

Linked Open Corpus Models, Leveraging the Semantic Web for Adaptive Hypermedia

Ian O’Keeffe^{a*} Dr. Alexander O’Connor^a Philip Cass^a Dr. Séamus Lawless^a
Prof. Vincent Wade^a

^a *Centre for Next Generation Localisation*

Knowledge and Data Engineering Group

School of Computer Science and Statistics, Trinity College Dublin, Ireland

E-mail: {Ian.OKeeffe,Alex.OConnor,casspm,Seamus.Lawless,Vincent.Wade}@scss.tcd.ie

Abstract. Despite the recent interest in extending Adaptive Hypermedia beyond the closed corpus domain and into the open corpus world of the web, many current approaches are limited by their reliance on closed metadata model repositories. The need to produce large quantities of high quality metadata is an expensive task which results in silos of high quality metadata. These silos are often under utilised due to the proprietary nature of the content described by the metadata and the perceived value of the metadata itself. Meanwhile, the Linked Open Data movement is promoting a pragmatic approach to exposing, sharing and connecting pieces of machine-readable data and knowledge on the WWW using an agreed set of best practices. In this paper we present a novel Linked Open Data based personalisation framework that leverages the Web of Data to satisfy the metadata requirements of a Adaptive Hypermedia System. Through the implementation of a prototype system we evaluate the effectiveness of this approach as well as identifying the potential issues that arise from building personalisation systems based on Linked Open Data.

Keywords: Personalisation, Linked Data, Semantic, Adaptation

1. Introduction

In the traditional, closed model approach to Adaptive Hypermedia (AH), it is necessary to define proprietary models that describe aspects of the system and the environment before any adaptation can be carried out. These models include the User Model, the Domain Model and the Adaptation Model (presentation strategy), which all need to be defined in advance. Such systems also require the content, which will be used in the generation of presentations, to be in a specific format and described using a specific metadata standard. This reliance upon bespoke, proprietary content and models restricts the ease of adoption, scalability and accessibility of such technologies. This paper describes an Open Model for AH and Adaptive Web (AW) systems. The aim of the paper is

to specify both the non-functional and functional properties of this model.

The Open Model for AH described by this paper proposes utilising the vast volume of data available on the WWW to address the issues described above. It is proposed to not only gather content from the web, but also to gather the data which describes that content. A number of web technologies have emerged in recent years which make this Open Model achievable. Linked Open Data (LOD) is a practical approach to exposing, sharing and connecting content via the WWW. The structure and links exposed in LOD repositories can be leveraged quickly. This structure has real value in AH as it describes both the content and the domain, which are both basic requirements of AH systems.

To further explore the application of this model, this paper presents a architectural description

and evaluation of a system that implements this open model using AH and LOD techniques and technologies. The system, the Personalised Movie Browser, is a movie recommender that suggests movies that match the user’s preferences, based upon indications of interest. The system uses data from LinkedMDB [1], a project which publishes an open semantic web database for movies, including links to other datasets on the open data cloud and references to related web pages.

1.1. Linked Open Data

The LOD movement is promoting a pragmatic approach to exposing, sharing and connecting pieces of machine-readable data and knowledge on the WWW using an agreed set of best practices [2]. LOD can be characterised using three key properties: (i) the data is published as an RDF data model; (ii) the data is linked to other relevant information through RDF links; and (iii) the data is made available using open, agreed standards [4]. RDF is a data format for defining a directed, labelled graph which represents information available on the WWW. Individual LOD Resources are identified using http URIs. LOD is intended to be data-centric; the focus is on providing large amounts of data in a highly accessible manner. The schema vocabularies used to structure the data should, where possible, reuse existing terms. The most commonly used vocabularies include:

- Dublin Core, which describes documents [5].
- Friend of a Friend, which describes individuals and their friends [6].
- Semantically Interlinked Online Communities, which describes user-generated content, such as forum posts [7].
- An RDF binding of Creative Commons, which is a model for licensing the distribution and reuse of works [8].

There are many existing Linked Data repositories available in a wide variety of disparate domains such as: general encyclopedic knowledge (DBPedia [9], YAGO [10]); Movies (linkedMDB [1]); and academic publications (DBLP RKB Explorer [11]).

SPARQL is the most common query language used with RDF and can be used to express queries across diverse data sources, whether the data is

stored natively as RDF or viewed as RDF via middleware. However, as the use of RDF data sources has increased, performance pertaining to SPARQL queries becomes a cause of concern, particularly where latency is an issue. The demonstrator system described in this paper implements a caching approach for executing SPARQL queries. This approach was evaluated using page load benchmarks and is described in more detail in section 4.2, below.

1.2. Adaptive Hypermedia

Adaptive Hypermedia (AH) combines concepts from Intelligent Tutoring Systems (ITS) and WWW Hypermedia. AH Systems utilise user and domain models to achieve the adaptive selection and sequencing of linked hypermedia content [12]. Early AH systems embedded adaptivity sequencing logic and navigational controls in the content used by the system. As AH systems have attempted to support more dynamic functionality, however, such as adaptivity and personalisation, the content requirements of these systems have become more stringent and the volumes of content consumed far greater. As a result, these systems are being hampered in efforts to deliver such dynamic functionality by their traditional closed corpus nature and reliance upon bespoke, proprietary content and models. To scalably support such system flexibility and dynamism, AH applications require access to large volumes of content which is varied in structure, language, presentation style, etc [13]. These applications also need to reduce the reliance upon bespoke, design time-authored user and domain models.

In an attempt to address the functionality and scalability issues described above, a “second generation” of AH systems emerged. These systems instigated a significant shift towards the detachment of adaptivity controls relating to content sequencing, and the physical content [14]. This development, in principle, enables the integration of externally developed, open corpus, content into these previously closed corpus systems. However, due to proprietary restrictions enforced by such systems on content format and structure, the majority still source all content from system-exclusive repositories of resources. The models used by such systems are also still typically authored at design time. The ability of AH applications to function using con-

tent from open corpus sources has been questioned in recent times, and has been labelled the “Open Corpus Problem” [13].

[3] describes a framework that semantically examines hyperlinked content on both the WWW and in digital content repositories, and creates concept specific caches of content which can be used in the generation of adaptive offerings. However, AH systems which attempt to reap the benefits of such approaches must be able to function without placing complex structural restrictions on content and without the need for design-time knowledge of available resources and the relationships between these resources.

Some educational hypermedia systems, such as KBS Hyperbook [15] and SIGUE [16], allow the incorporation of content sourced from the WWW. However, leveraging such content for use within these systems requires significant manual effort on the part of a domain expert or course designer, in advance of incorporation. The content must be sourced, not a trivial task, and then tagged using the LOM [17] metadata schema and associated with domain model concepts to enable integration [18]. SIGUE, similarly, requires the metadata generation for all new content and manual integration with a domain model. Knowledge Sea II also allows the integration of open corpus content, however, to add an open corpus resource to the system, a comparative analysis must be conducted between the new resource and every existing resource in the collection. A keyword-based similarity metric is used. Knowledge Sea II employs self-organising topic maps to categorise content and create a linked flow between resources [19]. While these AH systems are attempting to utilise open corpus content, none have yet attempted to implement the automated harvesting and use of open corpus knowledge models.

This paper continues by detailing related work in section 2. This related work includes RDF Browsing technologies, Ontology-driven AH and previous “speed of web” studies. The paper then presents an architectural overview of the proposed open model, personalisation framework along with a description of the Personalised Movie Browser system based on this architecture. An evaluation of this system is then presented. Finally, in section 6, a series of conclusions based upon this research are detailed.

2. Related Work

RDF Browsers are an area in which research is currently focused. One such system is /facet [41], a generic faceted browser whose primary features include easy browsing and selection of properties of other, semantically related, types handling of any RDFS dataset without extra configuration and effective search system for the dataset. Another system is LENA [42], which stands for LENS based Navigator. A “lens” represents a particular view onto RDF data and as described by the Fresnel Display Vocabulary. LENA enables viewing RDF data in your web browser and supports SPARQL selectors. Longwell [43], an RDF powered highly-configurable browser. While these browsers allow for effective browsing of RDF ontologies and metadata none deal with the more technical issues when dealing with remote data sources, primarily caching.

The application of semantic web techniques in Adaptive Hypermedia systems is a natural step as both aim to achieve personalisation. The degree to which AHS have integrated semantic web techniques varies from the use of ontologies to describe concept domains such as in ADAPT² [32] and the PNS [33] to the use of RDF as the medium for all of the metadata models that the AHS relies on, leveraging the power of the technology through semantic search techniques [34]. Where RDF has been used to such a degree it has typically required the RDF to be written specially for the application as the information was not available. The recent emergence of the Web of Data has seen a renewed interest in this area of research. In [40], the potential for the application of LOD techniques in User Modelling and Adaptation domains is discussed. The authors speculate about the utilisation of LOD for the realisation of the Domain Models and User Models, which form key components of any AHS. Beyond the AH domain, Linked Data has been used to support the delivery of personalised services as part of recommender systems [35] as well as for personalised television as part of the NoTube project [36].

There has been substantial research in the area of performance for Linked Data. In particular, there have been trials with extremely large collections of data, as well as specific data sources. The predominant methodologies have focused on comparing the server-side through reasoning [30], re-

trieval [31] or a combination of those features [23] on common data [29], or other query optimisations [24]. This kind of approach is of great value because it allows an objective perception of the performance of different Linked Data architectures. However, there has not been as much focus on in-situ measurement of performance in the Linked Open Data Web.

There have been studies on the nature of the content visible on the Linked Open Data Web and the Semantic Web [27], for example with a measurement of the change frequency of the information available [25], or the informative nature of particular URIs [26]. This is necessarily a more subjective study, since it depends on the complex blend of factors on both the client and server side, as well as network issues. Nonetheless, there is a need for application developers to have a good sense of the performance of the Linked Data repositories that are available. This paper will present a small initial study on the speed of Linked Data ‘in vivo’.

3. Architecture

In this section we present the architecture for the Open Model Personalisation Framework. As discussed previously, the aim of this system is to provide the user with a personalised composition while addressing some of the limitations of existing personalisation approaches such as those applied by Adaptive Hypermedia Systems. First we present the architecture of the Open Model Personalisation Framework followed by a discussion of an example application, the Personalised Movie Browser, which has been implemented based on this framework.

3.1. Open Model Personalisation Framework

The main component of the Open Model Personalisation Framework, shown in figure 1, is the Adaptive Engine (AE), this is a rule based system, based on JBoss Rules [44], that integrates metadata model access and storage functionality with the rule base, providing direct access to the metadata. In addition, the AE facilitates the execution of queries across LOD repositories as part of the rule base. The AE supports the adaptation process through the execution of an Adaptation

Strategy (AS). The AS can be considered similar to the Adaptation Model used in other AHS with the notable difference that the rules that make up the AS are not only scaffolded on the structure of the domain but also on a strategy that describes how the content composed by the system should be sequenced. The significant advantage of this approach is that the architecture can be applied in domains in which there is a need for the sequencing of a personalised composition to be influenced by a domain expert rather than being entirely based on the relationships that exist between concepts in the domain. This is a concept that has been adopted from the application of AHS in the eLearning domain [14]. The rules that make up the AS are designed to be independent of any physical resources, e.g. content. Being abstract facilitates the use of the AS over a wide range of content sources as well as the utilisation of content from a variety of open corpus sources. As such, the reusability of the AS is greatly enhanced.

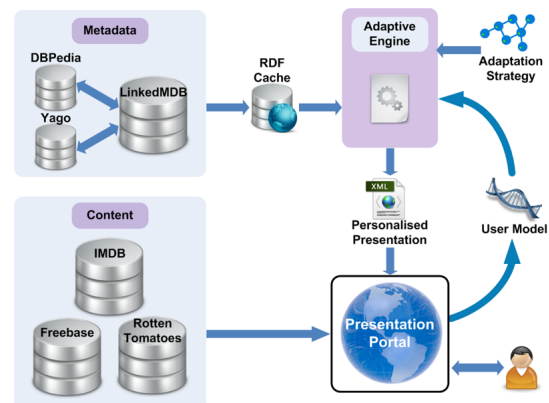


Fig. 1. Architecture Diagram for the Open Model Personalisation Framework

The adaptation process carried out by the AE is influenced through the use of the User Model (UM), which provides the AE with information about the learner that is pertinent to the adaptation process. This model is not intended to capture every detail of the user that the system could potentially adapt to but rather just the properties that are of interest, for example the user’s preferred movie genres. The UM is captured in an open XML format, allowing the AE to support the use of a wide range of standards and specifications. More concrete examples of the use of the

User Model can be found in the discussion of the Personalised Movie Browser application.

In a traditional AHS, the adaptation process would also involve the use of a content model and/or domain model to provide information about available resources and the relationships between concepts in the domain. In the framework presented here both of these metadata models are replaced by the use of LOD repositories that provide access, through a SPARQL query interface, to metadata that is encoded in RDF.

This approach provides two significant advantages, the first is that, by its very nature, the Linked Data describes the structure of the domain providing the system with information about the relationships between concepts. The second advantage is that it removes the need to generate metadata that describes the concepts or resources that are used by the application. This is arguably one of the most significant issues affecting the uptake of AH based systems as the manual generation of metadata for large content bases quickly becomes an expensive task [20] while automatic approaches to generating metadata typically depend on consistent document content, context or metadata [21].

As discussed later in this paper, see section 4.2, the use of LOD does require some special consideration. As all of the computation required to carry out the necessary SPARQL queries, which can be quite complex, is offloaded to the service provider. As such, the performance of the service can be difficult to guarantee, especially if the query spans across more than one LOD repository. To mitigate this issue, the implemented system utilises Memcached [39] as a RDF cache, which sits between the AE and the LOD repository. This provides improved performance for any queries that have previously been run.

The outcome of the adaptation process is a model describing the user’s personalised composition, which the user can access through a web based User Portal. The composition describes the specific resources that were selected for the user and how they should be sequenced for the user. These content resources can come from either closed content repositories or more open repositories available on the internet. The Presentation Portal also provides additional functionality such as tools through which the user can provide the

system with information about themselves, allowing them to influence the adaptation process.

3.2. Personalised Movie Browser

In order to validate the proposed architecture and to serve as a research platform for further experimentation in the application of Linked Data, the Personalised Movie Browser application was developed. This web based application is designed to provide personalised movie recommendations based on the user’s favourite movies. The Personalised Movie Browser makes use of the LinkedMDB LOD repository as its source of metadata about movies. LinkedMDB provides a significant amount of information, such as actors, directors, producers, synopsis, about a large collection of movies as well as providing links to content resources from sites such as IMDB and Freebase.

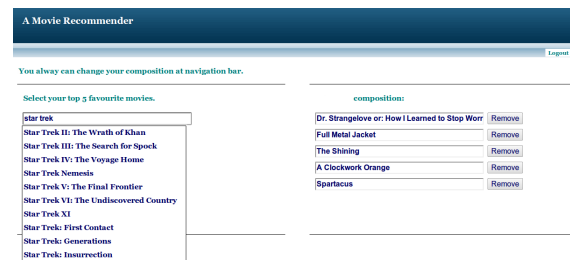


Fig. 2. Screenshot of the Personalised Movie Browser User Preferences Page

Figure 2 shows a screen capture of the user preferences page, which allows the user to specify a list of 5 movies (shown on the right hand side). The movies are selected using search box that provides auto-complete functionality based on the set of movies that the system knows about. This allows the user to select movies from the full range of movies described by LinkedMDB while ensuring that the task is not overly complex. In addition, this technique ensures that the movies selected by the user are valid entries for which searches of LinkedMDB will return results.

When the user is happy with their movie selections the system generates a personalised set of movie recommendations. The first step in this process is to identify the genres of films that the user is interested in, this is achieved by looking up the genres of the movies that the user has specified from linkedMD. An example of the SPARQL query

```

select ?genre where{
  ?movieURL dc:title [movieTitle].?movieURL
movie:genre ?url.?url
movie:film_genre_name ?genre.

```

Fig. 3. Example SPARQL Query

used to achieve this is provided in figure 3. The query shown is dynamically modified by the AE by changing the variable, denoted by square brackets [movieTitle]. The genres identified are ranked depending on the number of times they occur in the users selections. Based on the identified genres, the system then queries linkedMDB for movies that match the identified genres.

Figure 4 shows a screen capture of the personalised movie presentation for an individual user. On the left hand side of the page a synopsis of the movie is provided. This content is retrieved by dynamically querying linkedMDB using a HTTP endpoint. The user is also provided with a set of visualisation components on the right hand side of the screen that allow the user to view the information that the system has captured about them, their preferred movies, as well as the information that the system has inferred about them, their preferred genres. The degree to which the system believes a user likes a genre is shown using a 5 star rating. The user can also add films to their profile through a ‘like’ mechanism.

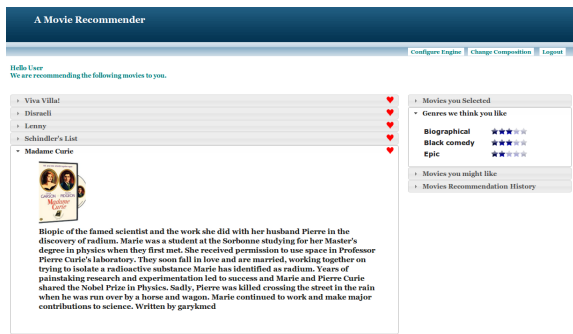


Fig. 4. Screenshot of the Personalised Movie Browser Content Page

4. Evaluation

The inclusion of distributed information sources in any user-facing application can reduce usability

by introducing delays, unreliability and distortion in the interface. In particular, because adaptive application seek to perform rich transformations on the web experience of the user based on their attributes, it is vital that the adaptive system be responsive and reliable. These concerns arise directly in the use of Linked Open Data for adaptivity, because such systems depend on querying remote knowledge bases with differing infrastructure and with often-complex query expressions. In order to address this concern, two evaluations were performed. The first was to attempt, at a qualitative level, to assess what kind of delay remote querying introduces. The second describes the effect of generic web caching techniques on improving linked data adaptive system performance. This shows that there are approaches to addressing the problem of speed and reliability that can be leveraged from existing tools, rather than necessarily needing to re-engineer specific LOD-AH methodologies.

4.1. Speed-of-web study

The relatively high cost of querying linked open data remotely presents an important design consideration, particularly for knowledge-intensive applications such as Adaptive Hypermedia environments.

The architecture presented in this paper demonstrates a system on one repository, the Linked Movie Database. However, the nature of the web dictates that different linked data endpoints for different content and platforms will have different performance profiles. In order to gain a qualitative perspective on the cost of remote linked data querying, two representative queries were chosen and executed on a number of repositories, chosen from across the LOD cloud. While these results are not statistically representative, they do provide an initial view of the “speed of the LOD web”.

The queries executed were designed to be representative of the kinds of operation relevant to an LOD AH environment, specifically two tasks: retrieving a long list of entities, and retrieving the detail of particular entities.

These queries are represented graphically in fig.5 One of the key challenges for creating representative queries was to define queries of approximately equivalent semantics. In order to achieve this, the three linked data sources chosen were ex-

amined to locate somewhat equivalent entities and attributes of similar cardinality and relationship to the entities.

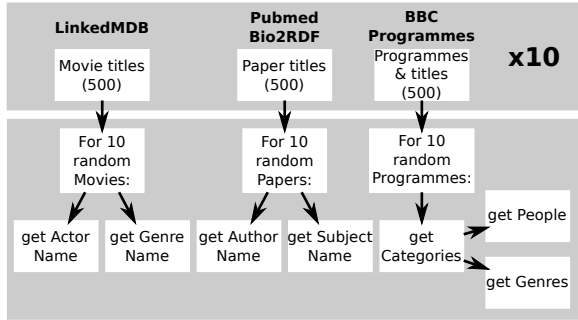


Fig. 5. Schematic representation of the queries used to measure the response of different LOD repositories. Note that each set of queries was run three times, and that the results were serialised as JSON

For the LinkedMDB data source, Movie entities were chosen as the basic object for retrieval. A query for 500 titles for entities of the Movie type were retrieved. This query was repeated ten times, and the aggregate time recorded.

From this list of retrieved titles, ten Movies were chosen at random, and the list of Actor names and Genre names for the Movie of a particular title were retrieved. Similarly, for Pubmed Bio2RDF [37], JournalArticle entities were chosen for retrieval, with paper authors and subject keywords were selected. These are approximately similar attributes in terms of cardinality and semantics to the Movie entities.

The BBC Programme data source [38] differed slightly in the way it represented the categorisations. The basic entity under investigation was the Episode, but the graph differed from the other two data sources by having one category relationship, and the different between Persons and Genres was decided by type inference. This also accounts for a high variance in the result timings because of a higher variance in the cardinality of the results. This is a good example of the challenges of attempting to uniformly assess different data sets: modelling differences can make equivalence only approximate.

4.1.1. Results

The query procedures on each Linked Data source were repeated three times, on separate days. The results were serialised as JSON. The ar-

chitecture for this test harness was based on the SPARQL Endpoint Interface for Python [22] and used Python’s built-in Timer library.

There are many inherent difficulties with effective performance testing, particularly with regard to http-based interfaces, where caching, proxies, network issues and other factors complicate the repetition of trials. Because of this, aggregate times for repetitions of individual times were recorded separately. This provides an informative, rather than objective guide to likely timing for queries. The results are shown in figures 6 and 7 below:




Fig. 6. Time (sec) for sequences of genres for 10 random items

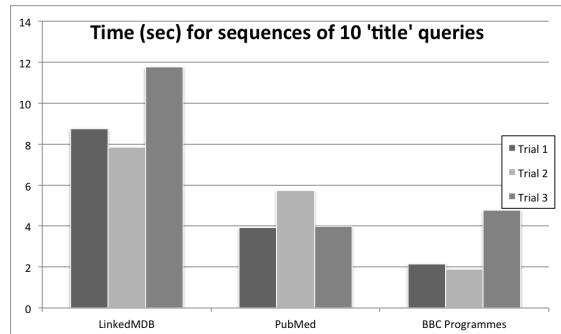


Fig. 7. Time (sec) for sequences of 10 title queries

These results show that the performance of remote semantic querying can vary depending on a complex series of known and unknown effects. This has implications for the design of adaptive systems, which point towards a need for caching strategies for complex or repetitive queries. Another factor observed was the limitation of connection rates by the end-points. One endpoint returned rate limits after 10 successive queries in a

12 second period. This points to another key non-functional design parameter: supporting rate limited access to end-points.

4.2. Cache

A caching mechanism for the results of SPARQL queries. This was implemented order to improve the performance of the Open Model Personalisation Framework. In the realm of medium to large scale online applications caching is often applied as a solution to improving application performance and while some work has been done on the subject with regard to local data sources under control of the developer, there has not been much research on the subject of caching of external resources, especially on the open web.

A caching system using a key-value associative array was implemented and invoked whenever a SPARQL query is made on a LOD repository. On a HTTP / SPARQL request the program generate a unique hashcode based on the request. It would search for this HTTP URL/SPARQL hashcode in the cache. If not found the request would continue as normal and the result stored in the cache using the hashcode of the request as a unique identifier. If found, the result would be requested from the cache.

In general, the model for caching was one which assumed that the data sources were relatively constant and unchanging. This is broadly representative of application areas such as movies, where the information does not change often on a particular movie. To keep results fresh a timestamp was added to the hashcode of the requests. This timestamp consisted of the current date. In this implementation, a 24-hour freshness measure was used to mark stale cache entries and refresh them. While arbitrary, this scale of delay appeared to prove an adequate trade-off between caching and freshness.

Figure 8 plots loading time of the main landing page, which contains content such as movie posters and film summaries extracted from IMDB and free-base.com. The bar chart reflects the difference in loading in milliseconds between cached and uncached requests, showing an 19.5 times decrease in load times (on average) when cached over five trials.

Figure 9 displays the difference in loading SPARQL requests from the cache and requesting




Fig. 8. Loading time for main landing page containing content

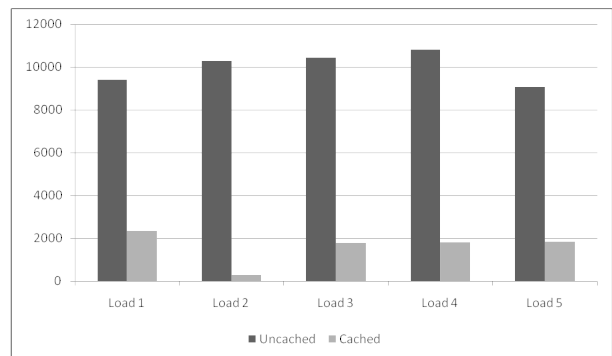


Fig. 9. The difference in cached SPARQL request times versus uncached.

them uncached, showing a 6.2 times decrease in load times (on average) when cached over five trials.

These results are based on generic genre queries e.g requesting a list of Action movies. For user specified requests, where there are specific constraints (like year made, or director) the likelihood of getting the same query (and thus a cache hit) is low. However caching of external additional data helps build a tolerance of external faults and thus improves the robustness of the system.

5. Conclusion

This paper has presented an architecture for driving multi-model adaptive personalisation leveraging Linked Open Data. There are several important advantages to using LOD for Open Model Adaptive Hypermedia. The first, and most important, is the large quantity of high quality, curated data available on the Linked Data Web. The avail-

ability of this data overcomes one of the key challenges of effective AH: the requirement for high-quality structured data.

However, there are also some important practical considerations to be taken into account in implementing effective Open Model AH systems. The high variability of the response performance of LOD repositories motivates further research into repository and client performance, particularly through caching.

It is important for future architectures for Open Model AH, and LOD in general, to be resistant to variable responses and which are able to avoid excessive load on LOD repositories.

The personalised nature of Open Adaptive Hypermedia means that it is an interesting example of the general problem of deciding how to store portions of large Linked Data sets to improve client performance. The challenges of personalisation are also applicable to other LOD use cases.

6. Acknowledgement

This research is supported by the Science Foundation Ireland (Grant 07/CE/I1142) as part of the Centre for Next Generation Localisation (www.cngl.ie) at Trinity College Dublin.

References

- [1] Hassanzadeh, O. and Consens, M. P. (2009). Linked movie data base. In Proceedings of the WWW2009 workshop on Linked Data on the Web (LDOW2009).
- [2] Bizer, C., Cyganiak, R., and Heath, T. (2007). How to publish linked data on the web. Retrieved June, 20:2008.
- [3] Lawless, S. "Leveraging Content from Open Corpus Sources for Technology Enhanced Learning". Ph.D. Thesis, Submitted to the University of Dublin, Trinity College, 2009.
- [4] O’Connor, A. "Context-Informed Semantic Interoperation". Ph.D. Thesis, Submitted to the University of Dublin, Trinity College, 2010.
- [5] Nilsson, M., Powell, A., Johnston, P., and Naeve, A. (2008). "Expressing Dublin Core metadata using the Resource Description Framework (RDF)". DCMI Recommendation.
- [6] Brickley, D. and Miller, L. (2005). "FOAF vocabulary specification". Namespace Document, 3.
- [7] Breslin, J., Harth, A., Bojars, U., and Decker, S. (2005). "Towards semantically-interlinked online communities". In The 2nd European Semantic Web Conference (ESWC’05), Heraklion, Greece, Proceedings, LNCS, volume 3532, pages 500-514. Springer.
- [8] Abelson, H., Adida, B., Linksvayer, M., and Yergler, N. (2008). "ccREL: The Creative Commons Rights Expression Language". Creative Commons Wiki.
- [9] Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., and Ives, Z. (2007). "Dbpedia: A nucleus for a web of open data". Lecture Notes in Computer Science, 4825:722.
- [10] Suchanek, F., Kasneci, G., and Weikum, G. (2007). "Yago: a core of semantic knowledge". In Proceedings of the 16th international conference on World Wide Web, pages 697-706. ACM New York, NY, USA.
- [11] Glaser, H. and Millard, I. (2007). "Rkb explorer: Application and infrastructure". Proceedings of Semantic Web Challenge.
- [12] Brusilovsky, P. "Adaptive Educational Systems on the Worldwide Web: A Review of Available Technologies". In the Proceedings of the Workshop "WWW-Based Tutoring", at the 4th International Conference on Intelligent Tutoring Systems, ITS’98, San Antonio, TX, USA. 1998.
- [13] Brusilovsky, P. and Henze, N. "Open Corpus Adaptive Educational Hypermedia". In The Adaptive Web: Methods and Strategies of Web Personalisation, Lecture Notes in Computer Science, vol. 4321, Berlin: Springer Verlag, pp. 671-696. 2007.
- [14] Conlan, O. "The Multi-Model, Metadata Driven Approach to Personalised TEL Services". Doctoral Thesis, Submitted to the University of Dublin, Trinity College, 2005.
- [15] Henze, N. and Nejdil, W. "Extendible Adaptive Hypermedia Courseware: Integrating Different Courses and Web Material". In the Proceedings of the International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems, AH2000, pp. 109-120, Berlin: Springer-Verlag, Trento, Italy. August 28th-30th, 2000.
- [16] Carmona, C., Bueno, D., Guzmán, E., Conejo, R. "SIGUE: Making Web Courses Adaptive". In Proceedings of 2nd International Conference on Adaptive Hypermedia and Adaptive Web Based Systems, AH2002, Malaga, Spain, 29-31 May, 2002. Lecture Notes on Computer Science, Vol. 2347. Berlin: Springer Verlag, pp. 376-379. 2002.
- [17] "IEEE Learning Object Metadata". Available online at <http://ltsc.ieee.org/wg12/>
- [18] Henze, N. and Nejdil, W. "Adaptation in Open Corpus Hypermedia". In the International Journal of Artificial Intelligence in Education, Special Issue on Adaptive and Intelligent Web-Based Systems, vol. 12, pp. 325-350. 2001.
- [19] Brusilovsky, P., Chavan, G., Farzan, R. "Social Adaptive Navigation Support for Open Corpus Electronic Textbooks". In Proceedings of 3rd International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems, AH 2004, P. DeBra, W. Nejdil (Eds.), Lecture Notes in Computer Science, Vol. 3137. Berlin: Springer Verlag, pp. 24-33. 2004.
- [20] Ben Steichen, Séamus Lawless, Alexander O’Connor, and Vincent Wade. 2009. "Dynamic hypertext generation for reusing open corpus content". In Proceedings of the 20th ACM conference on Hypertext and hyper-

- media (HT '09). ACM, New York, NY, USA, 119-128. DOI=10.1145/1557914.1557937
- [21] Melike Sah and Vincent Wade, “Automatic Metadata Extraction from Multilingual Enterprise Content”, ACM International Conference on Information and Knowledge Management (CIKM 2010), Toronto, Canada, 2010, pp. 1665-1668
- [22] Herman, I.; Fernández, S.; Tejo, C. “SPARQL Endpoint interface to Python”, <http://sparql-wrapper.sourceforge.net/>
- [23] Thakker, D., Osman, T., Gohil, S., Lakin, P, “A Pragmatic Approach to Semantic Repositories Benchmarking”. In Proc. of the 7th Extended Semantic Web Conference, ESWC 2010.
- [24] Hartig, O., Bizer, C. and Freytag, J.; “Executing SPARQL Queries over the Web of Linked Data THE SEMANTIC WEB” - ISWC 2009 Lecture Notes in Computer Science, 2009, Volume 5823/2009, 293-309, DOI: 10.1007/978-3-642-04930-9_19
- [25] “Towards Understanding the Changing Web: Mining the Dynamics of Linked-Data Sources and Entities”
- [26] Halpin, H and Presutti, “V: An Ontology of Resources for Linked Data”. Proceedings of the WWW2009 Workshop on Linked Data on the Web (LDOW 2009), Madrid, Spain, April 20, 2009.
- [27] L. Ding and T. Finin. “Characterizing the semantic web on the web”. In Proceedings of ISWC 2006, Athens, GA, USA, November 5-9, 2006, volume 4273 of LNCS, pages 242–257. Springer, 2006.
- [28] K. Wilkinson, C. Sayers, H. A. Kuno, and D. Raynolds. Efficient RDF Storage and Retrieval in Jena2. In SWDB 2003.
- [29] C. Bizer and A. Schultz. Benchmarking the performance of storage systems that expose SPARQL endpoints. In Proceedings of the ISWC Workshop on Scalable Semantic Web Knowledge- base systems. 2008.
- [30] Oren, E, Kotoulas, S, Anadiotis, Siebes,R, ten Teijea, A and van Harmelen, F.: Marvin: Distributed reasoning over large-scale Semantic Web data Web Semantics: Science, Services and Agents on the World Wide Web Volume 7, Issue 4, December 2009, Pages 305-316
- [31] Auer, S., et al.: Triplify Light-Weight Linked Data Publication from Relational Databases. In: Proceedings of the 18th International World Wide Web Conference (2009)
- [32] Peter Brusilovsky, Sergey Sosnovsky and Michael Yudelson. Ontology-based framework for user model interoperability in distributed learning environments. pages 28512855. AACE, 2005.
- [33] Conlan O, O.Keeffe I, Tallon S. Combining Adaptive Hypermedia Techniques and Ontology Reasoning to Produce Dynamic Personalized News Services. In: Wade V, Ashman H, Smyth B, eds. Adaptive Hypermedia and Adaptive Web Based Systems. Springer-Verlag; 2006:81-90.
- [34] N. Henze, “Personal Readers: Personalized Learning Object Readers for the Semantic Web,” Proceeding of the 12th International Conference on Artificial Intelligence in Education, IOS Press, 2005, pp. 274-281.
- [35] B. Heitmann and C. Hayes, “Using Linked Data to Build Open , Collaborative Recommender Systems”. Proceedings of the AAAI Spring Symposium Linked Data Meets Artificial Intelligence, 2010.
- [36] NoTube: making the Web part of personalised TV., Balthasar Schopman, Dan Brickley, Lora Aroyo, Chris van Aart, Vicky Buser, Ronald Siebes, Lyndon Nixon, Libby Miller, Vénique Malaise, Michele Minno, Michele Mostarda, Davide Palmisano and Yves Raymond. In Proceedings of the WebSci10: Extending the Frontiers of Society Online. Apr 2010. [Publication]
- [37] Nolin, M.-A., Ansell, P., Belleau, F., Idehen, K., Rigault, P., Tourigny, N., Roe, P., Hogan, J.M. and Dumontier, M. (2008) Bio2RDF Network of Linked Data. Semantic Web Challenge at International Semantic Web Conference (ISWC 2008). Karlsruhe, Germany
- [38] BBC Backstage programmes API <http://backstage.bbc.co.uk/>
- [39] Danga Interactive. Memcached. <http://memcached.org/>
- [40] L. Aroyo and G.-J. Houben, “User modeling and adaptive Semantic Web”. Semantic Web Journal, vol. 1, 2010, pp. 105-110.
- [41] Hildebrand, M. and van Ossenbruggen, J. and Hardman, L., “/facet: A browser for heterogeneous semantic web repositories”. The Semantic Web-ISWC 2006, pages 272-285
- [42] Koch, J. and Franz, T. and Staab, S. 2008. “Lena - browsing rdf data more complex than foaf”, In International Semantic Web Conference, Demo Session
- [43] SIMILE: Longwell RDF Browser (2003-2005) <http://simile.mit.edu/wiki/Longwell>
- [44] JBoss Rules (Drools), JBoss Community, <http://www.jboss.org/drools>