Using XForms to Create, Publish, and Manage Linked Open Data

Ethan Gruber

American Numismatic Society

<gruber@numismatics.org>

Abstract

This paper details the numismatic thesaurus, Nomisma.org, and its associated front-end and back-end features. The application's architecture is grounded in XML technologies and SPARQL, with XForms underpinning the creation, editing, and publication of RDF. Its public user interface is driven by the XML Pipeline Language in Orbeon, enabling transformation of RDF/XML and SPARQL XML responses into a wide array of alternative serializations, driving geographic visualizations and quantitative analyses in other digital numismatic projects.

Keywords: Semantic Web, XForms, Numismatics

1. Introduction

Nomisma.org is a collaborative project to define the intellectual concepts of numismatics following the principles of Linked Open Data: URIs for each concept with machine-readable serializations (RDF/XML, Turtle, JSON-LD, etc.) conforming to a variety of relevant ontologies (like SKOS: the Simple Knowledge Organization System) [SKOS]. What began as a prototype created in 2010 by Sebastian Heath and Andrew Meadows (then of New York University and the American Numismatic Society, respectively) to demonstrate the potential of applying semantic web technologies to numismatic research has evolved into the standard thesaurus for the discipline, driven by a scientific committee of scholars and information technologists, and adopted by a growing number of cultural heritage institutions. These Nomisma-defined concepts, and the software architecture built upon them, are the backbone for projects such as Coinage of the Roman Republic Online (CRRO) and Online Coins of the Roman Empire (OCRE), which seek to define all typologies of the Roman Republic and Empire, facilitating the aggregation of coins from museum and archaeological databases that are related to these typologies.

2. Numismatic Concepts as Linked Open Data: A Brief Introduction

Numismatics as a discipline emerged during the Medieval period and gradually became more scientific over the centuries. By the late 18th century, the classification methodology had evolved into system still used today [GRUBER]. Coins have historically been categorized by a variety of individual attributes: the manufacture process, material, monetary denomination, production place (or mint), date, entities responsible for issuing the coin (whether individual rulers or corporate organizations), and the iconography and inscriptions (or "legend" in numismatic terminology) on the front and back (obverse and reverse) of the coin. The combination of each of these individual attributes comprised a coin "type," and types were often uniquely numbered, thematically organized, and published in volumes of printed books. For example, Roman Republican coins have been published in numerous volumes over the last century, but the standard reference work for the period remains Michael Crawford's 1974 publication, Roman Republican Coinage (RRC). Collections of Republican coins therefore refer to standard type numbers from RRC, e. g., 244/1, a silver denarius minted in Rome in 134 B.C. These numbers were once printed in collection inventories or cards associated with each coin, but are now inserted into bibliographic fields in museum databases.

These databases, however, are authored in the native language of the collection. The Roman emperor, Augustus, is the same entity as Auguste in French or アウグストゥス in Japanese. In order to perform large-scale analyses of related coins across many different databases, each with its own terminology, the discipline needed to rethink authority control, and so the Linked

Open Data approach to taxonomies was adopted. Augustus could be represented by a URI, http://nomisma.org/id/augustus (with a CURIE of nm:augustus), defined as a foaf:Person in the Friend of a Friend ontology [FOAF]. This URI serves both as a unique, language-agnostic identifier for the entity, but also a web page where both human- and machine-readable information can be extracted. Below we discuss the RDF data models that comprise the Nomisma information system, organized into three broad divisions: concepts in the Nomisma.org thesaurus, coin types published in OCRE, CRRO, and other projects, and the model that defines physical coins.

2.1. The Data Models

2.1.1. Thesaurus

We implement a variety of data models for different types of data objects, mixing and matching classes and properties from numerous ontologies. The SKOS ontology was implemented for modeling the intellectual concepts of numismatics, which include not only the rulers responsible for issuing coinage, but each of the aforementioned categories: manufacture method, material, mint, denomination, field of numismatics (broad cultural areas, like Greek or Roman), and many others. Many of these categories are specific to the discipline, and are therefore defined by classes in a numismatic ontology (http://nomisma.org/ontology#, prefix: nmo) developed and published by the Nomisma.org scientific committee.

Other more generalizable types of data objects are bound to classes from other common ontologies. People and organizations carry foaf:Person and foaf:Organization classes, respectively, historical periods are defined by CIDOC-CRM, a conceptual reference model from the cultural heritage domain [CRM], the W3C Org ontology [ORG] has been implemented for defining the role a person plays with respect to the issuing of coinage (e.g., Roman emperor, king, or magistrates of various titles).

With respect to RDF properties, preferred labels and definitions may be inserted in as many languages as necessary to facilitate multilingual interfaces, and concepts may be linked hierarchically via skos:broader. And, importantly, Nomisma uses SKOS properties like exactMatch to link to identical concepts in other linked data systems (such as Geonames, OCLC's Virtual International Authority File (VIAF), Wikidata, and the Getty vocabularies), which enable the integration of coins into a wider array of cultural heritage projects, such as Pelagios Commons. While typical SKOS properties are implemented within instances within Nomisma, properties from other ontologies are implemented conditionally upon the class of object. Mints and regions may bear coordinates from the W3C basic geo (WGS84) vocabulary [WGS84] or from the UK's Ordnance Survey ontology in the form as geoJSON [OSGEO].

```
@prefix skos: <http://www.w3.org/2004/02/skos/core#> .
@prefix geo: <http://www.w3.org/2003/01/geo/wgs84_pos#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix dcterms: <http://purl.org/dc/terms/> .
@prefix nm: <http://nomisma.org/id/> .
@prefix nmo: <http://nomisma.org/ontology#> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .

nm:byblus a nmo:Mint;
        rdf:type <http://www.w3.org/2004/02/skos/core#Concept>;
        dcterms:isPartOf <http://nomisma.org/id/greek_numismatics>;
        skos:prefLabel "Byblus"@en,
                "Byblos"@fr;
        skos:definition "The mint at the ancient site of Byblus in Phoenicia."@en;
        skos:closeMatch <http://dbpedia.org/resource/Byblos>,
                <http://pleiades.stoa.org/places/668216>,
                <http://www.geonames.org/273203>,
                <http://collection.britishmuseum.org/id/place/x30547>,
                <http://vocab.getty.edu/tgn/7016516>,
                <http://www.wikidata.org/entity/Q173532>;
        geo:location <http://nomisma.org/id/byblus#this>;
```

```
        skos:broader <http://nomisma.org/id/phoenicia>;
        skos:altLabel "Byblos"@en.

nm:byblus#this a geo:SpatialThing;
        geo:lat "34.119501"^^xsd:decimal;
        geo:long "35.646846"^^xsd:decimal;
        dcterms:isPartOf <http://nomisma.org/id/phoenicia#this>
```

The Org ontology has been applied to connect people with dynasties and corporate entities, including their roles within these organizations and dates these offices have been held. Dublin Core Terms such as dcterms:isPartOf and dcterms:source have been applied for hierarchical linking and bibliographic references, respectively. The thesaurus models are stable, but do evolve to meet increased demands by our users. At the moment, the system is deficient in tracking data provenance, and we do plan to implement PROV-O soon [PROVO].

2.1.2. Coin Types

Just as individual concepts have been defined by URIs, so too are more complex coin types. RRC 244/1 is represented by http://numismatics.org/crro/id/rrc-244.1, an instance of an nmo:TypeSeriesItem in the Nomisma ontology, which contains a variety of properties connecting the type to individual categorical attributes (concepts in the Nomisma thesaurus) and literals for the obverse and reverse legends, symbols, and iconographic descriptions. Below is the RDF/Turtle representing RRC 244/1:

```
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix nmo: <http://nomisma.org/ontology#> .
@prefix dcterms: <http://purl.org/dc/terms/> .
@prefix nm: <http://nomisma.org/id/> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix skos: <http://www.w3.org/2004/02/skos/core#> .

<http://numismatics.org/crro/id/rrc-244.1>
        rdf:type <http://www.w3.org/2004/02/skos/core#Concept>;
        skos:prefLabel "RRC 244/1"@en;
        skos:definition "RRC 244/1"@en;
        dcterms:source nm:rrc;
        nmo:representsObjectType nm:coin;
        nmo:hasManufacture nm:struck>;
        nmo:hasDenomination nm:denarius;
        nmo:hasMaterial nm:ar;
        nmo:hasIssuer nm:c_abvri_gem_rrc;
        nmo:hasMint nm:rome;
        nmo:hasStartDate "-0134"^^xsd:gYear;
        nmo:hasEndDate "-0134"^^xsd:gYear;
        nmo:hasObverse <http://numismatics.org/crro/id/rrc-244.1#obverse>;
        nmo:hasReverse <http://numismatics.org/crro/id/rrc-244.1#reverse>.

<http://numismatics.org/crro/id/rrc-244.1#obverse>
        nmo:hasLegend "GEM X (crossed)";
        dcterms:description "Helmeted head of Roma, right. Border of dots."@en;
        nmo:hasPortrait <http://collection.britishmuseum.org/id/person-institution/60208>.

<http://numismatics.org/crro/id/rrc-244.1#reverse>
        nmo:hasLegend "C·ABVRI";
        dcterms:description "Mars in quadriga, right, holding spear, shield and reins in left hand
          and trophy in right hand . Border of dots."@en;
        nmo:hasPortrait <http://collection.britishmuseum.org/id/person-institution/59284>.
```

2.1.3. Physical Coins

The RDF model for physical coins implements properties from several ontologies, but the Nomisma ontology is the most prominent. The complexity of the model is variable depending on the condition of the coin ands the certainty by which it may be linked to a URI for a coin type. If the type RRC 244/1, http://numismatics.org/crro/id/rrc-244.1, is an nmo:TypeSeriesItem, then physical specimens that represent this typology are linked with the nmo:hasTypeSeriesItem property. By semantic reasoning, a physical coin of the type RRC 244/1 is a denarius minted in Rome and issued by the magistrate, C. Aburius Geminus, even if these URIs are not explicit within the RDF of http://numismatics.org/collection/1978.64.316. On the other hand, many archaeologically excavated coins are worn beyond certain attribution. In these cases, what information that may be ascertained are made explicit in the triples for a coin. The portrait may be identifiable as the emperor, Vespasian, but an illegible legend prevents linking to a specific coin type. This coin is still available for query alongside other positively identified coins of Vespasian within the Nomisma.org SPARQL endpoint.

In addition to recording the type URI or typological characteristics of a coin, metrological data, like weight and diameter, may also be included. These measurements are vital for large-scale (both in terms of time and space) economic analyses. If the coin has been photographed, foaf:thumbnail and foaf:depiction may be used to link to images. Finally, geodata (whether coordinates and/or a gazetteer URI) related to the find spot may be included, if known. In the event that the coin was found within a hoard (`dcterms:isPartOf ?hoard`), the find coordinates may be extracted from the triples for the hoard via the graph [NMDOCS].

2.2. Gradual Acceptance and Implementation by the Community

A growing body of institutions are adopting Nomisma URIs as standard identifiers for numismatics. The four of the five largest collections of ancient coins in world (American Numismatic Society, Bibliothèque nationale de France, British Museum, and Berlin Münzkabinett) are in various stages of implementation, as are a handful of European coin find databases (such as the UK's Portable Antiquities Scheme and the German Antike Fundmünzen Europa), and smaller museums such as the University of Virginia and University College Dublin. For a full list of contributors, see the Nomisma.org datasets list.

The landscape of digital numismatics has progressed significantly in the six years since the launch of Nomisma, and we are hopeful that, over the next decades, this information system will include millions of coins, enabling scholars to perform large-scale economic analyses of 800 years of Roman coinage distributed from Scotland to India. We hope (or expect) these technical methodologies will open pathways for similar research questions in other fields of numismatics. The foundation for these sophisticated query interfaces is a combination of XML and Semantic Web technologies, bound together in a collection of XML Pipelines, XSLT stylesheets, and Javascript/CSS for the public user interface, all of which are open source (on Github) and built on other open source server applications and open web standards.

3. Architecture

Nomisma.org's architecture is based on an adaptation of XRX (XForms, REST, XQuery), with SPARQL substituted for XQuery. It utilizes a variety of well-supported and documented open source Java-based server applications. Apache Solr facilitates faceted browsing, search-oriented web services, and a geospatial extension to the SPARQL endpoint. Apache Fuseki (part of the Jena project) is the triplestore and SPARQL endpoint. Orbeon XForm s is the core middleware that connects Solr, Fuseki, and external REST services—in both the front-end with XML pipelines and the back-end XForms engine. The SPARQL endpoint was deployed in production in early 2013, enabling us to aggregate Roman Imperial coins from the American Numismatic Society and Berlin Münzkabinett in an early version of OCRE (which contributed to a successful three-year, $300,000 bid with the National Endowment for the Humanities the following year). Orbeon was implemented before this–in early 2012—to deal with challenges surrounding the creation and maintenance of Nomisma IDs.

The architecture is modular. The greatest advantage of XForms, a W3C standard, is that the author may focus solely on the MVC functionality of the web form; the client-side Javascript and CSS and server-side code are inherent to the XForms processor, and so migration from one platform to another should, at least in theory, require little effort. Apache Fuseki was chosen for its ease of deployment, but it may be swapped with any SPARQL 1.1-compliant endpoint. Replacing Solr with another search index would be more difficult, but any application that supports REST interactions via the XForms engine is suitable.

3.1. XForms for CRUD operations of RDF

When Nomisma.org launched in 2010, it was published entirely in Docuwiki, an open source wiki framework. The URI for the mint of Rome, http://nomisma.org/id/rome, was created by hand-editing XHTML embedded with RDFa properties in a single textarea element in an HTML web form. This presented some problems: there was no check for well-formedness of XHTML fragments, so malformed XML could break a web page, and the particular XHTML document model might be inconsistent and not translate into the the appropriate RDF model using the W3C's RDFa 1.1 distiller.

Ensuring data consistency was a primary concern, and so we developed an XForms application to handle the editing of these XHTML+RDFa fragments. By curbing direct human editing of XTHML, we eliminated malformed XML problems outright. XForms bindings enabled stricter validation based on XPath, for example:

- To require one English SKOS Preferred Label and Definition
- Restrict labels and definitions in other languages to a maximum of one
- Latitudes and longitudes must be decimal numbers between -180 and 180
- A variety of conditionals that restrict certain RDF properties to particular classes of data object; e. g., that geographic coordinates apply only to skos:Concepts that represent regions or mints

Since linking Nomisma URIs to concepts in other systems is a vital feature of Five-Star Linked Open Data, we implemented a variety of simple lookup mechanisms that pass search keywords from the web form to XForms submissions in order to query external APIs (for example, of the Getty Vocabularies and British Museum SPARQL endpoints, Wikidata's XML response, or VIAF's RSS feed). These widgets simplified the process by which skos:exactMatch or skos:closeMatch URIs might be integrated into Nomisma data, reducing errors in manually transcribing URIs into the web form. Wikidata's REST API is especially useful, as we are able to extract article titles in many languages to rapidly assemble a list of skos:prefLabels (Figure 1, “Creating nm:dirham, after importing from Wikidata”). These lookup mechanisms are authored in the form of XBL components in Orbeon, and can be (and have been) easily reimplemented in other XForms applications, such as xEAC and EADitor , which are archival publishing frameworks.

In addition to external lookups, the editing interface includes internal lookup mechanisms that query

Figure 1. Creating nm:dirham, after importing from Wikidata

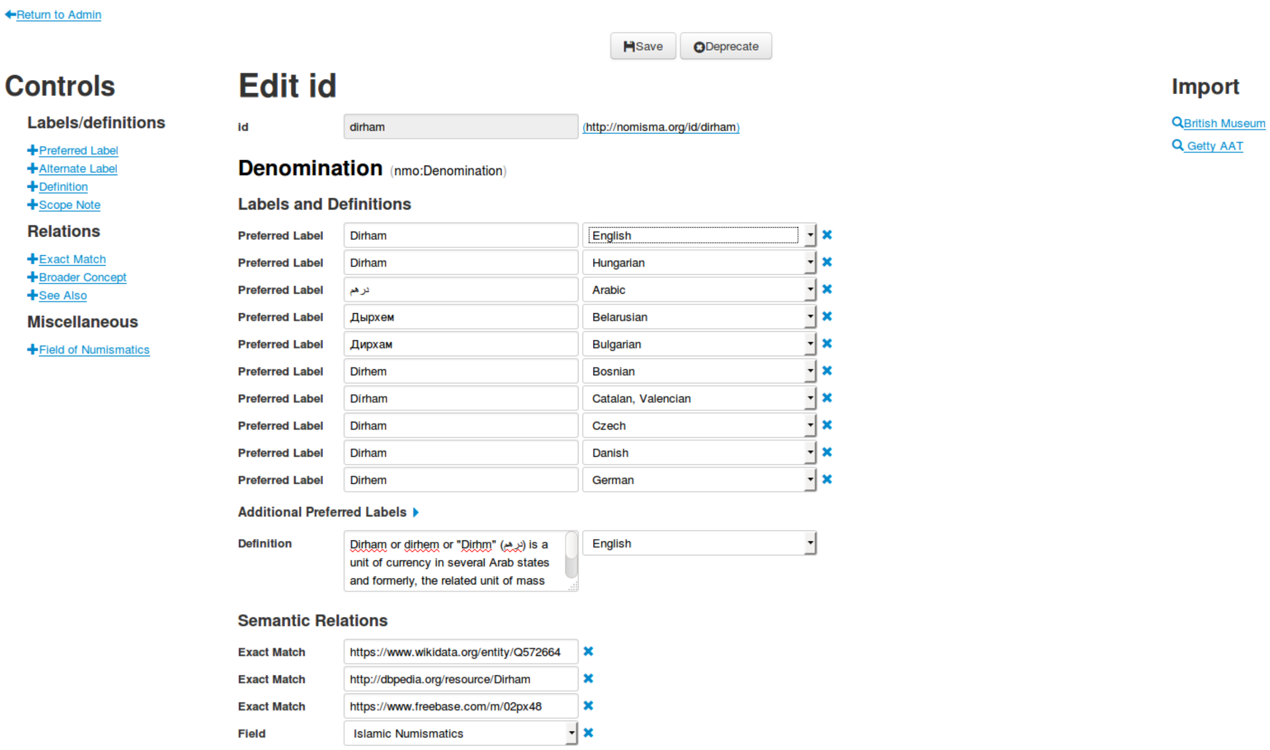

Nomisma's own SPARQL endpoint or Atom feed to link to broader concepts, to associate a concept with a field of numismatics (e.g., to say that a denarius is part of Roman numismatics; `nm:denarius dcterms:isPartOf nm:roman_numismatics`), or to associate a person with a dynasty or organization.

Upon clicking the 'Save' button, the data model was serialized into an XML file and written to the filesystem (as opposed to an XML or NoSQL database), in order that commits could be made nightly to a Github repository for our data [NMDATA]. The Github backups remain an integral part of the data publication workflow today.

The framework eventually grew to incorporate Solr and a SPARQL endpoint. The 'Save' button then hooked into several additional XForms actions and submissions. The source model (XForms instance) is transformed by XSLT into an XML document conforming to Solr's ingestion model, and posted into Solr's REST interface. Next, the XForms engine constructs a SPARQL query to purge triples associated with the URI from the endpoint, which is posted to the endpoint via the SPARQL/Update protocol. Finally, the updated RDF is posted into the endpoint.

In 2014, we begin the transition of migrating from Heath and Meadows' original XHTML+RDFa document model into RDF/XML that conforms proper RDF ontologies, including the introduction of the formal Nomisma ontology developed by Karsten Tolle, a computer scientist at the University of Frankfurt. The model was separated from the view, making it easier to extend the functionality of the public user interface without interfering with the RDFa distillation or exports into other serializations of linked data, like Turtle or JSON-LD. The APIs were rewritten and made dramatically more efficient, having eliminated the need to preprocess XHTML+RDFa into RDF/XML before delivering content to OCRE or other projects.

The XForms editing interface was adapted to RDF/XML, which led to improved consistency of the data, since our export APIs (XSLT transformations of XHTML into other formats) did not always account for every imaginable permutation of RDFa properties jammed into a document model. Today, Nomisma.org's APIs serve out large quantities of RDF data to other American Numismatic Society or partner projects for web maps, multilingual interfaces, etc., as efficiently as possible.

3.2. Batch Concept Creation

Nomisma's XForms back-end functions well for managing individual skos:Concepts. It is an intuitive system used primarily by academics to create or update identifiers required for existing digital numismatics projects. When creating new projects, however, we had a need to create potentially hundreds of new identifiers for mints, rulers, denominations, etc. Academics are accustomed to working in spreadsheets, and the author was often tasked with writing one-off PHP scripts to read spreadsheets as CSV and process data into RDF.

In summer 2015, this functionality was ported into a new XForms application [NMBATCH]. This application requires a Google spreadsheet that conforms to some basic requirements, which must be published to the web so that it can be made available as an Atom feed through the Google Sheets API. All IDs in the spreadsheet must be the same class of data object (for example, a mint or denomination). Each column heading will be read from Atom in the XForms engine, and the user may choose to map a column to a permitted list of RDF properties (Figure 2, “Mapping spreadsheet columns to RDF properties”). There are some XPath controls on the properties that are available for selection—there must be an English preferred label and definition, latitude and longitude are only available for mints, and other conditions described in the Github wiki. If the mapping itself is valid, the user may click a button to validate each row of the spreadsheet. The XForms engine ensures that latitudes and longitudes are valid decimal numbers, that each row has a preferred label and definition, that a URI under a skos:broader mapping does conform to the appropriate RDF Class. After this phase of validation, the application will display a list of validation errors, or if there are none, present the user with a button to publish the data into Nomisma. The publication workflow transforms each Atom entry element into RDF/XML, writes and updates the file on the disk, publishes the data to Solr, and creates or updates the triples in the SPARQL endpoint.

The publication process includes an additional feature where any SKOS matching property with a Wikipedia URL or DBpedia URI is parsed, and a series of API calls are executed to extract preferred labels and alternative authority URIs from Wikidata. This feature has enabled us to enhance concepts that have already been published; we can execute a SPARQL query of all Nomisma IDs with a DBpedia URI to download CSV, upload the CSV to Google Sheets, and then re-run the spreadsheet through the XForms import mechanism to pull labels and other URIs from Wikidata into Nomisma. This is

Figure 2. Mapping spreadsheet columns to RDF properties

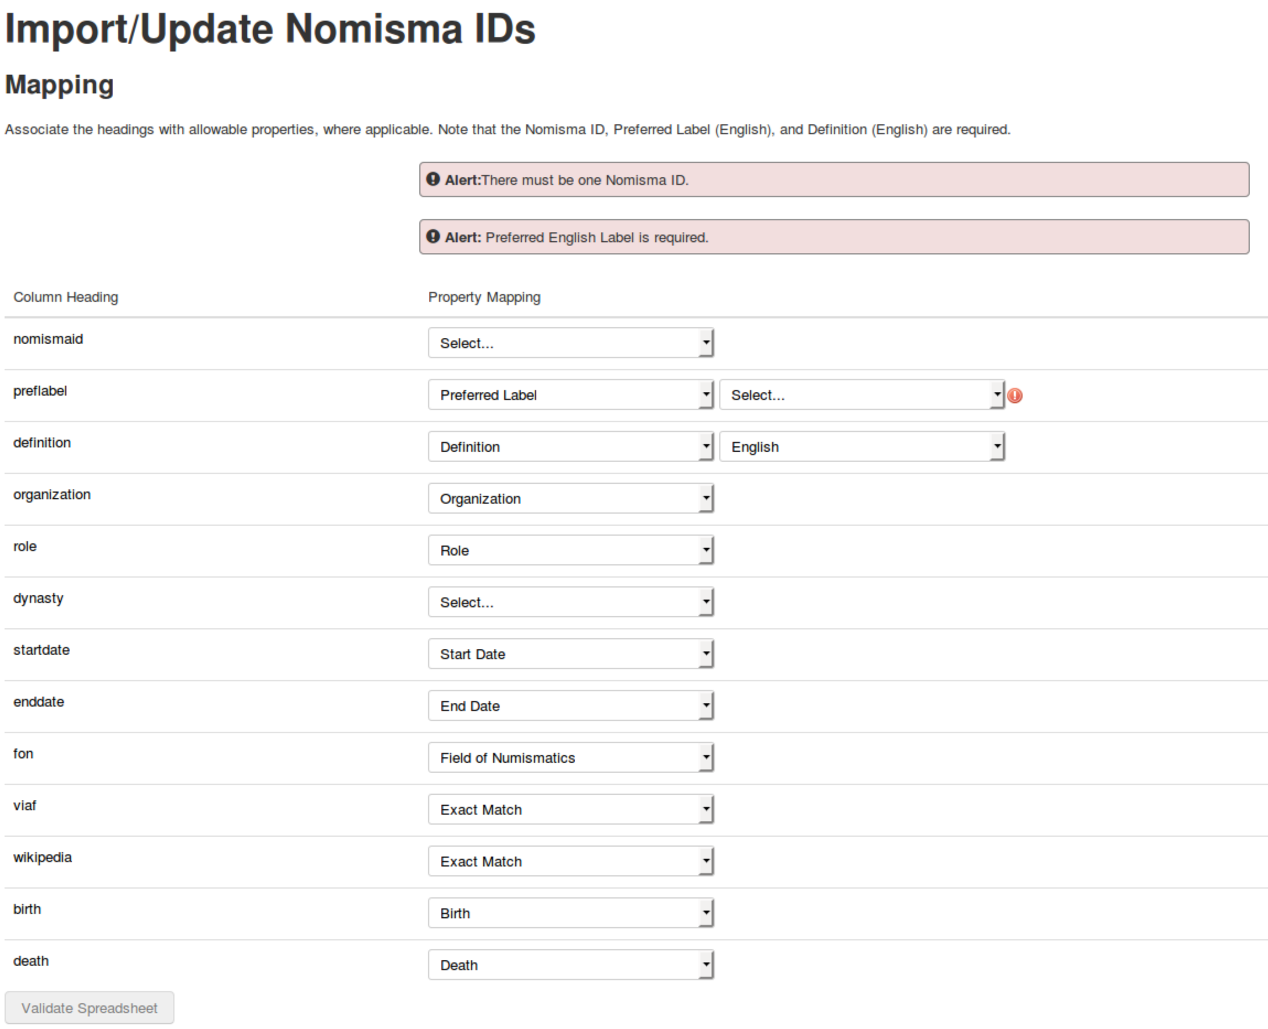

one reason projects like OCRE and CRRO are available in Arabic.

3.3. Public User Interface

The front-end of Nomisma.org is delivered through Orbeon's Page Flow Controller and XML Pipeline Language (XPL). URIs for concepts are constructed by a pipeline that aggregates the source RDF file from the filesystem with, and, depending on whether the concept is mappable, two SPARQL query responses to ascertain whether there are mints or findspots connected with the skos:Concept. This aggregate XML model is passed through XSLT in order to generate an HTML5 web page. The URI for Julius Caesar, http://nomisma.org/id/julius_caesar, therefore includes the RDF transformed into HTML5+RDFa, plus a map rendered with the Javascript library, Leaflet, which shows a layer for mints that struck coins issued by Casesar, a heatmap showing the geographic distribution of all locations where coins of Caesar have been found, and two additional layers (off by default) that show points for hoards (three or more coins found in the same archaeological context) or individual coin finds. These layers are generated by Nomisma APIs that interface with the SPARQL endpoint, passing the XML response from the endpoint through a pipeline to transform it into geoJSON. Additionally, when available, the URI web page will display a list of coin types associated with the skos:Concept, including thumbnails of coins from partner institutions.

The browse page and Atom feed are both generated by XPL which send request parameters to Solr's REST interface, and pipe the XML response through XSLT into the appropriate view (HTML or Atom). There are additional pages for the current and previous versions of the formal ontology, APIs, documentation, and the SPARQL endpoint. We strive to make data available in as many formats through as many protocols as possible. A user may request alternative serializations by appending an extension on a concept URI, e. g., http://nomisma.org/id/augustus.jsonld, to receive JSON-LD by REST, but concept URIs, the browse page, ontology URI, and SPARQL endpoint offer interactivity by Content Negotiation. Content Negotiation is vital for conforming to linked data framework standards, and these advanced HTTP features are possible with Orbeon's XPL.

4. Results

According to a 2014 survey by OCLC, Nomisma.org is one of the most heavily-used Linked Open Data systems in the Library, Archive, and Museum sector, behind OCLC's own VIAF and Worldcat.org services (100,000+ requests per day), and in the same range as the British Library's British National Bibliography service, at between 10,000-50,000 requests per day [OCLC]. Nomisma's load has doubled since then, serving between 40,000-50,000 API calls per day (including about as many SPARQL queries, though there is some overlap between API and SPARQL requests), nearly all of which are from non-search robot machines that facilitate dynamic mapping and multilingual interfaces in a wide variety of numismatic projects. The architecture uses both SPARQL and Solr to their natural advantages while minimizing SPARQL's known scalability limitations. Nomisma.org has suffered approximately three minutes of downtime in three years, which is remarkable considering it runs on a Rackspace cloud server with only 4 GB of RAM and an annual budget of little more than $1,000.

We are aware that we will likely need to upgrade our infrastructure eventually, as our triplestore will grow exponentially in the coming years as more institutions become involved in the numismatic linked data cloud. We strive to continually build more sophisticated query and visualization interfaces that, in turn, require greater server resources. Numismatics has technologically evolved dramatically over the last half-decade, and its march toward Big Data is inevitable, finally making it possible to conduct research in ways that have been dreamed about for centuries.

Bibliography

[WGS84] Dan Brickley, ed. W3C Semantic Web Interest Group: Basic Geo (WGS84 lat/long) Vocabulary. 2003. World Wide Web Consortium (W3C). https://www.w3.org/2003/01/geo/

[FOAF] Dan Brickley and Libby Miller. FOAF Vocabulary Specification 0.99. 14 January 2014. http://xmlns.com/foaf/spec/

[CRM] Nick Crofts, Martin Doerr, Tony Gill, Stephen Stead, and Matthew Stiff, eds. Defntion of the CIDOC Conceptual Reference Model. November 2011. ICOM/CIDOC Documentation Standards Group. http://www.cidoc-crm.org/html/5.0.4/cidoc-crm.html

[GRUBER] Ethan Gruber. Recent Advances in Roman Numismatics. 15 May 2013. MA Thesis, University of Virginia. doi:10.5281/zenodo.45328

[NMDOCS] Ethan Gruber. How to Contribute Data. http://nomisma.org/documentation/contribute

[NMBATCH] Ethan Gruber. Import/Update IDs. 28 July 2015. https://github.com/nomisma/framework/wiki/Import-Update-IDs

[PROVO] Timothy Lebo, Satya Sahoo, and Deborah McGuinness, eds. PROV-O: The PROV Ontology. 30 April 2013. World Wide Web Consortium (W3C). https://www.w3.org/TR/prov-o/

[SKOS] Alistair Miles and Sean Bechhofer, eds. SKOS Simple Knowledge Organization System. 18 August 2009. World Wide Web Consortium (W3C). https://www.w3.org/2004/02/skos/

[NMDATA] Nomisma.org. Nomisma Data. Created 2012. Github.org repository.

https://github.com/nomisma/data

[OSGEO] Geometry Ontology. Ordnance Survey.
http://data.ordnancesurvey.co.uk/ontology/geometry/

[ORG] Dave Reynolds, ed. The Organization Ontology. 16 January 2014. World Wide Web Consortium (W3C).
https://www.w3.org/TR/vocab-org/

[OCLC] Karen Smith-Yoshimuri. Linked Data Survey results 1 – Who's doing it (Updated). 4 September 2014. OCLC.
http://hangingtogether.org/?p=4137