006/07 and the first three quarters of the financial year 2007/08).

In order to limit the size of the knowledge base we do not describe **before** and **after** relationships between every instance, which will cause a geometric explosion in the number of statements. Instead we limited the kind of relationships explicitly described to the topological ones (i.e. `time:intervalContains` and `time:intervalDuring`) and the ones useful to recreate the lattice of entities (i.e. `time:intervalMeets`, `time:intervalMetBy`, `time:intervalStarts`, and `time:intervalStartedBy`). It is noteworthy that in Figure 2 we did not represent the overlapping relationships between the solar year 2007 and the financial year 2007. In the UK in fact, the financial year starts the first of April of the solar year, therefore every solar year overlaps its financial year.

The rational for limiting the kind of relationships to the topological ones is twofold: primarily if all the temporal relationships were to be reported, as stated before, there could be a potential problem of space complexity. Secondarily, a topological relationships' semantics is easier to model when dealing with data integration and aggregation. In fact, as argued by [2], *during* relationships are the best candidates to define a hierarchy of intervals where properties can be inherited. For example, if a condition P holds during a time interval T and we know that a time interval t happens *during* T, then we can conclude that P holds also during t.

4.1 Dynamic temporal URI encoding

The space of discrete URIs described so far is complemented by a RESTful interface for defining dynamically time instants and intervals that happen in the same timeline. When resolving such *dynamic* URIs a document is returned that provides a description of the instant/interval. The format of the document is either decided via a content negotiation mechanism or by directly stating the desired format. Encoding time entities via URIs allows users to refer explicitly to particular timelines (e.g. for UK financial periods) even if the actual URIs are not known or not yet existent, enabling therefore a future evolution of the timeline.

The API interface for creating dynamic temporal URIs follows the `http://placetime.com` encoding[6], and it is defined as follows:

```
http://time.psi.enakting.org/{type}/{YYYY}-{MM}-{DD}T{hh}:{mm}:{ss}{TZ}
P[{y}Y][{m}M][{d}D][T[{h}H][{n}M][{s}S]][/{format}]
```

Where the first mandatory part is defined as follows:

- `type` is either `interval` or `instant`
- `YYYY` Four digits (from 0001 to 9999) represent the year (only AD dates are valid).
- `MM` Two digits (from 01 to 12) represent the month of the year.
- `DD` Two digits (from 01 to 28, 29, 30 or 31 depending on the month) represent the day of the month.
- `hh` Two digits (from 00 to 23) representing the hour of the day.
- `mm` Two digits (from 00 to 59) representing the minute of the hour.
- `ss` Two digits (from 00 to 59) representing the second of the minute.
- `TZ` A string representing the timezone. This can be either a `Z` for UTC timezone or a string in the format: `[+/-]hh:mm` representing an offset of `hh` hours and `mm` minutes from the UTC timezone.

[6] `http://placetime.com`

If `type` is `interval`, then a duration must be given and least one of the following parameters must be provided:

- `y` A number of years in the interval
- `m` A number of months in the interval
- `d` A number of days in the interval
- `h` A number of hours in the interval
- `n` A number of minutes in the interval
- `s` A number of seconds in the interval

The last part is optional and defines the format of the returned document:

- `format`, if present, can be one of: `doc`, `rdf`, or `ttl` and defines the format of the document returned (HTML, RDF/XML, or Turtle respectively). The content negotiation will return the same document if this parameter would be missing and the `Accept:` header in the HTTP request would be set to the wanted MIME type (e.g. `application/x-turtle` for the `ttl` format).

As an example of usage of this API, if we consider the following instant URI: `http://time.psi.enakting.org/instant/2007-01-10T10:00:00Z/ttl`, it will return the following Turtle document:

```
@PREFIX timepsi: <http://time.psi.enakting.org/def/> .
<http://time.psi.enakting.org/instant/2007-01-10T10:00:00Z>
 a time:Instant ;
 rdfs:comment "The instant 10:00:00Z, of the tenth day of the
              month of January in the year 2007 of the
              Gregorian calendar." ;
 timepsi:hasXSDStart "2007-01-10 10:00:00+00:00"^^xsd:dateTime ;
 rdfs:label "2007-01-10T10:00:00+00:00" ;
 time:inside <http://time.psi.enakting.org/id/2007> ,
    <http://time.psi.enakting.org/id/2006/07> ,
    <http://time.psi.enakting.org/id/2006/07/Q4>.
```

As per the previous example, for defining intervals, if we consider the following URI: `http://time.psi.enakting.org/interval/2007-01-01T12:00:00ZP4M/ttl`, it will return the Turtle document reported as follow:

```
@PREFIX timepsi: <http://time.psi.enakting.org/def/> .
<http://time.psi.enakting.org/interval/2007-01-01T00:00:00ZP4M>
 a time:ProperInterval ;
 rdfs:comment "A time-interval of exactly   4 month(s),
              beginning at 0:00:00Z, on the first day
              of the month of January in year 2007 of
              the Gregorian calendar." ;
 rdfs:label "2007-01-01T12:00:00+00:00P4M" ;
 timepsi:hasXSDStart "2007-01-01 00:00:00+00:00"^^xsd:dateTime ;
 timepsi:hasXSDDuration "P4M"^^xsd:duration ;
 time:intervalDuring <http://time.psi.enakting.org/id/2007> ;
 time:intarvalContains <http://time.psi.enakting.org/id/2006/07/Q4> .
```

As illustrated in the two code examples above, the RDF documents describe the temporal entity resolved with a label and a description in natural language. Moreover, the documents describe the topological temporal relationships that hold between the resolved entity and the entities managed in the timeline. The rationale for returning only the topological relationships has been discussed in a earlier part of this section.

By default then, when resolving a dynamic URI in a *Linked Timeline*, the following entities belonging to it are returned, depending on the kind of temporal entity (i.e. instant or interval):

instant: all the temporal intervals that contain the instant encoded by the resolved URI (`time:inside` property is here used to state that an instant is contained within a temporal interval)

interval: all the temporal intervals that contain or are contained by the interval encoded by the resolved URI (`time:intervalContains` and `time:intervalDuring` are used here respectively to state that an interval contains or is contained by another interval)

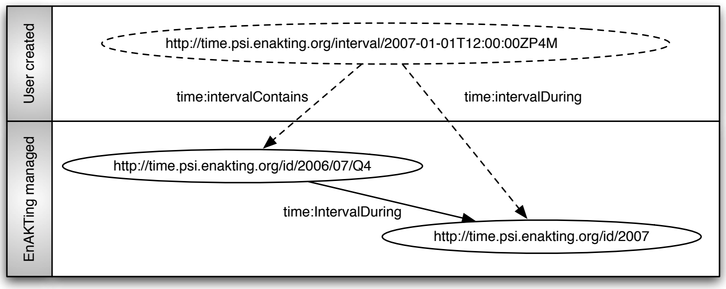

Fig. 3. URI dynamic resolution

The dynamic resolution of temporal entities, together with the temporal reasoning that puts it in context with the managed ones, allows the reuse of time URIs in foreign data sets. Users can in fact create their URIs programmatically (*User created* in Figure 3) making reference to a *Linked Timeline* that will automatically contextualize them with reference to a managed set of temporal entities (*EnAKTing managed* in Figure 3).

4.2 EnAKTing Linked Timeline implementation

Within the EnAKTing project, we translated public sector information in Linked Data format, aligning the different dimensions whenever possible. Geographical dimensions were aligned to the Ordnance Survey administrative ontology but for temporal dimensions we had to create an ad hoc solution. The *Linked Timeline* instance created for PSI data is accessible at `http://time.psi.enakting.org` and it is composed by two sections. The first section (upper part of Figure 4) is a timeline widget[7] that illustrates the discrete entities managed by the timeline.

[7] For this we have used the **Timeline** widget from SIMILE project, `http://www.simile-widgets.org/timeline/`.

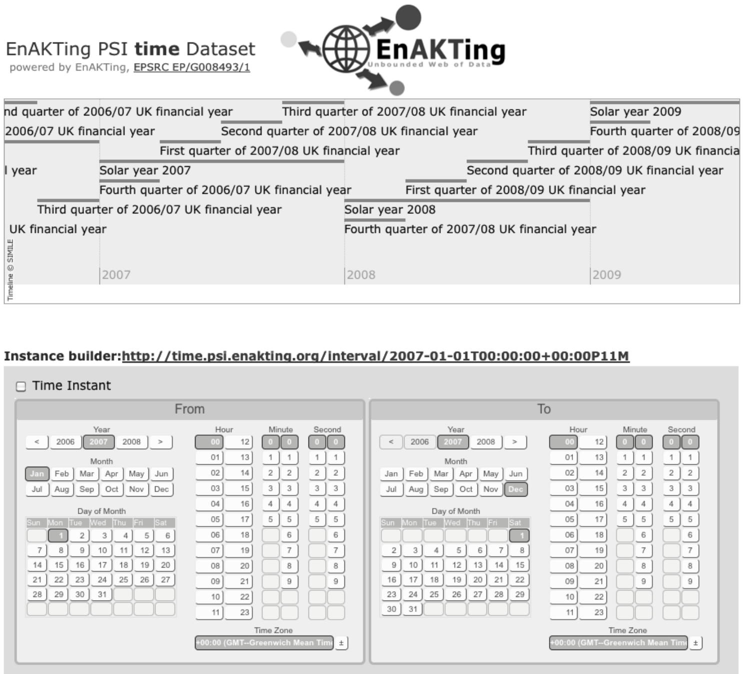

Fig. 4. EnAKTing Linked Timeline

The second part of the interface (bottom part of Figure 4) is composed by two calendars that allow to select the starting and ending day of a temporal interval. Alternatively, the user can check the **Time Instant** check box and the second calendar will be disabled. While the user chooses the temporal extent of its temporal entity the URI just above the two calendars will change accordingly, creating a URI that can be resolved immediately or copied and paste for later use.

4.3 Backlinks Integration

The PSI Backlinking service provides an access point to retrieve backlinks from Foreign URIs [11]. Foreign URIs make data discovery difficult because it is not possible to navigate the RDF documents of the WoD bidirectionally. `http://backlinks.psi.enakting.org` provides an API to retrieve collections of backlinks for a given URI.

In order to improve the connectivity of our data sets we have aligned the temporal entities in a number of our data sets (the process is still ongoing) to our linked Timeline. Then we have collected the back links from the temporal entities and extended our backlinking service in order to exploit the temporal reasoning. The data sets aligned for this evaluation are three: `http:`

//parliament.psi.enakting.org) is an historical record of the UK parliamentary voting, MPs and their constituencies; http://nhs.psi.enakting.org is a temporal series of NHS performance statistics collected over the years (from 2005 to 2008); and http://education.data.gov.uk is a directory of present and past educational institutions reporting ancillary information (e.g. address, number of pupils, opening and closing date).

The Backlinking and the Linked Timeline services run as separate services and the first service performs HTTP requests to get the temporal entities from the latter that employ temporal containments (see Figure 3) in order to retrieve all the relevant URIs. Entities returned by the temporal reasoner to the Backlinking service are created by a previous scraping of the data sets and are not usually returned to normal users that will be otherwise cluttered with non contextualized temporal entities. When the backlinking service recognizes a URI from the Linked Timeline namespace it gets the list of contained entities for the input URI and returns the backlinks connected to any URI contained in the temporal interval. For example, consider the URI used in Section 4 for defining temporal intervals: http://time.psi.enakting.org/interval/2007-01-01T00:00:00ZP4M. Calling the backlink service on this URI will return (as shown in Figure 5) the following backlinks: 3 ministerial offices, 299194 NHS performance statistic items, and 246 divisions.

Fig. 5. Backlinks Service result

5 Conclusions

In this paper we have presented a pattern for temporal URIs exposure and a tool based on Allen's temporal algebra for creating such URIs within Linked Timelines that supports the contextualization task. The usage of HTTP resolvable URIs and the decoupling of the information publishing and retrieval phase make this approach suitable for large scale data publishing infrastructures. Use of external links for contextualization of local information instead of using local statements means that local data querying and retrieval can be eventually delegated to external services.

There are many benefits to be realised by including or retrospectively adding Linked Timelines to data sets, not least of which are enhanced discovery and querying potential providing linkage and exposure for information that may not previously have been accessible. Moreover, the adoption of Linked Timelines does not preclude data publishers from providing local solutions for temporal indexing and information retrieval.

Acknowledgements

This work was supported by the EnAKTing project funded by the Engineering and Physical Sciences Research Council under contract EP/G008493/1.

References

1. Data elements and interchange formats – information interchange – representation of dates and times. ISO 8601:1988(E), International Organization for Standardization (June 1988)
2. Allen, J.F.: Maintaining knowledge about temporal intervals. Commun. ACM 26(11), 832–843 (1983)
3. Allen, J.F., Hayes, P.J.: Moments and points in an interval-based temporal logic. Comput. Intell. 5(4), 225–238 (1990)
4. Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., Ives, Z.: DBpedia: A nucleus for a web of open data. In: Proceedings of the 6^{th} International Semantic Web Conference (ISWC). Lecture Notes in Computer Science, vol. 4825, pp. 722–735. Springer (2008)
5. Berners-Lee, T.: Design issues: Linked data. `http://www.w3.org/DesignIssues/LinkedData.html` (2006)
6. Correndo, G., Salvadores, M., Yang, Y., Gibbins, N., Shadbolt, N.: Geographical service: a compass for the web of data. In: Linked Data on the Web (LDOW2010) (April 2010)
7. Glaser, H., Jaffri, A., Millard, I.: Managing co-reference on the semantic web. In: WWW2009 Workshop: Linked Data on the Web (LDOW2009) (April 2009)
8. Goodwin, J., Dolbear, C., Hart, G.: Geographical linked data: The administrative geography of great britain on the semantic web. Transaction in GIS 12(1), 19–30 (February 2009)
9. Gutierrez, C., Hurtado, C.A., Vaisman, A.: Introducing time into RDF. IEEE Transactions on Knowledge and Data Engineering 19, 207–218 (2007)
10. Hobbs, J.R., Pan, F.: An ontology of time for the semantic web. ACM Transactions on Asian Language Information Processing (TALIP) 3(1), 66–85 (2004)
11. Salvadores, M., Correndo, G., Szomszor, M., Yang, Y., Gibbins, N., Millard, I., Glaser, H., Shadbolt, N.: Domain-specific backlinking services in the web of data. In: Web Intelligence (September 2010)
12. Snodgrass, R.T. (ed.): The TSQL2 Temporal Query Language. Kluwer (1995)
13. Straccia, U., 0002, N.L., Lukacsy, G., Polleres, A.: A general framework for representing and reasoning with annotated semantic web data. In: Fox, M., Poole, D. (eds.) AAAI. AAAI Press (2010)
14. Tappolet, J., Bernstein, A.: Applied temporal RDF: Efficient temporal querying of RDF data with SPARQL. In: 6^{th} Annual European Semantic Web Conference (ESWC2009). pp. 308–322 (June 2009)