

Investigations as research objects within facilities science

Brian Matthews, Vasily Bunakov, Catherine Jones and Shirley Crompton

Scientific Computing Department, Science and Technology Facilities Council, Harwell
OX11 0QX, United Kingdom

*{Brian.Matthews, Vasily.Bunakov, Catherine.Jones,
Shirley.Crompton}@stfc.ac.uk*

Keywords: Research Data, Research Lifecycle, Data Publication, Digital Preservation, Linked Data, Research Object

Abstract. We consider the notion of data publication in the context of large-scale scientific facilities. Dataset publication allows access to and citation of data, but do not provide sufficient context. We propose instead to publish an investigation, a more complete record of the experiment, including details of the context and parameters of the experiment. We relate this investigation to the emerging concept of a research object, and consider how investigation research objects can be constructed to support the more complete publication of facilities science.

1 Introduction

Data publication is becoming an increasingly accepted part of the future data ecosystem to support research. This involves enabling public access to data by other researchers, with appropriate guarantees of integrity in the management and persistence of the data, and encouraging researchers to cite the use of the data within publications. The intentions behind data publication include: assigning credit and recognition to the collectors of data; encouraging the inspection of data by peers to assess the quality of the data, and to validate the assertions of scientific insights claimed in published articles arising from the analysis of the data; enabling the reuse of the data by other researchers to re-analyse to discover new insights and reportable results, thus furthering the value of the research which arises from the data collection. As a consequence, a number of different approaches and infrastructures have been advocated for data publication (for example [1, 2]).

This is also becoming recognised in the field of “facilities science”. We define facilities science as that science which is undertaken at large-scale scientific facilities, in particular in our case neutron and synchrotron x-ray sources, although similar characteristics can also apply for example to large telescopes, particle physics experiments, environmental monitoring centres and satellite observation platforms. In this type of science, a centrally managed set of specialised and high value scientific instruments is made accessible to a community of users to run experiments which re-

quire the particular characteristics of those instruments. The facilities have their own dedicated staff and funding to supply a scientific service.

In this paper, we concentrate on neutron and x-ray sources, central facilities which supply beams of particles of a particularly intense (such as synchrotron x-ray radiation) or rare (such as beams of neutrons) nature which are otherwise unobtainable for individual research teams. These are then used to analyse the structure of matter at a micro- or nano-scale. These types of facilities differ from other “big iron” [2] science projects in that whilst the facility itself has the characteristics of “big science”, including large long term investments, specialised support teams, large quantities of data, high-performance computing analysis requirements, the science itself is more characteristic of “small science” (or bench science), with many small experiments undertaken by small research teams taking readings of many samples, with diverse funding sources and intellectual objectives. This mixture of characteristics has influenced how facilities are approaching data publication.

In particular, the institutional nature of the facilities, with the provision of support infrastructure and staff, has allowed the facilities to support their user communities by systematically providing data acquisition, management, cataloguing and access, thus providing some of the advantages of “big science” to a small science community. This has been successful to date; however, as the expectation of facilities users and funders develop, this approach has its limitations in the support of validation and reuse, and thus we propose to evolve the focus of the support provided. We propose that instead of focussing on traditional artefacts such as data or publications as the unit of dissemination, we elevate the notion of “investigation” as an aggregation of the artefacts and supporting metadata surrounding a particular experiment on a facility to a first class object of discourse, which can be managed, published and cited in its own right. By providing this aggregate “research object”, we can provide information at the right level to support validation and reuse.

In this paper, we discuss the facilities approach to managing and publishing data, concentrating on our approach which we have been developing at the STFC’s facilities, in particular the ISIS Neutron Spallation Source¹. We then discuss the limitations of this approach, and introduce the concept of an Investigation as a research object, as the unit of publication and access for facilities data. We discuss how this may be represented as Linked Data, comparing it with other similar approaches to research object in the literature. We then further consider how this Investigation may be used, and the tools support which would be required to collate, maintain and preserve such a research artefact.

2 Supporting data management and publication

The neutron and synchrotron radiation facilities support a wide range of different experimental techniques (e.g. crystallography, tomography, spectroscopy, small-angle scattering), and experiments are undertaken within a wide range of different disciplines, including chemistry, bio-chemistry, materials science, earth science, biology, metallurgy, engineering and archaeology. However, from a data management perspective, they all follow similar processes. User scientists apply for an allocation of time on an instrument supported by a science case, which, if accepted, is followed

¹ <http://www.isis.stfc.ac.uk>

by one or more visits to the facility's site where a number of samples, prepared by the user in advance, are placed in the target area, and then exposed to the beam of particles for a desired period of time. During the exposure the beam particles are then blocked or deflected by the sample and then detected by banks of sensors arranged around the target area. These sensors then generate data on such parameters as particle counts, angle of deflection, time-of-flight of the particle, energy, or frequency. This raw data is then streamed off via data acquisition and data management systems which collect, aggregate and move the data to short or long term storage to await further analysis.

Traditionally, this process has been carried out with standard file systems; however, it has been recognised for some time that with the ever increasing data rates and volumes, and increase throughput of experiments, this approach was becoming increasingly hard to manage by hand with the accompanying risk of data loss or corruption. Consequently, we have systematised the process of data management by developing a data catalogue system, ICAT [3]. This cataloguing component, based on an information model capturing a view of a facilities experiment or "investigation" (the Core Scientific MetaData (CSMD) model [4]), within a relational database, provides a common point of gathering information about the experiment. This captures information on the experimental team and intent from the proposal system, and when the experimental visit takes place, will register the data sets, their locations in storage, and experimental parameters. This information is then exposed via an API, either for users to use for browsing and data download on or off site via a web interface (the "TopCat" tool), or else integrating with analysis tools and frameworks so that they can search for and access the data directly. This approach has been successful, and ICAT is being both augmented with additional components and also promoted as an open source tool for use in other similar facilities across Europe and beyond, and has been adopted as the reference data catalogue for the Pan-Data consortium of 13 European facilities².

The image shows a screenshot of a web page titled "ISIS Data" for the DOI "RB820232". The page has a blue header with navigation tabs: "About STFC", "Business & Innovation", "Funding & Grants", "Sites & Facilities", "Our Research", and "Public & Schools". Below the header, there are sub-tabs: "How we operate", "Collaborate with STFC", "Calls, rules & statistics", "A guide to STFC", "Overview of programmes", and "Engaging the public". The main content area is divided into two columns. The left column features a photograph of a laboratory instrument with the text "Data collected on the CRISP instrument at the ISIS facility". The right column contains the following information: "ISIS Data", "RB820232", "Investigation title: Magnetic moment of EuO in spin filtering magnetic tunnel structures.", "DOI: 10.5286/ISIS.E.24066298", "Date of Experiment: Thu Feb 19 13:34:31 GMT 2009", "Publisher: STFC ISIS Facility", "Data format: RAW/Nexus", and a "DOWNLOAD" button with the text "Select the data format above to find out more about it." Below this is a "Data Citation" section with the text: "The recommended format for citing this dataset in a research publication is as: [author], [date], [title], [publisher], [doi]" and "For Example: Griffin. et al. (2009): RB820232, STFC ISIS Facility, doi:10.5286/ISIS.E.24066298". At the bottom of the page, there is a footer with the Research Councils UK logo, the text "Science and Technology Facilities Council", "ISIS User Office: +44 (0) 1235 445592", and a navigation bar with links: "GLOSSARY : SITE-MAP : ACCESSIBILITY : PRIVACY POLICY : ACCESS TO INFORMATION : TERMS OF USE : WEBMASTER".

Figure 1: DOI Landing Page for ISIS

² www.pan-data.eu

Changes to data policy within ISIS have recognised the value of releasing data publicly; data is released for general use after an embargo period of exclusive use to the user. This can be done via the TopCat interface. However, to encourage citation of data and thus attribution and credit for data collection, ISIS issues Digital Object Identifiers (DOIs) for data issued via the DataCite consortium³. Thus for each investigation within ISIS, a DOI is issued, a minimal amount of metadata deposited with the DataCite search engine, and a suitable landing page produced as the “front page” of the data, as for example in Figure 1. From this landing page, given suitable permissions for embargoed data, the data can be accessed. ICAT can provide a stable and quality source of metadata, and a route to archival storage. Thus this provides a suitable data publication channel for ISIS data.

2.1 The changing landscape of facilities science

This established process has been successful for data management and the data publication method via DOIs and landing page, whilst still evolving, should provide a mechanism to support basic data discovery, and support citation of data via DOI and a suitable recommended citation format, thus allowing credit to be attributed to experimenters in traditional publications, and following this, allow the facility to via citation tracking to monitor the value of the use of data generated. However, the landscape of facilities science is changing. We summarise some factors [5].

- Instrumentation and data analysis have become more user friendly than in early days of facilities science. This has led to a lesser significance of the instrumentation “gurus” with a current trend of not including them as the authors of papers; the estimate for biology papers is that about half of them do not now include any facility staff members as co-authors [6] so that new methods and forms may be required for the fair and inclusive attribution of research output.
- The advances of instrumentation and Internet have also led to services allows users to send their samples for remote investigation according to one of the service plans. The sample exposure on a large facility may be just one of the experimental techniques included in the service plan. The service provider then collects the experimental data and supplies them to the user in pre-agreed formats. This implies considering service providers the legitimate agents of facilities science with their inclusion in data management policy.
- Facilities use more than one service to collect data. The user monitoring exercise performed by PaNdata initiative showed that about 7000 (22% of the total) of visitor researchers across Europe have used more than one neutron or synchrotron radiation facility for their investigations⁴. This makes actual the development of common user authentication and user authorization services, as well as experimenting with “virtual laboratories” for the collaborative data analysis.
- New experimental techniques like neutron tomography, or using robots for manipulating multiple samples, or studies of dynamics of materials. The new techniques produce larger volumes of data; they also raise potential opportunities for researchers to perform comparative and multi-aspect studies for the same samples using different experimental techniques, or using the same experimental technique for

³ www.datacite.org

⁴ <http://wiki.pan-data.eu/CountingUsers>

much wider variety of different samples. This scales up all three V's of Big Data: Volume, Velocity, and Variety, and makes their analysis more demanding from modelling and from computational points of view.

- Publishers and scholarly institutions such as the International Union of Crystallography are increasingly requiring traceability of published results through final result dataset to the raw data collected at the facility instrument, so that peers can test the validity of the claimed result.

Thus there is an increasing need to reuse and combine results from different sources; to provide sufficient detail to reviewers so that they can reconstruct the experiment to validate results; and to provide mechanisms to allow credit for various participants in the experimental process, suitable for their role, as in for example [7]. Much of this needs to be mediated via automated tools, so the record of the experiment needs to be available in a machine readable format. The current data publication mechanism based on DOIs and landing pages does not support this well as the context of the data collection, the relationships between various research artefacts, and the different roles of individuals in the process is not captured adequately, so we need to rethink what data publication means in this context.

3 Investigations as Research Objects

Our starting point is to consider the research lifecycle in facilities science, given in schematic form in Figure 2 and given in more detail in [8]. From the point of view of the *Facility* (the user scientist may have a different view of their scientific process) investigations tend to go through the same stages of proposals, preparation, experimental visit, data management, data analysis and visualisation, and publication.

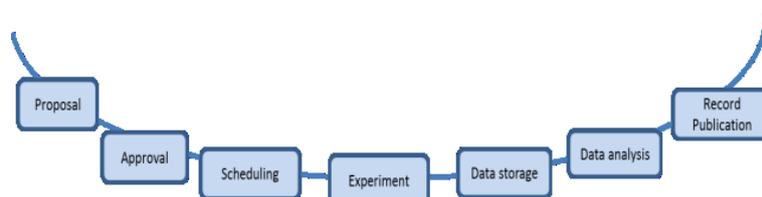


Figure 2: Generic research lifecycle in facilities science

The different stages of research lifecycle produce data artefacts (research proposals, user records, datasets, publications etc.) that are similar across research facilities. Different actors are also involved at the various stages. We also need to record the details of the experiment; which sample was analysed under which experimental conditions, to collect data representing which parameters. Thus by following through the lifecycle of a successful beam time application, we can collect all the artefacts and objects related to it, with their appropriate relationships. As this is strongly related to allocation of the resources of the facility, this is a highly appropriate unit of discourse for the facility; the facility want to record and evaluate the scientific results arising from the allocation is its scarce resources. Thus we propose that the appropriate unit of publication for facilities science is the Investigation.

At one level this is what we already do when we present a landing page for an investigation. Much of the information which is required can be recorded within the

ICAT system. It can support describing which sample was used on which instrument to generate which data set under which experimental conditions to measure which parameters. However, the DataCite metadata does not include these, and while some of this information can be found on the landing page (e.g. instrument) and much more can be found by exploring the detailed metadata in TopCat itself, this is human accessible only, not straightforward to find or navigate, and is not distributed in a machine readable form. Further, related artefacts (derived data, publications, provenance information) is not systematically collected or presented, although now ICAT has the capability to collect this information [9, 10]. What we propose to do is publish the investigation as a single aggregated unit which can be identified and delivered to the user in a machine readable format and contain sufficient contextual information to support discovery of all the components of the investigation and their relationships, so they are available for validation and reuse; that is publish the investigation as a *research object*.

The notion of Research Objects has been explored in a number of projects in recent years (e.g. [11, 12, 13]), and Research Objects have been defined as:

... semantically rich aggregations of resources that bring together data, methods and people in scientific investigations. Their goal is to create a class of artefacts that can encapsulate our digital knowledge and provide a mechanism for sharing and discovering assets of reusable research and scientific knowledge⁵

Research Objects (ROs) as implemented can be seen to have the following characteristics.

- Information about research artefacts and their attributes and relationships are represented as Linked Data; thus RDF is used as the underlying model and representation, with URI used to uniquely identify artefacts. As ROs are linked data objects, they can link into to the existing Linked Data cloud to provide additional context information and be managed by the standard tools of Linked Data and the Semantic Web.
- Standard vocabularies are used to represent relationships describing the research process, such as workflow (workflow4ever⁶), provenance (e.g. Prov-O⁷), and citation (e.g. cito [14]). Use of standard vocabularies encourages shared understanding, enables reuse and allows the use of tools which are tailored for their specialised semantics.
- A bound is provided on the object as an aggregation, so we can determine membership of the research object; typically, OAI-ORE⁸ is used for this purpose.
- The whole research object can be identified via a URI, so its own history and attributes can be related as a first class research artefact in its own right.

The notion of the boundary of a RO is particularly important. A research artefact can be linked to a number of research artefacts. An investigator or instrument can participate in a number of investigations; a publication may use the output of several inves-

⁵ <http://www.researchobject.org>

⁶ <http://www.wf4ever-project.org/>

⁷ <http://www.w3.org/TR/prov-o/>

⁸ <http://www.openarchives.org/ore/>

tigations to support its results. If this is represented as a simple web of linked data, then it would be difficult to distinguish which artefacts and relationships are members of which research object. We need a notion of defining a boundary to determine membership of the RO; OAI-ORE, with its notions of Aggregation and Resource Map provides such a boundary. Research Objects are thus highly suitable as a mechanism to represent and publish Investigations.

4 Building an Investigation Research Object

We outline the major steps of building a research object to represent facility's investigations.

4.1 Representing CSMD in RDF

We can represent the CSMD as an OWL ontology. This will allow us to represent metadata as RDF triples within triple stores (or provide a triple based front end onto metadata databases such as ICAT via for example a SPARQL endpoint) and allows us to publish data about investigations into Linked Open Data. Figure 3 gives a sample of the OWL representation; the full model can be found on the ICAT Google Code site⁹. The OWL representation has a base URI: <http://www.purl.org/net/CSMD/4.0#>

```
<owl:Class rdf:about="csmd:Investigation">
  <rdfs:label>Investigation</rdfs:label>
  <rdfs:comment>An investigation or experi-
ment</rdfs:comment>
</owl:Class>

<owl:Class rdf:about="csmd:Facility">
  <rdfs:label>Facility</rdfs:label>
  <rdfs:comment>An experimental
facilty</rdfs:comment>
</owl:Class>

<owl:Class rdf:about="csmd:Dataset">
  <rdfs:label>Dataset</rdfs:label>
  <rdfs:comment>A collection of data files and part
of an investigation</rdfs:comment>
</owl:Class>

<owl:Class rdf:about="csmd:Datafile">
  <rdfs:label>Datafile</rdfs:label>
  <rdfs:comment>A data file</rdfs:comment>
</owl:Class>
```

Figure 3: A fragment of the CSMD Ontology

⁹ <https://code.google.com/p/icatproject/>

4.2 Constructing an investigation research object

As the facilities lifecycle is enacted within an experiment, we can then construct the research object. Thus, immediately after an investigation has been approved, we can initialise the research object, assigning a DOI at this early stage, and providing some basic information from the proposal, such as instrument used and investigator, as in Figure 4, which also includes a prototypical fragment in RDF-Turtle of the investigation object at this stage.

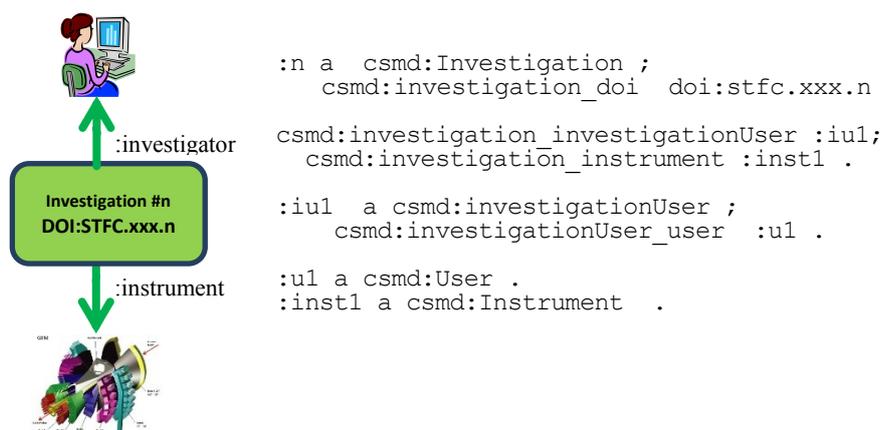


Figure 4: Initialising the Investigation Object

As the experiment is undertaken, we can add further information to the investigation object, to build a more complete picture of the collection of raw data on a sample, again as in a simplified view in the figure below. This step captures the information presented on the current DOI landing page.

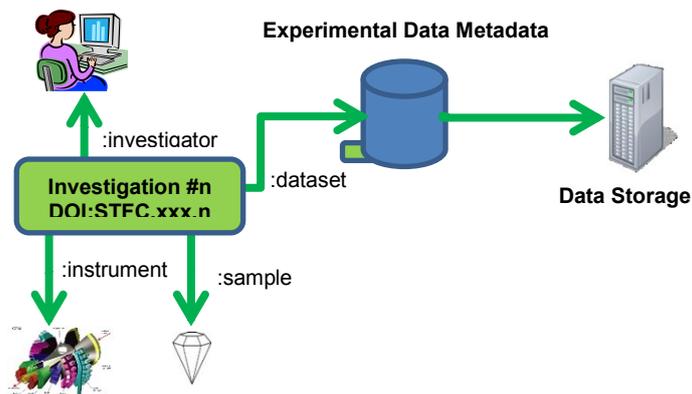


Figure 5: Investigation Object after the Experiment

As the experimental lifecycle goes on, as for example analysis of the data through software packages, and publications and other auxiliary content is added to the investigation, together with the parameters and configurations used, and provenance infor-

mation collected, we can continue to add to the Investigation object, building an eventual object which may contain references to objects in different repositories, ownerships and locations, brought together in a single linked structure as in Figure 6.

Thus this provides a complete picture of the full investigation. This is a dynamic object; further entities could be added it, further derived datasets, publications, or annotations for example as further reuse is undertaken of the research object.

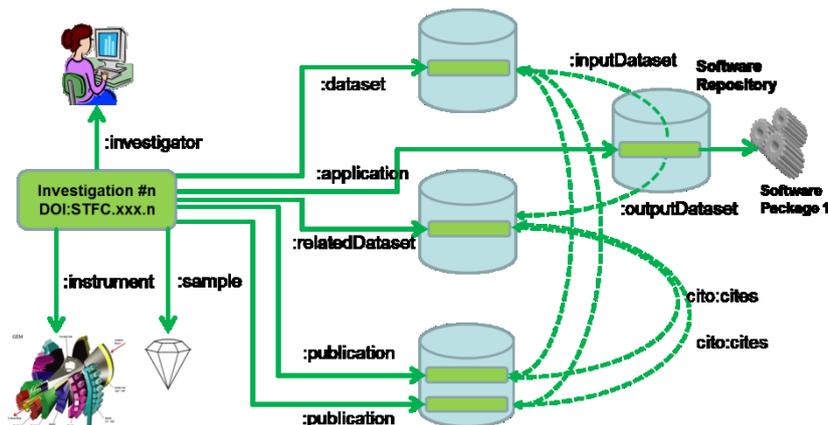


Figure 6: Investigation Object after a complete lifecycle

4.3 Using OAI-ORE as an aggregation constructor

Using the aggregation defined above we have described how an investigation research object can be constructed. However, the research resources within the linked data graph can also be connected to other objects. For example, a publication could use data from several investigations. The publication should be included in each investigation object, but any particular investigation should not include fully the other investigations. Thus we need to provide a boundary. As mentioned above, other approaches have used OAI-ORE to provide a boundary of what is included within the research object, and we propose to follow a similar approach.

OAI-ORE provides some core constructs for capturing aggregations. The class `ore:Aggregation` provides an abstract concept for aggregating resources (`ore:AggregatedResources` in OAI-ORE), with an object property `ore:aggregates` as the combining mechanism. `ore:ResourceMap` describes the aggregation, the resources and the relationships between them. Thus to represent an Investigation Research Object, which is an aggregation, we declare that the Investigation class is a subclass of `ore:Aggregation`:

```
csmd:Investigation rdfs:subClassOf ore:Aggregation .
```

This follows the approach of the Core Research Object Model¹⁰, and thus we can also declare:

```
csmd:Investigation rdfs:subClassOf ro:ResearchObject .
```

¹⁰ <http://www.researchobject.org/ontologies/>

Further, we can declare the core relationships between Investigations and other resources in the CSMD using sub-properties:

```
csmd:investigation_dataset rdfs:subPropertyOf
                           ore:aggregates .
```

We can thus use OAI-ORE to construct the investigation research object with minimal changes to our information model.

5 Using Investigation Research Objects

5.1 Supporting multiple viewpoints

Regardless of discipline there is an acknowledged “life cycle” of research, which is realised in many ways depending on the audience and purposes; a researcher, a funder, a research organisation, a publisher or a preservation institution will focus on different aspects of this life cycle and bring additional contextual links relating to their business process and requirements. For these different stakeholders the central object to which context is added will be different as their world viewpoint is different, for example a publisher will want to establish links from the publication; a funder may wish to do the same for grants. We have described building the links to the investigation from our viewpoint as a facility which is responsible for the creation, discovery and curation of the investigation undertaken at the facility.

The use of research objects supports well this notion of different points of view. Publishing data within a linked open data context in particular makes notions of what constitutes a coherent viewpoint of relevant resources and relationships hard to capture. By providing boundaries and criteria for membership, research objects can support multiple points of view within one data infrastructure. Thus different stakeholders can construct, use and reuse the context relevant to them, and also be credited to the portion of the object which is appropriate to their contribution.

5.2 Data publication

We would propose to use investigation research objects as the unit of publication for our facilities data. Thus we would identify investigation and their related resource maps by persistent identifiers, and use them to generate a landing page. This would be extensible to provide access to the research object in its entirety and include related entities to provide more information in context which could be accessed by other automated agents. Metadata associated with the DOI would need to be changed. Currently, the Datacite metadata field Resource Type supports Dataset and Collection (amongst others), neither of which is correct¹¹ in this context. We would propose that the list of allowed values for this field is extended to include the notion of experiment, study or investigation.

Using the notion of research object as a more open ended bounded object raises the notion of what exactly is being published persistently in this case. If we add additional information are we maintaining stability? Research Objects are well suited to notions of versioning, where we can relate objects together as they change, thus keep-

¹¹ http://schema.datacite.org/meta/kernel-2.2/doc/DataCite-MetadataKernel_v2.2.pdf

ing the old boundary stable. Further, we would propose to have different levels of assurance in our case. The core information on the experiment (sample, instrument, parameters, raw dataset) would remain constant, with other information being secondary and subject to possible extension; this would be made clear in the presentation.

5.3 Data preservation

Shifting the focus from the data to the investigation makes the data preservation activity a more complex one, as it moves from activities relating to the preservation of a well-defined digital object to include not only the digital object but also activities to ensure that the complex linked data, OAI-ORE resource map maintains its integrity and meaning, and links still point to resolvable objects. For preservation purposes it is important that these links are permanent to ensure the integrity of the object.

6 Discussion

The work presented in this paper represents a work in progress. Further discussions are required to agree the correct representation of Investigations as research objects, and design and implementation work to provide tools support so that investigation research objects can be constructed, maintained and published as linked data. However, we see that this could form the basis of a data publication route for facilities data via enhanced landing pages.

Despite the potential for Investigations to become daily “commodities” of the research discourse in facilities science, there are challenges for them, too:

- Universal IDs for Investigations are still a novelty: there is not many of them
- Lack of IDs for other components like instruments, or experimental techniques
- Proto-objects most circulate within a “native” facility (although projects of PANDATA collaboration raise hopes)
- Many researchers, data practitioners, publishers and policy makers are unaware of the potential of Research Objects as intellectual entities

There is also a specific psychological challenge related to the Linked Data representation for Investigations that we mentioned as the most appropriate modelling technique. The flexibility of Linked Data allows re-use and re-combination of its granular parts in other information context so the boundaries of Linked Data information entities tend to be socially defined, and depend upon the role and the interests of the data modeller. Sharing Investigations as Linked Data implies the data publisher agreement for the modifications and derived objects to circulate along with what publisher thought was the best model for Investigation description.

Also the socially defined boundaries of Linked Data objects raise the importance of best practices that should accompany data modelling and system implementation effort. The best practices should be included in structured data curation framework for facilities science that will help the information departments of large facilities to perform a role of a conscious data curator helping to increase data value across the entire research data lifecycle for the variety of stakeholders [15]. Information technologies and services will be then an important means to underpin the data curation role but not the end in themselves.

ACKNOWLEDGEMENTS

This paper is related to the projects of PaNdata (www.pan-data.eu) collaboration, and SCAPE project supported by the EU 7th Framework Programme for Research and Technological Development. The authors would like to thank their colleagues for their input for this paper although the views expressed are the views of the authors and not necessarily of the collaboration.

REFERENCES

1. Lawrence B., Jones C., Matthews B., Pepler S., Callaghan S. Citation and peer review of data: Moving towards formal data publication. *Int. Journal of Digital Curation*, 6(2) (2011)
2. Parsons, M.A., Fox, P.A. Is Data Publication the Right Metaphor? *Data Science Journal* Vol. 12 (2013)
3. Flannery, D., et.al.. ICAT: Integrating Data Infrastructure for Facilities Based Science. e-Science: Fifth IEEE International Conference on e-Science (2009)
4. Matthews, B., et. al. Using a Core Scientific Metadata Model in Large-Scale Facilities. 5th International Digital Curation Conference, London, UK, (2009)
5. Bunakov, V., Matthews, B. Data curation framework for facilities science. 2nd International Conference on Data Management Technologies and Applications (2013)
6. Mesot, J. A need to rethink the business model of user labs? *Neutron News*, 23 (4), (2012)
7. Marcos, E. et al. Author order: what science can learn from the arts. *Communications of the ACM*, , 55(9),39-41 (2012)
8. Matthews, B. et al., Model of the data continuum in Photon and Neutron Facilities. *PaNdata ODI, Deliverable D6.1*. (2012). <http://pan-data.eu/sites/pan-data.eu/files/PaNdataODI-D6.1.pdf>
9. Yang, E. Matthews, B., Wilson, M. Enhancing the Core Scientific Metadata Model to Incorporate Derived Data. *Future Generation Computer Systems* 29 (2) 612-623 (2013)
10. Fisher, S.M., Phipps, K., Rolfe, D. ICAT Job Portal: a generic job submission system built on a scientific data catalogue. 5th International Workshop on Science Gateways, (2013).
11. Bechhofer, S. et al. Why linked data is not enough for scientists. *Future Generation Computer Systems*, 2013, 29(2),599-611 (2013)
12. Shaon, A., Callaghan, S., Lawrence, B., Matthews, B., Osborn, T., Harpham, C. Opening up Climate Research : a linked data approach to publishing data provenance. 7th International Digital Curation Conference (DCC11), Bristol, England, (2011)
13. Belhajjame K, et. al. Workflow-Centric Research Objects: A First Class Citizen in the Scholarly Discourse. In proceedings of the ESWC2012 Workshop on the Future of Scholarly Communication in the Semantic Web (SePublica2012), Heraklion, Greece, (2012)
14. Shotton, D. CiTO, the Citation Typing Ontology J Biomed Semantics. 1(Suppl 1): S6. (2010). doi: 10.1186/2041-1480-1-S1-S6
15. Wilson, M.. Meeting a scientific facility provider's duty to maximise the value of data. In DataCite Summer Meeting, Digital Research Data in Practice (DataCite2012), Copenhagen, Denmark. (2012) <http://epubs.stfc.ac.uk/work-details?w=62852>