

Linked USDL: A Vocabulary for Web-scale Service Trading

Carlos Pedrinaci¹, Jorge Cardoso^{2,3}, and Torsten Leidig⁴

¹ Knowledge Media Institute, The Open University, Milton Keynes, United Kingdom
`carlos.pedrinaci@open.ac.uk`

² CISUC/Dept. Informatics Engineering, University of Coimbra, Coimbra, Portugal

³ Technical University of Dresden, Dresden, Germany
`jcardoso@dei.uc.pt`

⁴ SAP Research Karlsruhe, Germany
`torsten.leidig@sap.com`

Abstract. Real-world services ranging from cloud solutions to consulting currently dominate economic activity. Yet, despite the increasing number of service marketplaces online, service trading on the Web remains highly restricted. Services are at best traded within closed silos that offer mainly manual search and comparison capabilities through a Web storefront. Thus, it is seldom possible to automate the customisation, bundling, and trading of services, which would foster a more efficient and effective service sector. In this paper we present Linked USDL, a comprehensive vocabulary for capturing and sharing rich service descriptions, which aims to support the trading of services over the Web in an open, scalable, and highly automated manner. The vocabulary adopts and exploits Linked Data as a means to efficiently support communication over the Web, to promote and simplify its adoption by reusing vocabularies and datasets, and to enable the opportunistic engagement of multiple cross-domain providers.

Keywords: Services, Vocabulary, Linked Data, USDL, eCommerce

1 Introduction

The importance of real-world services, that is business activities of a mostly intangible nature (e.g., life insurance, consulting), has grown over the last 50 years to dominate economic activity [1]. Because of their intangible nature, services can often be bundled, adapted, and traded in an automated manner. In an attempt to exploit the Web as a service trading platform a number of service marketplaces have emerged, ranging from purely technical registries like UDDI [2], to business-oriented marketplaces such as Google Helpouts. Technical registries have for the most part focussed on the computer science aspects of services which is limiting as it ignores fundamental characteristics of services including the economic, social, and business contexts [3]. Business-oriented marketplaces on the other hand have focussed on providing silos that offer limited search and

comparison capabilities through an essentially human-oriented storefront [4]. As a result, common and essential economic activities in the service sector such as the generation of customised offerings, the creation and trading of possibly cross-domain and multi-provider service bundles, or simply the communication between customer and provider remain largely manual activities [4].

Supporting the trading of services over the Web in an open, scalable, and automated manner enabling the opportunistic engagement of multiple cross-domain providers requires a shared means for capturing and reasoning upon the economic, social, and technical aspects governing service exchanges [1, 3, 4]. The Unified Service Description Language (USDL) is the most comprehensive attempt in this direction but it has received limited adoption due to its complexity, while it also exhibited limitations with respect to the level of extensibility and automation supported. In this paper we present Linked USDL⁵, a new vocabulary which builds upon the results and experience gained with USDL combined with prior research on Semantic Web Services, business ontologies, and Linked Data to better support Web-scale automated service trading. We present the methodology and main decisions adopted for transforming the complex USDL specification into a network of vocabularies that is anchored on simplicity as well as on vocabulary and data reuse. The resulting vocabulary is thoroughly evaluated in terms of domain coverage, suitability for purpose, and its current level of adoption.

2 Related Work

Service Science aims to reach a better understanding of services, service networks, value co-creation and service innovation, to name a few of the main research topics [1]. These efforts, which encompass several disciplines, are geared towards establishing solid foundations for advancing our ability to design, create, and analyse service systems with both business and societal purposes in mind.

Relevant work in Computer Science includes service-oriented systems, which approach the development of complex applications by integrating networked software components called Web services [2]. This area has been prolific in terms of both tooling and specifications including a number of approaches for describing technical services semantically, e.g., OWL-S, SAWSDL, and WSMO [5, 6]. Although (semantic) Web services work provides advanced support for discovering or composing technical services, it disregards the fundamental socio-economic context of real-world services (e.g., value chains and offerings), and does not cover the widespread manual services (e.g., consulting) [3]. Complementary work on Workflow and Business Process Management has focussed on the operationalisation of the processes within enterprises [2, 3, 5], which has more recently also incorporated human activities [7]. This work is, however, centred on a procedural view on how activities are carried out within an organisation which is orthogonal to the business characteristics of the services offered (e.g., speed of internet connection offered) which are essential to service trading.

⁵ See <http://linked-usdl.org/>

The most notable effort able to represent and reason about business models, services, and value networks is the e³ family of ontologies which includes the e³service and e³value ontologies [3, 8]. This research has, however, not been much concerned with the computational and operational perspectives covering for instance the actual interaction with services. Likewise, the technical issues related to enabling a Web-scale deployment and adoption of these solutions were not core to this work. GoodRelations [9] (GR) on the contrary is a popular vocabulary for describing semantically products and offerings. Although GR originally aimed to support both services and products, it is mostly centred on products to the detriment of its coverage for modelling services, leaving aside for instance the coverage of modes of interaction, or the support for value chains.

USDL [4, 10] is, to date, the most comprehensive approach to supporting the description of services for automated processing. USDL consists of 9 modules modelled in eCore capturing services, interaction interfaces, pricing models, service level agreements, and related legal issues⁶. Despite its comprehensive support, this effort underestimated the need for such an all encompassing model to be widely open, highly flexible and extensible, and yet simple in nature [11]. On the one hand, the rather centralised and controlled nature of the approach led to an overly complex model hard to grasp and apply. On the other hand, eCore exhibited technical limitations towards its extensibility and its use as a lingua franca on the Web where Linked Data and light semantics are currently considered a more adequate technology.

3 Requirements Analysis

Informed by research carried out on services, including the related work covered earlier, we have elicited a number of requirements that Linked USDL and any other language or vocabulary with such an ambitious purpose should address. This includes notably coverage requirements, which we shall cover first. We also present additional criteria that we identified during the standardisation activities of USDL as potential issues and limitations for its Web-scale adoption [11].

3.1 Description Requirements

One of the essential difficulties when dealing with services beyond mere technical interfaces, is the fact that they are at the intersection of many diverse disciplines that range from technical aspects, to operational ones, socio-economic concerns, or even legal issues. Being able to move across each of these domains is essential to support the trading of services online. We detail the main dimensions next.

Functionality Services are business activities that normally take place through (possibly technology mediated) interactions between stakeholders, resulting in benefits to the actors involved. Fundamental to the notion of service is

⁶ See the full specification at http://www.internet-of-services.com/fileadmin/IOS/user_upload/pdf/USDL-3.0-M5-Archive.zip

therefore its functionality in terms of what it does, requires, and provides. Given the highly diverse nature of services this should cover the entire spectrum from fully automated provisioning (e.g., Spotify) to those essentially manual (e.g., car repair service). Depending on the stakeholder, the level of abstraction could vary from a detailed operational view (provider), to a high-level one for customers.

Agents and Networks Services delivery engages several stakeholders in (possibly ephemeral) ad-hoc business networks, e.g., banks often engage in partnerships with insurances to provide accounts with integrated travel insurance. The modelling of services should seamlessly support both the emergence and analysis of such networks in order to enable the dynamic co-creation of value through Web-wide service trading. Important aspects to be covered are thus the agents involved in a certain network and the role(s) they play.

Service Relationships Thanks to their intangible nature, services can be combined, repurposed, and adapted to better fulfil customer needs. Services are often related to other services and products. For instance, services can often be *enhanced* with others, or there can be *variations* over established types. Services are often *bundled*, i.e., aggregated and offered jointly in packages like broadband and TV services. And in the case of automated services, services may be *composed* according to specific data and control flow to achieve a complex objective out of simpler components.

Operational and Delivery The delivery of services is often subject to restrictions or conditions. These may range from geographical concerns (e.g., the insured individual should live in the UK), temporal availability, legal issues, variable pricing, and so on. From a service provider operational perspective, there may well be limitations due to the resources required, e.g., staffing, that need to be tracked as they determine the costs and the capacity for providing a service.

Consumption Services are most often accessed or “consumed” through interactions by means of designated *communication channels*. For example, making an insurance claim may require the customer to phone the insurance company, or fill up a form online. These communication channels may vary during the service delivery process (e.g., initially claim by phone and check the progress online), and there may exist restrictions on how interactions should take place. For instance a car repair service may require you to bring the car to the garage whereas in other cases the service may take care of sending a mechanic within some geographical boundaries.

3.2 Language Requirements

In addition to the aforementioned coverage requirements, research in the area has highlighted further requirements that the language should meet. First and foremost given the complexity of the domain and the fact that the aim is to maximise to the extent possible the level of automation that can be achieved during the life-cycle of services, the modelling of services needs to rely on a conceptual model with formal foundations that can enable automated processing [3,

10]. Nonetheless the language should be modular and extensible in order to be able to accommodate different domains and the many facets of services while minimising the complexity for users and tool developers.

Our subsequent work on standardisation highlighted that although necessary, these requirements did not appear to be sufficient for Web-scale adoption:

An Open Solution To support the engagement of any business entity across any domain the technological approach should be open. It should be open so as to allow anybody to engage and trade services online, as well as towards its evolution in order to cater for new requirements, accommodate new ways of doing business, or support new domains.

A Web-based solution A scenario like the one envisaged requires an approach that can support the engagement of millions of service providers and consumers in exposing, locating, interpreting, and contracting services. This necessarily calls for highly interoperable and scalable solutions in terms of data sharing, data processing, and communication protocols.

Promoting take up While providing an open solution is likely to have a positive impact on technology take up, adoption will largely be determined by the simplicity with which any business entity could adopt a solution based on these technologies and the compatibility with existing legacy systems.

4 Linked USDL Vocabulary

Driven by the aforementioned requirements and informed by the drawbacks exhibited by USDL we worked on Linked USDL focussing essentially on reducing the complexity underlying USDL and fostering its wider adoption through the use of Web-centric technologies that are more amenable to extension, modification, and automation at large scale.

4.1 Design Decisions

First, due to the success, scale, growth, and current adoption of the Web for world-wide telecommunication and electronic commerce we believe that any technology hoping to enable service trading online should necessarily embrace and build upon the Web principles and technologies [12]. Notably Linked USDL should also embrace principles like i) the establishment of global identifiers, e.g., by using URIs to identify services and providers; ii) the use of links to other resources on the Web to enrich a particular datum with reusable and externally provided information, e.g., pointing to complementary services; iii) the use of HTTP as a simple uniform protocol for supporting interactions; and iv) the decoupling between resources and their representation. Doing so brings a technology stack that has proven to support large scale, efficient, multi-party interactions, as well as it directly provides an integration point with open, standard technologies that are already widely used and supported.

Second, to enable effective interactions at the business level, we need to provide standards that go beyond data transportation and syntactic representation [1].

To this end, Linked USDL embraces the use of formal ontology representation languages to capture the semantics of services such that they are amenable to automated reasoning. Linked USDL goes one step forward in the adoption of Web technologies to embrace the emerging standard approach for data sharing online, namely Linked Data [13]. Adopting these principles enables Linked USDL to capture, share, and interlink data about services of highly heterogeneous nature and domains, in an open, scalable, and uniform manner. Linked Data principles promote and support reuse which in turn helps to reduce the data modelling overhead (e.g., by reusing conceptual models and existing data sets), and maximise the compatibility with existing tooling. Both aspects are major challenges earlier versions of USDL faced which this work aims to alleviate.

4.2 Design Methodology

Following common Knowledge Engineering best practices [14], we aimed at creating a modular solution based on well-designed, widely adopted vocabularies that did not introduce substantial ontological commitments away from the core topics of interest. Thus, considerable effort was devoted to identifying and evaluating reusable ontologies.

First, we identified the main topics to be covered given the original USDL specification and determined some core terms characterising each of these topics. Informed by the topics and terms identified, we carried out both a manual and semi-automated search to determine potentially relevant reusable ontologies. On the one hand, we performed a state of the art analysis to identify ontologies that were relevant for the modelling of services, see [11] and Section 2. On the other hand, we used Swoogle [15], Watson [16], LOD Stats [17], and the Linked Open Vocabularies (LOV)⁷ engines to search for ontologies covering the main terms identified. For each of the queries asked, we kept the top 10 results. The resulting list was eventually enriched with widely-used general purpose vocabularies such as Dublin Core (DC) and Simple Knowledge Organisation Scheme (SKOS).

Second, for each of the vocabularies identified, we used both LOD Stats and LOV to figure out the number of datasets using these terms, the number of instances of the main concepts of interest present in datasets on the Web, and the number of times the vocabulary is reused elsewhere. The search for reusable ontologies provided us pointers to existing vocabularies of potential interest together with indications regarding their use and popularity. Table 1 shows the results obtained for the vocabularies for which there was at least one instance found on the Web⁸. Indeed, the statistics should not be taken as an exact value of the overall use of these vocabularies (e.g., GR is used more frequently than what is reflected by this analysis), but rather as a relative indication. Indeed we also took into account the properties defined by these vocabularies which are in some cases (e.g., DC Terms) the main constructs reused.

The design of Linked USDL was driven by these statistics, and a manual assessment of the quality, coverage, and potential alignments of the vocabularies.

⁷ <http://lov.okfn.org>

⁸ These statistics were last retrieved in November 2013.

Table 1. Top Vocabularies per Topic.

Topic	Vocabulary	# Datasets		# Instances		LOV Reuse
		LOD	LOV	LOD	LOV	
Service	GR	6	45	146	0	6
	MSM	2	0	41,368	0	0
	OSLC	2	0	2	0	0
	COGS	N/A	5	N/A	0	3
Offering	GR	6	8	824	656	4
Location	vCard (v3 & v4)	5	0 + 2	3,684	3,686 + 3	0 + 2
	WGS84	11	1	3,204	1,7651	1
	AKT Signage	18	0	11,789	0	0
	DC Terms	1	9	39	39	6
	Schema.org	-	1	-	5	1
Business Entities	Schema.org	2	4	1,570,778	1,570,778	3
	FOAF	60	135	14,613	14,557	29
	GR	1	N/A	3,918	N/A	N/A
	W3C Org.	1,050	11	2	1,050	2
Time	W3C Time	9	N/A	236,433	N/A	N/A

4.3 Model

Informed by the aforementioned analysis, Linked USDL, which is publicly available together with further examples in GitHub⁹, builds upon a family of complementary networked vocabularies that provide good coverage of necessary aspects and are widely used on the Web for capturing their particular domains. In particular Linked USDL builds upon:

- DC Terms¹⁰ to cover general purpose metadata such as the creator of a certain description, its date of creation or modification, etc.
- SKOS providing low-cost support for capturing knowledge organisation systems (e.g., classifications and thesauri) in RDF.
- Time Ontology (Time)¹¹ for covering basic temporal relations. The ontology allows us to capture temporal relationships such as *before* and *during*.
- vCard vocabulary¹² a vCard 4 compatible vocabulary to support providing location and contact information for people and organisations.
- Minimal Service Model¹³ (MSM) [18] to provide coverage for automated service-based interactions including Remote Procedure Call solutions (e.g., WSDL services) and RESTful services.
- GR¹⁴ [9] to provide core coverage for services, business entities, offerings, and products.

The vocabulary has been modelled mostly using RDF/RDFS constructs and we have limited the inclusion of abstract foundational concepts, so as to attain a

⁹ <https://github.com/linked-usdl/usdl-core>

¹⁰ <http://purl.org/dc/terms/>

¹¹ <http://www.w3.org/TR/owl-time>

¹² <http://www.w3.org/TR/vcard-rdf/>

¹³ <http://iserve.kmi.open.ac.uk/ns/msm>

¹⁴ <http://purl.org/goodrelations/>

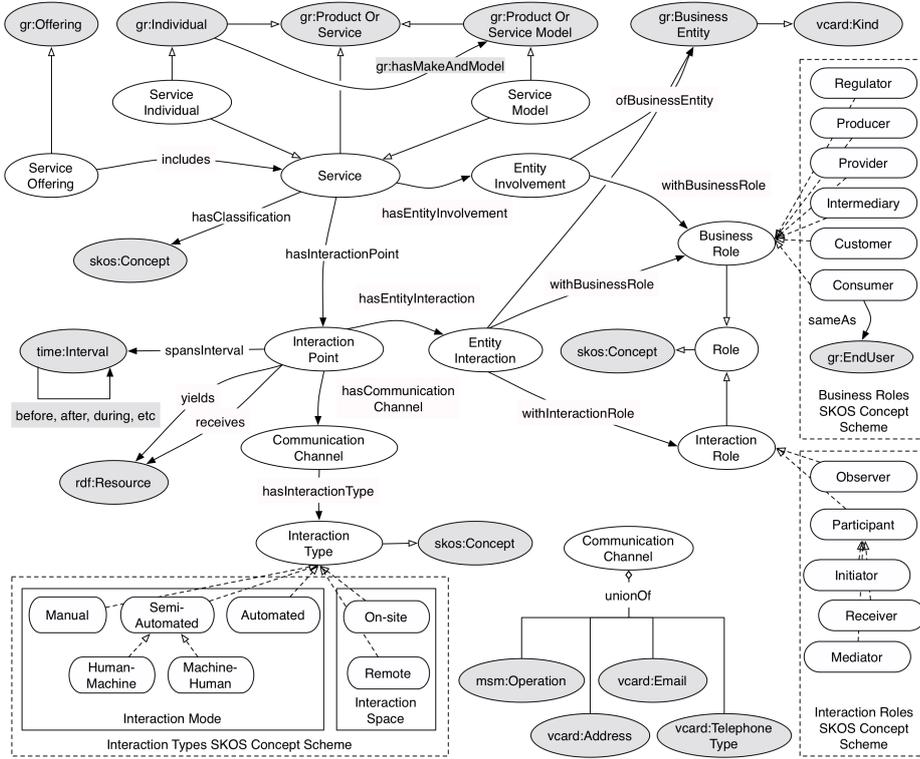


Fig. 1. Linked USDL Core.

model that is simple enough for its adoption online. The reader is referred to [19] for indications on how this model could be mapped to a foundational ontology.

As the core and initial module of a set of vocabularies for supporting service trading online Linked USDL Core, see Figure 1, aims to cover four essential aspects: offerings, services, the business entities involved in the delivery chain, and the actual interaction points allowing consumers to contract or trigger the benefits of contracted services.

Linked USDL extends GR which is nowadays the de-facto standard vocabulary for publishing semantic descriptions for products. It is worth noting that although services are accommodated within GR, their coverage is rather basic at this stage. Extending GR enables linking services and products descriptions which is particularly useful since many products are often sold in combination with a service, e.g., a repair or replace service. Additionally, it also ensures that an initial alignment with the increasingly popular vocabulary Schema.org is in place, for GR is already largely aligned to it.

The most important concepts provided by Linked USDL are:

Service is a refinement of *gr:ProductOrService* and subsumes all classes describing service types. Examples of subclasses of *Service* could be “internet provisioning service” and “insurance service”. Instances of *Service* may define i) prototypical services part of a portfolio, e.g., “BT unlimited broadband service”, as covered by *ServiceModel*, ii) one-of services custom tailored for a potential customer, or iii) actually contracted services, e.g., “your concrete life insurance provided by AXA”, as covered by *gr:ServiceIndividual*.

ServiceModel is a refinement of *gr:ProductOrServiceModel* which specifies common characteristics (e.g., download speed) of a family of services. *ServiceModel* thus defines families of *Services* sharing common characteristics, e.g., “*BT unlimited broadband services* share the characteristic of supporting unlimited download”. An actual service instance shares the properties of its service model. This is a feature that requires non-standard reasoning which specific implementations should take care of.

ServiceIndividual is a subclass of *gr:Individual* and *Service*. Instances of *ServiceIndividual* are actual services that are creating value to a network of business entities. For instance, “your concrete life insurance provided by AXA” is a *ServiceIndividual* which is providing value to yourself and AXA.

ServiceOffering is a subclass of *gr:Offering* and represents essentially offerings by a business entity *including* at least one *Service*. *ServiceOffering* may have limited validity over geographical regions or time.

EntityInvolvement is introduced in Linked USDL in order to enable capturing service value networks. In a nutshell, Entity Involvement allows capturing a ternary relationship expressing that a business entity, e.g., “AXA”, is involved in a service, e.g., “basic life insurance” playing a business role, e.g., “provider”. Linked USDL provides a reference SKOS taxonomy of basic business roles that covers the most typical ones encountered such as regulator and intermediary.

InteractionPoint link services to interactions that may be possible or required between the members of a service value network and the service during its life cycle. This allows answering questions such as “*what is the sequence of interactions I may expect if I want to make an insurance claim and what communication channels are available to that end?*”.

CommunicationChannel is the class of all different communication channels that business entities could use for communication. Linked USDL covers the most widely used channels by means of 2 vocabularies: vCard (e.g., email, phone), and MSM (e.g., Web services, and RESTful services). Communication channels are additionally characterised by their interaction type. Linked USDL provides 2 reference SKOS taxonomies covering the main modes (e.g., automated) and the interaction space (e.g., on-site).

EntityInteraction links interaction points to business entities or types (e.g., provider), and the role they play within the interaction (e.g., initiator). *EntityInteraction* allows expressing things like “*to make a claim, the consumer should first contact the insurance provider and provide the policy number*”.

Classifications Classifications or taxonomies of entities are most often used when describing services to capture, for instance, service types, business entity roles, e.g., “provider”, as well as interaction related issues, e.g., “manual vs automated”. We also expect that classifications will be needed in forthcoming modules addressing strategic issues or the internals of delivery chains.

This could be approached directly using subclassing which is directly supported by RDFS. However, the use of a hierarchy of classes establishes strict relationships which may not adequately match existing organisation schemes. For this reason, in Linked USDL we have accommodated the use of SKOS, which enables capturing classification schemes and taxonomies. Indeed, this mechanism does not prevent users from providing their own domain-specific categorisations through subsumption if they wish to. This approach thus enriches Linked USDL with a powerful, yet flexible and extensible means for creating categorisations.

The current version of Linked USDL includes three SKOS schemes with reference categorisations for *BusinessRoles*, *InteractionRoles*, and *InteractionTypes*, see Figure 1. These schemes have been, however, kept as separate modules so that different schemes can be used if necessary.

5 Evaluation

We have evaluated Linked USDL using three well-known and recommended techniques [20] including domain coverage, suitability for an application or task, and vocabulary adoption.

5.1 Coverage Evaluation

Ontologies are often evaluated by comparing them to a gold standard ontology [20]. In our case, we have done such an evaluation by comparing the resulting model to USDL, the most comprehensive model available for describing services. Doing so allows us to get a clear indication of the overall coverage of the domain, and to identify as well the main deviations from USDL.

A fundamental goal of this work is providing a conceptual model that would be easy to grasp, populate, process, and ultimately be adopted for Web-scale use. Thus, out of the 9 modules of USDL we have essentially deferred covering the following modules: Service¹⁵, Legal, Service Level, and Pricing. Nonetheless, for every module we have checked the coverage of the main concepts defined in order to get an indication of both module-specific and the overall coverage of Linked USDL. The results of this analysis are summarised in Table 2.

This analysis shows that thanks to integrating an reusing existing vocabularies we have managed to cover the vast majority of USDL, by providing a vocabulary consisting of 12 concepts and 3 complementary SKOS categorisations. In particular, from an original specification with 125 concepts we cover 74%, if we limit ourselves to the specific modules we targeted, and 60% overall, which

¹⁵ The Service module covers the internal details of a service which are often private.

Table 2. Evaluation of Linked USDL coverage of USDL (version M5).

USDL Module	Topic	Vocabulary	Comments	Classes	Covered	Ratio
Foundation	Time	Time	Advanced temporal reasoning provided	46	35	76%
	Contact Details	GR & vCard				
	Agents	GR & vCard				
	Conditions	±	Deferred to modules, e.g., Technical			
	Resources	X				
Technical	Interfaces	MSM	Higher automation through semantics	10	8	80%
	Protocols	HTTP & MSM	HTTP & SOAP/WSDL			
	Access Profile	X				
Interaction	Simple Protocols	Linked USDL		6	3	50%
	Complex Protocols	Linked USDL	Partial. Conditions at the operations level.			
Participants	Roles	Linked USDL	Business Roles SKOS	7	6	86%
	Target Consumers	X				
Functional	Parameters & Faults	MSM		4	2	50%
	Functions	GR	Basic coverage			
Approximate Coverage of Main Addressed Modules of USDL M5				73	54	74%
Service	Single Services	Linked USDL		11	5	45%
	Service Variants	Linked USDL & GR	Partial with Service Model			
	Service Types	Linked USDL	Interaction Types SKOS			
	Composite Services	X	Offering bundles supported			
Pricing	Basic Pricing	GR	Payment types, taxes, cost	19	7	37%
	Variable Pricing	X				
Service Level	Metrics & Conditions	GR & MSM		9	4	44%
	Guarantees	X				
Legal	Basic Legal	GR	License, Validity, etc	13	5	38%
	Rights, Requirements	X				
Approximate Total Coverage of USDL M5				125	75	60%

shall contribute towards reducing the overhead related to understanding and adopting Linked USDL. It is worth noting that out of the concepts not explicitly covered several are sometimes redundant (e.g., Condition is subclassed in many modules), or were seldom properly understood and used (e.g., Functions, Phases of interactions, Service Level Agreements).

5.2 Suitability for Tasks and Applications

Given that Linked USDL does not cover all concepts present in USDL it is worth assessing the impact of such decisions. Table 2 shows the main aspects and their current coverage. In qualitative terms, the decisions adopted are such that Linked USDL does not currently provide support for capturing how providers deliver services in terms of resources needed, complex internal workflows, or strategic decisions (e.g., targeted markets). The reason for this is two-fold. First, such aspects are often not automated and when they are, providers already have mechanisms in place to this end. Second, these are private concerns that are orthogonal to the trading of services. Similarly, Linked USDL does not currently include support for conceptualising complex agreements including legal requirements and guarantees as these were barely used or understood by users. Finally, we have opted for a simple mechanism for capturing prices and have

deferred to a separate module the modelling of more complex dynamic pricing that are less often used and usually remain private to the provider.

Despite these changes, Linked USDL provides advanced support for modelling, comparing, discovering, and trading services and service bundles. It provides means for tracking and reasoning about the involvement of entities within delivery chains which informs the trading and comparison of services as well as it enables the tracing and analysis of service value networks. It provides advanced support for automating the interactions between actors during the life-cycle of services. Additionally it includes support for capturing service offerings, for combining services and products (e.g., a car often comes with a warranty), and for applying temporal reasoning, which were not previously available. Finally, and most importantly, these activities can be achieved with a greater level of automation benefitting from automated reasoning and they can be performed on a Web-scale across Web-sites and service providers thanks to capturing and sharing the semantics of services as Linked Data.

Empirically, the suitability of the language for supporting the automation of key tasks has been evaluated by two main means. On the one hand, we have reused and developed tools that provide support for these tasks, and, on the other hand, we are continuously applying Linked USDL in a number of domains. In terms of reuse, thanks to the adoption of existing Linked Data vocabularies, Linked USDL benefits from general purpose tooling, e.g., SPARQL engines and RDF stores, but also from vocabulary-specific solutions. This notably concerns existing advanced machinery for discovering, composing, and invoking technical services (i.e., RESTful and WSDL services) described in terms of MSM [18].

Additionally, general purpose infrastructure has been developed specifically for Linked USDL. A Web-based Linked USDL editor is currently available to help providers to easily generate Linked USDL descriptions¹⁶. There is also an advanced multi-party dynamic and open service marketplace¹⁷ developed in the context of the FI-WARE project¹⁸, able to gather, combine, and exploit rich service descriptions from distributed providers to help match offer and demand. Notably the marketplace supports consumers in searching for service offerings, comparing them, and contracting them.

Finally, from the perspective of its suitability for supporting service trading across domains, Linked USDL is currently being applied in a variety of domains. For instance, in the field of Software as a Service we have explored the use of Linked USDL in conjunction with TOSCA[21]. Linked USDL was used to formalise, structure, and simplify the discovery and selection of services of the Web-based customer relationship management (CRM) platform SugarCRM, while TOSCA supported the automated deployment and management of the services. Additionally this work helped us evaluate the extensibility of Linked USDL by integrating it with complementary third party specifications such as TOSCA. In the FI-WARE project Linked USDL is used to support a service infrastructure

¹⁶ See <https://github.com/linkedin-usdl> for existing tooling and model extensions.

¹⁷ <http://store.testbed.fi-ware.org/>

¹⁸ <http://www.fi-ware.eu/>

supporting service ecosystems in the cloud covering both the technical and business perspectives. The FINEST¹⁹ project aims to support the transport and logistics (T&L) ecosystem, in which many service providers collaborate in order to transport goods over what is referred to as a “chain of legs”. Therein Linked USDL is being exploited in the planning of chains of legs to support searching and matching transport service offerings in a transparent, distributed, and multi-party manner.

Across the diverse domains where Linked USDL is being applied (see list of projects next), it has proven to be a valuable resource as a means to provide shared and globally accessible service descriptions integrating both technical and business aspects. The genericity, modularity, and extensibility of the approach has enabled extending the vocabulary with dedicated domain-specific vocabularies in the areas of SaaS and T&L, while generic software infrastructure was easily reused across domains.

5.3 Vocabulary Adoption and Use

When evaluating ontologies and vocabularies, one aspect that is often taken into account is their adoption and use. This evaluation may be carried over the ontology itself and/or over the different ontologies that are imported. The former gives an indication of the acceptance and adoption of the ontology in its entirety whereas the latter provides a more granular assessment over the reused ontologies. In this section we mainly address the latter but also provide preliminary indications of the overall adoption of Linked USDL.

The methodology that was followed, see Section 4.2, was centred on the reuse of widely adopted vocabularies. Table 1 presented earlier shows the main vocabularies that were identified through search engines, together with core indicators of their use on the Web. These figures highlight that Linked USDL is based on vocabularies that are the most used in their respective domains of interest. Only two exceptions exist, AKT Signage which was not adopted for it was not dereferenceable, and Schema.org which is indirectly aligned via GR. This approach in turn reduces the potential overhead one would incur when using Linked USDL: frequently reused vocabularies are likely to have greater acceptance and support by people and existing systems.

Additionally, the availability of datasets with instances in terms of the vocabularies reused guarantees that new descriptions could reuse and link to existing resources, e.g., allowing the reuse of descriptions of companies. Doing so provides clear benefits from the perspective of data acquisition which was one of the main concerns Linked USDL was trying to address. Additionally, by linking to existing instances the data provided is enriched which may in turn enable further advanced processing as well as it may increase the discoverability of services.

Providing a substantial account of the adoption of Linked USDL would require a reasonable wait from its first release, which coincides with this publication. Nonetheless, Linked USDL is currently already in use within more than 10

¹⁹ <http://www.finetest-ppp.eu/>

research projects, namely FI-WARE, FINEST, Value4Cloud, Deutsche Digitale Bibliothek, MSEE, FIspace, FITMAN, FI-CONTENT, ENVIROFI, OUTSMART, SMARTAGRIFOOD, IoT-A, Broker@Cloud, and GEYSERS. These projects are using Linked USDL as the core vocabulary for describing services, contributing to validate the suitability, genericity, and extensibility of Linked USDL for different domains. This also highlights that despite its youth, Linked USDL is already witnessing a promising adoption.

6 Conclusion

Despite the importance of services in developed economies, the widespread adoption of world-wide electronic commerce over the Web, most service trading is still essentially carried out via traditional and often manual communication means. A fundamental reason for this is the difficulty for capturing the abundant information and knowledge governing services and their related transactions in a way amenable to computer automation. Out of the wealth of work around services, USDL is the most comprehensive solution proposed thus far for enabling (semi)automated service trading. Yet, work on its standardisation highlighted a number of limitations for Web-scale service trading.

We have presented Linked USDL, the next evolution of USDL centred on fostering its wider adoption and better automation support through the (re)use of Linked Data. Linked USDL has been developed following a methodology centred on maximising the reuse of existing vocabularies and datasets and minimising the complexity. The resulting vocabulary has been evaluated along in terms of domain coverage, suitability for purpose, and vocabulary adoption.

Despite the good evaluation results obtained, Linked USDL is to be regarded as one step towards enabling Web-scale service trading, albeit a fundamental one. Further work is required for covering aspects such as complex dynamic pricing models and agreements which are common in certain domains such as Cloud services. Additionally, from the tooling perspective, developing advanced mechanisms able to support steps such as the negotiation between service providers and consumers, or the bundling of services would also be necessary. We expect in this last regard to take inspiration and adapt solutions developed for the e³ family of ontologies.

Acknowledgment

This work was partially funded by DFG under project agreements SFB 912/1 2011, and by the COMPOSE (FP7-ICT-317862) and FI-WARE (FI-PPP-285248) EU projects. We also thank all the members of the W3C USDL Incubator Group.

References

1. Chesbrough, H., Spohrer, J.: A research manifesto for services science. *Communications of the ACM* **49**(7) (July 2006) 35

2. Papazoglou, M.P., Traverso, P., Dustdar, S., Leymann, F.: Service-Oriented Computing: State of the Art and Research Challenges. *Computer* **40**(11) (2007) 38–45
3. Akkermans, H., Baida, Z., Gordijn, J., Peña, N., Altuna, A., Laresgoiti, I.: Value Webs: Ontology-Based Bundling of Real-World Services. *IEEE Intelligent Systems* **19**(4) (August 2004) 57–66
4. Cardoso, J., Barros, A., May, N., Kylau, U.: Towards a Unified Service Description Language for the Internet of Services: Requirements and First Developments. *Services Computing (SCC)*, 2010 IEEE International Conference on (2010) 602–609
5. Cardoso, J., Sheth, A.: Semantic e-workflow composition. *Journal of Intelligent Information Systems (JIIS)* **21**(3) (2003) 191–225
6. Pedrinaci, C., Domingue, J., Sheth, A.: Semantic Web Services. In: *Handbook on Semantic Web Technologies. Volume Semantic Web Applications*. Springer (2010)
7. Oppenheim, D.V., Varshney, L.R., Chee, Y.M.: Work as a service. In: *ICSOC'11: Proceedings of the 9th international conference on Service-Oriented Computing*, Springer-Verlag (December 2011)
8. Gordijn, J., Yu, E., van der Raadt, B.: e-service design using i* and e3value modeling. *IEEE Software* **23** (2006) 26–33
9. Hepp, M.: GoodRelations: An Ontology for Describing Products and Services Offers on the Web. In: *Knowledge Engineering: Practice and Patterns*. Springer (2008) 329–346
10. Oberle, D., Barros, A., Kylau, U., Heinzl, S.: A unified description language for human to automated services. *Information Systems* (2012)
11. Kadner, K., Oberle, D., Schaeffler, M., Horch, A., Kintz, M., Barton, L., Leidig, T., Pedrinaci, C., Domingue, J., Romanelli, M., Trapero, R., Kutsikos, K.: Unified Service Description Language XG Final Report. Technical report (October 2011)
12. Jacobs, I., Walsh, N.: *Architecture of the World Wide Web, Volume One*. W3C Recommendation (December 2004)
13. Bizer, C., Heath, T., Berners-Lee, T.: Linked Data - The Story So Far. *International Journal on Semantic Web and Information Systems (IJSWIS)* (2009)
14. Suárez-Figueroa, M.C., Gómez-Pérez, A., Motta, E., Gangemi, A., eds.: *Ontology Engineering in a Networked World*. Springer (2011)
15. Ding, L., Finin, T., Joshi, A., Pan, R., Cost, R.S., Peng, Y., Reddivari, P., Doshi, V.C., Sachs, J.: Swoogle: A Search and Metadata Engine for the Semantic Web. In: *CIKM '04: Thirteenth ACM international conference on Information and Knowledge Management*. (November 2004)
16. d'Aquin, M., Motta, E.: Watson, more than a Semantic Web search engine. *Semantic Web* **2**(1) (2011) 55–63
17. Auer, S., Demter, J., Martin, M., Lehmann, J.: LODStats — an extensible framework for high-performance dataset analytics. In: *EKAW'12: Proc. of the 18th international conference on Knowledge Engineering and Knowledge Management*, Springer (2012)
18. Pedrinaci, C., Domingue, J.: Toward the Next Wave of Services: Linked Services for the Web of Data. *Journal of Universal Computer Science* **16**(13) (2010) 1694–1719
19. Ferrario, R., Guarino, N., Janiesch, C., Kiemes, T., Oberle, D., Probst, F.: Towards an ontological foundation of services science: The general service model. *Wirtschaftsinformatik* (February 2011) 16–18
20. Sabou, M., Fernandez, M.: Ontology (network) evaluation. In Suárez-Figueroa, M.C., Gómez-Pérez, A., Motta, E., Gangemi, A., eds.: *Ontology Engineering in a Networked World*. Springer (2012) 193–212
21. Cardoso, J., Binz, T., Breitenbücher, U., Kopp, O., Leymann, F.: Cloud Computing Automation: Integrating USDL and TOSCA. In: *25th Conf. on Advanced Inf. Systems Engineering (CAiSE 2013)*. Volume 7908 of LNCS., Springer (2013) 1–16