

A Proof Markup Language for Semantic Web Services

Paulo Pinheiro da Silva Deborah L. McGuinness
Richard Fikes

Knowledge Systems Laboratory, Stanford University
Stanford, CA 94305, USA.
e-mail: {pp,dlm,fikes}@ksl.stanford.edu

Abstract

Web services propose that they provide the means for remote inter-operable access of components and software systems. However, successful inter-operation between components that do anything more than the simplest information retrieval is dependent upon those components having a shared understanding of the results that have passed between them. In this paper, we address the issue of understanding and trusting results generated by web services. We introduce a proof markup language (PML) that provides an interlingua for capturing the information agents need to understand results and to justify why they should believe the results. We also introduce our Inference Web infrastructure that uses PML as the foundation for providing explanations of web services to end users. We additionally show how PML is critical for and provides the foundation for hybrid reasoning. Our contributions in this paper focus on technological foundations for capturing formal representations of term meaning and justification descriptions thereby facilitating trust and reuse of answers from web agents.

1 Introduction

Web services are expected to make use of automated reasoners and retrieval engines to perform their services. In order for Semantic Web services to explain their results, they need to be able to generate justifications of their results in an exchangeable, combinable format. We refer to a sequential trace of information manipulations used to generate an answer as a proof of the answer. These proofs (possibly provided by multiple question answering systems) may be combined to form proofs of compound results. Proofs

may be abstracted, summarized, or transformed into more abbreviated complex aspects or more understandable justifications for the answers. For the purposes of this paper, we will view an explanation as any transformation (abstraction or summaries included) of a proof that makes it more understandable to users.

In order to provide the foundation for proof manipulations - presentations, abstractions, summaries, transformations, and explanations - we introduce a proof markup language as a proof representation interlingua. This article introduces the Proof Markup Language (PML) that can be used as follows:

- To represent different kinds of proofs ranging from formal natural deduction derivations to proof traces typically produced by highly optimized theorem provers. These representations are aimed at services that leverage inference.
- To represent proofs resulting from services aimed at information retrieval or entity extraction. These proofs may provide information concerning sources and extractors use;
- To represent different kinds of explanations varying from summaries of assertions or knowledge sources used to derive an answer to a more elaborate abstraction of a derivation path;
- To support the exchange of proof information among automated reasoners in hybrid reasoning systems.

The rest of the article is organized as follows. Section 2 describes the Inference Web that is a Semantic Web infrastructure for explanations. Section 3 introduces the PML specification. Section 4 demonstrates practical uses of the PML format for explanations and hybrid reasoning. Section 5 presents related work. Section 6 concludes the article and describes future work.

2 Inference Web

Inference Web [9] (IW) is a framework for explaining answers produced from the Semantic Web Services. Inference Web provides access to proofs and explanations published anywhere on the Web. It provides a framework and tools for building, maintaining, presenting, exchanging, combining, annotating, filtering, comparing, and rendering proofs and proof fragments. Inference Web is composed of the following portions:

- *Specifications*: Question answering components may generate answers and justifications for their answers using the PML format described in this article. PML provides a proof interlingua representation. Inference Web includes a specification of PML in DAML+OIL [2] and OWL [11].
- *Data*: PML documents published on the Web become a portion of the Inference Web data used when browsing and summarization tools are presenting explanations, summaries, and other viewing options to users. Inference Web also processes the PML documents to identify when portions of them may be combined to form more complex conclusions. IW also uses rewrite rules to transform the proofs into more understandable explanations. The IW data includes the PML proofs and explanations along with a registry of information used for proof presentation. The registry, called IWBBase, is a distributed repository of meta-data including sources, inference engines and inference rules. This information is used to provide follow-up information concerning the elements presented in proofs and explanations.
- *Tools and Services*: Inference Web includes a tool suite including a browser for displaying proofs and explanations, an explainer for abstracting PML proofs into PML explanations, a registrar for submitting and maintaining the evolving IWBBase entries, a proof generation service for facilitating the creation of PML proofs by inference engines, and the IWBBase registry for storing information used in proofs and explanations.

The IWBBase [10] (formerly known as the IW Registry) is an interconnection of distributed repositories of meta-information relevant to proofs and explanations, including knowledge provenance information [13]. Every entry in these repositories is an instance of an IWBBase concept as described in Section 3.3. For example, *Knowledge Source* is an IWBBase concept that is useful for entries such as ontologies, knowledge bases, thesauri, etc. The knowledge source entry for an ontology describes stores of assertions about the ontology such as its original creator(s), date of creation, data of last update, version, URL (for browsing), description in English, etc. IWBBase's provenance information is expanding on an as-needed basis driven by application demands.

Every entry has a URI and is stored both as a file written in OWL and as a set of tuples in a database. IWBBase files are mainly used by PML proofs to annotate their content as described throughout this article.

3 PML Specification

PML classes are OWL classes (thus they are subclasses of `owl:Class`). They are used to build OWL documents representing both proofs and proof provenance information. Thus, PML concepts can be considered to be either proof level concepts or provenance level concepts. Primitive types mentioned in the article are from the XML schema specification¹.

3.1 Proof Level Concepts

`NodeSet`², `InferenceStep`, and `Expression` are the main constructs of proofs and explanations.

A `NodeSet` represents a step in a proof whose conclusion is justified by any of a set of inference steps associated with the `NodeSet`. PML adopts the term “node set” since each instance of `NodeSet` can be viewed as a set of nodes gathered from one or more proof trees having the same conclusion.

- The `URI`³ of a node set is the unique identifier of the node set. Every node set has one well-formed URI.
- The `Conclusion` of a node set represents the expression concluded by the proof step. Every node set has one conclusion, and a conclusion of a node set is of type `Expression`.
- The `ExpressionLanguage` of a node set is the language in which the conclusion is represented. Every node set has one expression language, and that expression language is of type `ExpressionLanguage`.
- Each `InferenceStep` of a node set represents an application of an inference rule that justifies the node set’s conclusion. Every node set has at least one inference step, and each inference step of a node set is of type `InferenceStep`.

An `InferenceStep` represents a justification for the conclusion of a node set. Inference steps are anonymous OWL classes defined within node sets. For this reason, it is assumed that applications handling PML proofs are able to identify the node set of a inference step. Also for this reason, inference steps have no URIs.

¹<http://www.w3.org/TR/xmlschema-2/>

²PML concept names are typed in **sans serif** style and PML attribute names are typed in **courier** style.

³<http://www.ietf.org/rfc/rfc2396.txt>

- The **Rule** of an inference step is the rule applied to produce the conclusion. Every inference step has one rule, and that rule is of type `InferenceRule` (see Section 3.3.3). Rules are in general registered in the `IWBase` by engine developers. However, PML specifies three special instances of rules: *Assumption*, *DirectAssertion*, and *UnregisteredRule*. When specified in an inference step, the *Assumption* rule says that the conclusion in the node set is an explicit assumption.
- The **Antecedents** of an inference step is a sequence of node sets each of whose conclusions is a *premise* of the application of the inference step's rule. The sequence can contain any number of node sets including none. The fact that the premises are ordered may be relevant for some rules such as *ordered resolution* [14] that uses the order to match premises with the schemas of the associated rule. For other rules such as modus ponens, the order of the premises is irrelevant. In this case, **Antecedents** can be viewed as a set of premises.
- Each **Binding** of an inference step is a mapping from a variable to a term specifying the substitutions performed on the premises before the application of the step's rule. An inference step can have any number of bindings including none, and each binding is of type `VariableBinding`. For instance, substitutions may be required to unify terms in premises in order to perform resolution.
- Each **DischargedAssumption** of an inference step is an expression that is discharged as an assumption by application of the step's rule. An inference step can have any number of discharged assumptions including nodes, and each discharged assumption is of type `Expression`. This property supports the application of rules requiring the discharging of assumptions such as natural deduction's *implication introduction*. An assumption that is discharged at an inference step can be used as an assumption in the proof of an antecedent of the inference step without making the proof be conditional on that assumption.
- Each **Source** of an inference step refers to the an entity representing where original statements from which the conclusion was obtained. An inference step can have any number of sources including none, and each source is of type `Source` as described in Section 3.3.1. An inference step's source supports the justification of the node set conclusion when the step's rule is a *DirectAssertion*.
- The **Engine** of an inference step represents the inference engine that

produced the inference step. Each inference step has one `Engine`, which is of type `InferenceEngine`.

- The `TimeStamp` of an inference step is the date when the inference step was produced. Time stamp is of the primitive type `dateTime`. Every inference step has one time stamp.

An inference step is said to be well-formed if:

1. Its node set conclusion is an instance of the conclusion schema specified by its rule;
2. The expressions resulting from applying its bindings to its premise schemas are instances of its rule's premise schemas;
3. It has the same number of premises as its rule's premise schemas; and
4. If it is an application of the *DirectAssertion* rule, then it has at least one source, else it has no sources.

Further proof verification may be performed by checking side conditions on declarative rules and running verification methods on method rules. However, a discussion of this level of verification is beyond the scope of this article.

PML node set schemas and PML inference step schemas used later in the article are defined as follows. A PML **node set schema** is a PML node set which has a conclusion that is either a sentence schema⁴ or a sentence; which has a set of variable bindings that map free variables in the conclusion to constants; which has zero or more inference steps; and whose inference steps are either inference steps or **inference step schemas**. An inference step schema is an inference set of a node set schema whose antecedents are node set schemas.

An `Expression` is a PML representation of well-formed logical expressions written in accordance with a given `ExpressionLanguage`.

3.2 Proofs

A PML node set represents a directed acyclic graph of proofs of its conclusion. In this section we describe the proofs that are represented by a PML node set.

⁴A sentence schema is a sentence optionally containing free variables. An instance of a sentence schema S is a sentence that is S with each free variable replaced by a constant.

We begin by defining a *proof* as a sequence of “proof steps”, where each proof step consists of a conclusion, a justification for that conclusion, and a set of assumptions discharged by the step. “A proof of C ” is defined to be a proof whose last step has conclusion C . A proof of C is conditional on an assumption A if and only if there is a step in the proof that has A as its conclusion and “assumption” as its justification, and A is not discharged by a later step in the proof. An unconditional proof of C is a proof of C that is not conditional on any assumptions. (Note that assumptions can be made in an unconditional proof, but each such assumption must be discharged by a later step in the proof.) Finally, proof $P1$ is said to be *subproof* of $P2$ if and only if the sequence of proof steps that is $P1$ is a subsequence of the proof steps that is $P2$.

A PML node set N having conclusion C represents a set of proofs of C defined as follows. If N is a node set, we will say that P is a “proof from N ” if and only if:

1. The conclusion of the last step of P is the conclusion of N ;
2. The justification of the last step of P is one of N ’s inference steps S ;
and
3. For each antecedent A_i of S , exactly one proof from A_i is a subproof of P .

If N is a node set having conclusion C , then a proof from N is a proof of C .

3.3 Provenance Level Concepts

Inference Web stores provenance information about proofs and explanations in the IWBase. This section describes the concepts supported by IWBase and that are part of the PML specification.

3.3.1 Provenance Element

`ProvenanceElement` represents a information unit describing the origin of some PML proof level concept introduced in Section 3.1. `ProvenanceElement` is a superclass of the PML concepts at the provenance level and the `ProvenanceElement` attributes are described as follows:

- The URI of a provenance element is the unique identifier of the provenance element. Every provenance element has one well-formed URI, which is an instance of the primitive type `anyURI`.

- The **URL** of a provenance element describes a URL used to browser an element’s web document. For instance, if the provenance element is an organization named the New York Times, then the `http://www.nytimes.com` URL can be used to access a web document about the organization. I In this case is the URL points to the organization’s web site. A provenance element can have zero or one URL.
- The **Name** of a provenance element described a short name (or “nick-name”) for the element within the IWBBase. Every provenance element has one name, and the name is an instance of the primitive type `string`.
- The **Submitter** of a provenance element represents the team of people responsible for the registration of the provenance element in IWBBase. Every provenance element has one submitter, and the submitter is an instance of the type `Team`, which is a subclass of `Source` as described in Section 3.3.2.
- The **DateTimeInitialSubmission** of a provenance element is the date when the provenance element was first registered in IWBBase. Every provenance element has one **DateTimeInitialSubmission**, and that is and instance of the primitive type `dateTime`.
- The **DateTimeLastSubmission** of a provenance element is the last date when the provenance element was registered in IWBBase. Every provenance element has one **DateTimeLastSubmission**, and that is an instance of the primitive type `dateTime`.
- The **EnglishDescription** of a provenance element is a description in English of the provenance element. A provenance element can have zero or one description in English, and **EnglishDescription** is an instance of the primitive type `string`. The description in English is intended to be used by tools to present provenance elements to human agents. For example, the description in English of the *Modus Ponens* inference rule may be a better presentation and more informative for most users browsing a PML document than the presentation of the formal specification of the rule.

3.3.2 The Source Concept

A `Source` is a `ProvenanceElement` representing an entity which is the source of the original data. Current types of Inference Web sources include: `InferenceEngine`, `KnowledgeSource`, `Organization`, `Person`, `Publication`, `RepresentationLanguage`, `ExpressionLanguage` and `Team`. A full specification of sources

is available at <http://www.ksl.stanford.edu/software/IS/spec>. We provide a description of two source types here.

- An `InferenceEngine` represents an engine that is able to produce a justification for a given conclusion. Note that the use of the term “inference engines” in this article is not limited to engines with reasoning capabilities. For example, search engines retrieving information may serve as an inference engine and provide a justification of their answer by a direct assertion inference step. Similarly extraction modules may be viewed as inference engines.
- A `ExpressionLanguage` represents a language used to write conclusions of node sets. Any language can be registered as a representation language in `IWBase`. Expression languages are typically able to represent logical sentences, although PML does not place restrictions on the expression language.

3.3.3 The Rule Concept

An `InferenceRule` represents a `ProvenanceElement` specialization describing rules applied to premises deriving node set conclusions. An `InferenceRule` can be either a `PrimitiveRule` or a `DerivedRule`.

A `PrimitiveRule` is a type of an `InferenceRule` that is implemented by one or more inference engines. A given rule R_1 may not be called primitive until it become associated with one or more inference engines. Thus, assuming that R_1 is implemented by an inference engine E_1 , the inference engine may declare R_1 to be a primitive rule. The notion that R_1 is a primitive rule for one specific engine is relevant since R_1 may also be derived from R_2 , which is a primitive rule for another engine E_2 . In this case, R_1 may be registered once as a primitive rule and zero or more times as a derived rule, depending on how many combinations of rules are used to derive R_1 . For example a natural deduction reasoner E_1 may define *modus ponens* as a primitive rule and another reasoner E_2 may register Robinson’s *resolution* rule as a primitive rule. The E_2 reasoner may be able to derive a *modus ponens* rule using its primitive resolution rule. Thus, any rule R_1 can be registered multiple times, once as a primitive rule and multiple other times as derived rules depending on the number of different combinations of rules used to derive R_2 .

- The `ConclusionSchema` of a primitive rule represents the conclusion of the primitive rule and it can be either a sentence schema or a sentence.

Every primitive rule has one conclusion schema, and that conclusion schema is an instance of the type `Expression`.

- Each `PremiseSchema` of a primitive rule represents a premise of the primitive rule and it can be either a sentence schema or a sentence. A primitive rule have zero or more premise schemas, and each of them is an instance of the type `Expression`.
- Each `SideCondition` of a primitive rule represents a condition specified in terms of the attributes of the conclusion’s node set, the conclusion’s inference step applying the primitive rule; and premises’ node sets.
- The `ExpressionLanguage` of a primitive rule is the language in which the conclusion schema, premise schemas and side conditions of the primitive rule are represented. Every node set has one expression language, and that expression language is of type `ExpressionLanguage`.

A `DeclarativeRule` is a `PrimitiveRule` and is well-formed if and only if all side conditions for the primitive rule hold.

A `MethodRule` is a `PrimitiveRule` whose side conditions cannot be completely specified in terms of the attributes of the conclusion’s node set, the conclusion’s inference step applying the primitive rule; and premises’ node sets. Inference rules that are “procedural attachments” are examples of method rules. The current work on PML does not involve checking if method rules are well-formed although future plans include checking support.

A `DerivedRule` is an `InferenceRule` specified from a PML node set schema with the restriction that each node set schema must have one and only one inference step. The `Specification` of a derived rule represents a proof (as defined in Section 3.2) from a given PML node set schema since each PML node set must have one and only one inference step. Moreover, the derived rule’s proof is a **proof schema** since the PML node sets of the proof are PML node set schemas.

4 Using PML Proofs

4.1 Support for Hybrid Reasoning

Experience with automated reasoners has made clear that in order to effectively determine answers to complex real-world questions, general-purpose reasoners need to be augmented with special-purpose reasoners that embody

both domain-specific and task-specific expertise. That is, effective deductive answer determination requires *hybrid reasoning*.

We have developed an object-oriented modular architecture for hybrid reasoning (called the JTP architecture), a library of general-purpose reasoning system components (called the JTP library) that supports rapid development of reasoners and reasoning systems using the JTP architecture, and a multi-use reasoning system (called the JTP system) employing the JTP architecture and library [3]. The JTP architecture and library is intended to enable the rapid building, specializing, and extending of hybrid reasoning systems. Each reasoner in a JTP hybrid reasoning system can embody special-purpose algorithms that reason more efficiently about particular commonly-occurring kinds of information. In addition, each reasoner can store and maintain some of the system's knowledge, using its own specialized representations that support faster inference about the particular kinds of information for which it is specialized.

Proofs represented in PML play a central role in the JTP architecture in that they are used to represent both queries and reasoning results as they are sent to and received from reasoners during the hybrid reasoning process. In this section we describe JTP's use of PML proofs in hybrid reasoning.

4.1.1 JTP System Architecture

The JTP architecture assumes that there is a single initially empty knowledge base (KB) with respect to which all processing is done. A KB is considered to be a representation of a logical theory and to contain a set S of symbolic logic sentences and a set of justifications for each sentence in S . The architecture supports commands for loading a KB, adding an axiom to a loaded KB, removing an axiom from a loaded KB, and asking what (partial or full) instances of a sentence schema are entailed by a loaded KB.

The primary work of the system is assumed to be performed by modules called **reasoners**. There are "telling" reasoners that are invoked when a sentence is being added to the KB and "asking" reasoners that are invoked when the KB is being queried. Reasoners produce **reasoning steps**, each of which is a partial or completed portable proof. The reasoning steps produced by telling reasoners are completed proofs of additional sentences that are inferred from the reasoner's input. The reasoning steps produced by asking reasoners are partial or completed proofs of candidate answers to a query.

Since the set of answers to a query may be of unpredictable size and may require an unpredictable amount of time to derive, the output of a reasoner

is an **enumerator** that can be **pulsed** to obtain the next reasoning step produced by the reasoner. Enumerators enable a reasoner to provide output reasoning steps as they are derived and for additional derivations to be attempted on an as needed basis.

A reasoning system using the JTP architecture needs some means of determining to which of its arbitrary number of reasoners to route its inputs. That capability is provided by reasoners in the system that act as “**dispatchers**” of an input to other reasoners that the dispatcher determines may be able to process the input. Each dispatcher has a set of child reasoners associated with it and serves as a transparent proxy for those child reasoners.

4.1.2 Reasoning Steps

Reasoners produce enumerations of reasoning steps, and take reasoning steps as input. A reasoning step is a PML node set schema that represents a partial or completed proof of a symbolic logic sentence.

A reasoning step that is a node set schema N having conclusion C , variable bindings B , and no inference steps specifies a query to find proofs of instances of the sentence schema C/B (i.e., the sentence schema produced by applying the bindings B to C). A reasoner given such a reasoning step as input can produce either partial or complete proofs of instances C/B . Each partial or complete proof to be returned by the reasoner can be represented by adding inference steps and variable bindings to a copy of the input reasoning step. Each node set in a reasoning step produced by a reasoner that does not have an inference step is an unproven subgoal for which a proof is needed in order to complete the proof.

Thus, reasoning steps are used to represent both queries and reasoning results as they are sent to and received from reasoners during the hybrid reasoning process.

4.1.3 The Tell and Ask Commands

The *Tell Command* takes as input a sentence S and adds it to the KB. The command processor does that by forming a reasoning step R representing a proof of S justified as a direct assertion, and then calling a telling reasoner with P as input.

A telling reasoner takes as input a reasoning step that is a proof. The proof may represent either a sentence that is being told to the system (justified as a direct assertion), or a result of a previous inference, justified by an inference rule, that the reasoner is to build upon. The reasoner may assert

the sentence to one or more knowledge stores, produce additional inferences in the form of new proofs, or signal that a contradiction has been found. The output of a telling reasoner is an enumerator whose output when pulsed is a proof representing the result of a new inference.

The *Ask Command* takes as input a sentence schema S , and produces as output an enumerator whose output when pulsed is an unconditional proof of an instance of S . The command processor produces its output by forming a reasoning step having conclusion S , and then calling an asking reasoner with that reasoning step as input.

An asking reasoner accepts as input a reasoning step R having conclusion C , and no inference step. R represents a query whose answers are instances of C . The reasoner attempts to produce reasoning steps having the same conclusion as R and a variable binding set that is a superset of R 's variable binding set. The reasoner's output is an enumerator whose output when pulsed is such a reasoning step. If no reasoning steps can be produced, then the enumerator is empty.

For example, consider a query to find bindings for v_1 and v_2 such that " $(P v_1 v_2)$ " is true. Assume the knowledge base contains the axioms " $(\Rightarrow (\text{and } (Q x y) (R y z)) (P x z))$ " and " $(Q a b)$ ". An asking reasoner that receives this query as a node set schema could return a node set schema with " $(P v_1 v_2)$ " as the conclusion, $\{(v_1 a)\}$ as the variable bindings, and an inference step schema whose rule is generalized modus ponens; whose bindings are $(x a)$ and $(y b)$; and whose antecedents are A_1 , A_2 , and A_3 . The conclusion of A_1 would be " $(\Rightarrow (\text{and } (Q x y) (R y z)) (P x z))$ ", and A_1 would have an inference step whose rule is "direct assertion". The conclusion of A_2 would be " $(Q a b)$ ", and A_2 would have an inference step whose rule is "direct assertion". The conclusion of A_3 would be " $(R b z)$ ", and A_3 would have no inference steps. Thus, " $(R b z)$ " would be an unproven subgoal to be dispatched to an appropriate reasoner in an attempt to complete the proof.

4.2 Support for Explanations

Any presentation of the proof of a conclusion that is more understandable to the user than the proof is considered an explanation in this paper. Thus, the presentation of PML proofs as discussed in Section 4.2.1 is a strategy to explain conclusions. Most users in the Semantic Web, however, may be unable to understand logical proofs. For these users, the presentation of the ground assertions used to conclude C as discussed in Section 4.2.2 may be a better type of explanation for C since it gives them an understanding of what the information depended upon without going into the details of

the inference. In general, any abstraction of the proof of C is also an explanation for C . The use of rewriting rules as discussed in Section 4.2.3 is an useful way of abstracting proofs and is one that we leverage to generate explanations that provide some notion of the inferences without providing all of the details.

4.2.1 Browsing PML Documents

The IW browser is a web application used to render PML documents into human-friendly HTML documents. Given a node set's URI, the browser can render a proof by traversing from the node set's inference steps the graph composed of node sets and their inference steps. The use of the PML format provides two immediate benefits for browsing proofs:

- *Direct Access to Proof Nodes:* Proofs can easily be composed of hundreds or thousands of node sets. Users, however, may not need to browse all nodes to understand a large proof. In fact, we often see that information from few key node sets may be enough for most users to understand a proof. Further, not only may it be enough, it is preferable to only view a few critical components of the proof instead of viewing most or all of the proof. Thus, PML proofs allow users to refer to URIs of specific node sets of proofs rather than to a monolithic proof.
- *Lightweight Loading of Proofs:* The browser's **proof lens** is a metaphor of a magnifier visualizing parts of proofs. The lens **focus** is the conclusion of a given node set. The lens **magnitude** is the maximum number of levels of inference steps traversed in order to render the lens. The **refocusing** of a proof lens is a user selection of a conclusion in the lens that is not the the current focus of the lens. Thus, using PML node sets, the lens metaphor implements a lightweight loading of proofs since node sets are loaded on demand when the lens magnitude changes or the lens is refocused.

When interacting with the browser, users may select from a number of *proof styles* and *sentence formats* for displaying PML documents. Proof style is the layout used to place proof elements on the display when presenting a proof. For example, a “textbook” logic-style uses a bar to separate the conclusion of a node set from the premises of that conclusion.

The name of the inference step is placed on the right side of bar. Sentence format identifies the preferred way of formatting the conclusions of node sets. The “raw” format means that conclusions are presented as they are

represented in the PML document. Other sentence formats rely on the browser capability of translating conclusions from their original languages into the requested format. Since many users prefer to see a form of natural language output, we have provided a simple logic to English presentation module that will present nodes labeled in KIF [5] in a limited English format.

4.2.2 Browsing Knowledge Provenance from PML Documents

Users and reviewers of our work have consistently been interested in having simple ways to obtain the ground assertions that were used to find a particular conclusion C . A summary of the statements used in the proof provides one course level of abstraction of the proof. Beyond simple collections of the statements or knowledge bases used, Inference Web also provides access to the meta-information available for these sources. Thus a user may find that the particular conclusion C relied on exactly two knowledge bases (and potentially a particular set of sentences in those two knowledge bases) and additionally may learn that those two knowledge bases were considered authoritative sources by a particular verification body and the two knowledge bases were updated within the last week. That may be as much information as some users would like to see about a conclusion at a particular time. One further refinement to this service is that users may ask for the source and associated meta information just within a particular lens focus (thus just the sources used in the particular portion of the proof that is being displayed on the screen at the given time).

In practical scenarios such as those used on the ARDA AQUAINT⁵ and NIMD⁶ projects, presentation of knowledge provenance information can be large lists of assertions entailing the conclusion in the lens focus. In terms of PML proof concepts, the ground assertions are the conclusions of node sets justified by direct inference steps applying the direct assertion rule. The list of sources is a consolidation of the sources of the ground assertions' inference steps justifying the ground assertions. The meta-information of the sources are the entries of the sources in the IWBase. Thus, PML proofs are the artifacts relating knowledge provenance information in the IWBase to explanations of Semantic Web services.

⁵<http://www.ic-arda.org/InfoExploit/aquaint/>

⁶http://www.ic-arda.org/Novel_Intelligence/

4.2.3 PML Explanations

When Inference Web provides abstraction techniques for proofs, it will hide potentially a lot of information about how a conclusion was reached. We have observed that hiding many of the primitive core rules in reasoners is useful for many users. For example, users may not want to see many applications of “modus ponens” in the JTP reasoner and may instead prefer to see one application of a courser grained inference such as class transitivity. The reasoner was built using primitive rules because they were useful for efficient implementation of the reasoner, but not necessarily because they are useful for human understanding. Typically primitive rules are at the wrong level of granularity.

The rewriting of proofs based on primitive rules into proofs based on derived rules is one way of abstracting primitive rules. However, syntactic manipulations of proofs may also be insufficient for abstracting machine-generated proofs into some more understandable proofs [6]. Proofs can become more understandable if they are rewritten using IW **tactics**, that are rules derived from axioms from language descriptions such as the DAML [4] axiomatic set. In tactics, axioms are the elements responsible for aggregating some semantics in order to make the rules more understandable.

The IWBase registrar has an editor of derived rules that can be used to specify tactics. Thus, the IW explainer algorithm generates explanations in a systematic way using IWBase derived rules. Thus, for a conclusion C , the IW explainer can abstract away both a number of node sets of proof of C . The user may always ask follow-up questions and still obtain the proof of C , however the explanation of C provides abstracted explanations. The general result is to hide primitive rules and expose higher-level derived rules.

The ability to provide portions of proofs and provide support for follow-up questions has been found to be a critical component for some users ability to view explanations and proofs[8]. Inference Web follows this architectural design of being able to present stand alone components of proofs and then provide follow-up questions that are appropriate for the context in that proof.

5 Related Work

Automated reasoners have different ways to represent proofs [15]. Moreover, a reasoner can have multiple ways of representing proofs. For example, a reasoner can have an internal representation of proofs that has a number of features to handle optimizations and an external representation of proofs

used to present a trace of how conclusions are derived. Proofs exchanged between reasoners are usually external representation of proofs. Indeed, many proof elements such as optimization properties can be useless for other reasoners because the optimizations are tied to the internals of a particular engine.

External representations of proofs have been developed for several reasons. For example, most automated reasoners are able to produce a simple trace of their proofs in order to debug the reasoners themselves. The need to check proofs is a more sophisticated reason to have an external representation of proof. For example, Watson [16] created a technique to represent and check proofs produce by Isabelle [12].

Most external representations of proofs in use by more than one automated reasoning were developed within the context of hybrid reasoning systems and logical frameworks. But even in the context of hybrid reasoning systems, few are the efforts to create a general representation of proofs rather than a representation that is able to conciliate few reasoner-specific representations. PDS [1] is an example of a general representation of proofs that is used on hybrid reasoning systems such as MBase [7].

PML and PDS share the ability of representing proofs at different levels of abstraction. For instance, the PML ability to generate explanations from proofs corresponds to the PDS notion of “third dimension” for is proof representation. The MBase use of PDS is also of particular interest for the PML since it demonstrated the need of a web artifact to represent proofs. Indeed, MBase use an XML version of PDS.

We have provided a few example web services that use PML proofs. The KSL Wine Agent⁷, the DAML Query Language Client⁸, and the OWL Query Language Client all use PML proofs in order to provide interoperable explanations of their answers. The IW Browser⁹ is a web service that renders PML proofs and presents them in multiple formats for humans. All of these agent’s use the Stanford JTP hybrid reasoner but we are extending SRI’s SNARK theorem prover¹⁰ to produce portable proofs and simultaneously integrating ISI’s Prometheus¹¹ query planner with Inference web. We are also pursuing discussions with designers of other reasoning systems including

⁷<http://www.ksl.stanford.edu/people/dlm/webont/wineAgent/>

⁸<http://onto.stanford.edu:8080/dql/servlet/DQLFrontEnd>

⁹<http://belo.stanford.edu:8080/iwbrowser>

¹⁰<http://www.ai.sri.com/~stickel/snark.html>

¹¹<http://www.isi.edu/info-agents/Prometheus/>

W3C's CWM¹² and UT's KM¹³ and SRI's SPARK.

6 Conclusion and Future Work

In this article we have introduced a proof markup language – PML – that is used to support web services inter-operation and trust. PML has been integrated into the Inference Web infrastructure and is used as the interlingua that allows IW to present, combine, summarize, abstract, and explain answers generated by web services. In our discussion of hybrid reasoning and our integration efforts with JTP, we presented how PML supports hybrid reasoning and provided one implemented example of how the integration with a reasoner works. PML provides the representational foundation and Inference Web provides the tools and infrastructural support for enabling proof annotation, knowledge provenance presentation, and explanation browsing for agents and humans. This work provides the foundation from which web services may realize their promise of providing remote, interoperable access across components of distributed software systems.

Acknowledgments. The authors would like to thank Pat Hayes for many thought provoking conversations concerning explanation and for the PML name.

References

- [1] Lassaad Cheikhrouhou and Volker Sorge. PDS – A Three-Dimensional Data Structure for Proof Plans. In *Proceedings of the International Conference on Artificial and Computational Intelligence (ACIDCA '2000)*, Monastir, Tunisia, March 2000.
- [2] Dan Connolly, Frank van Harmelen, Ian Horrocks, Deborah L. McGuinness, Peter F. Patel-Schneider, and Lynn Andrea Stein. DAML+OIL (March 2001) Reference Description. Technical Report Note 18, World Wide Web Committee (W3C), December 2001.
- [3] Richard Fikes, Jessica Jenkins, and Gleb Frank. JTP: A System Architecture and Component Library for Hybrid Reasoning. Technical Report KSL-03-01, Knowledge Systems Laboratory, Stanford University, Stanford, CA, USA, 2003.

¹²<http://www.w3.org/2000/10/swap/doc/cwm.html>

¹³<http://www.cs.utexas.edu/users/mfkb/km.html>

- [4] Richard Fikes and Deborah L. McGuinness. An Axiomatic Semantics for RDF, RDF-S, and DAML+OIL (March 2001). Technical Report Note 18, W3C, December 2001.
- [5] Michael R. Genesereth and Richard Fikes. Knowledge interchange format, version 3.0 reference manual. Technical Report Logic-92-1, Computer Science Department, Stanford University, Stanford, CA, USA, 1992.
- [6] Xiaorong Huang. Reconstructing Proofs at the Assertion Level. In *Proceedings of CADE-94*, LNAI-814, pages 738–752. Springer, 1994.
- [7] Michael Kohlhase and Andreas Franke. MBase: Representing Knowledge and Context for the Integration of Mathematical Software Systems. *Journal of Symbolic Computation*, 32(4):365–402, September 2001.
- [8] Deborah L. McGuinness. *Explaining Reasoning in Description Logics*. PhD thesis, Rutgers University, 1996.
- [9] Deborah L. McGuinness and Paulo Pinheiro da Silva. Infrastructure for Web Explanations. In D. Fensel, K. Sycara, and J. Mylopoulos, editors, *Proceedings of 2nd International Semantic Web Conference (ISWC2003)*, LNCS-2870, pages 113–129, Sanibel, FL, USA, October 2003. Springer.
- [10] Deborah L. McGuinness and Paulo Pinheiro da Silva. Registry-Based Support for Information Integration. In *Proceedings of IJCAI-2003 Workshop on Information Integration on the Web (IIWeb-03)*, pages 117–122, Acapulco, Mexico, August 2003.
- [11] Deborah L. McGuinness and Frank van Harmelen. OWL Web Ontology Language Overview. Technical report, World Wide Web Consortium (W3C), December 9 2003. Proposed Recommendation.
- [12] Lawrence C. Paulson. Isabelle: the next 700 theorem provers. In P. Odifreddi, editor, *Logic and Computer Science*, pages 361–386. Academic Press, 1990.
- [13] Paulo Pinheiro da Silva, Deborah L. McGuinness, and Rob McCool. Knowledge Provenance Infrastructure. *IEEE Data Engineering Bulletin*, December 2003. To appear.

- [14] J. Reynolds. Unpublished seminar notes. Stanford University, Stanford, CA, 1966.
- [15] Geoffrey N. Watson. Proof Representations in Theorem Provers. Technical Report 98-13, Software Verification Research Centre, The University of Queensland, Queensland, Australia, September 1998.
- [16] Geoffrey N. Watson. *A Generic Proof Checker*. PhD thesis, The University of Queensland, Queensland, Australia, 2002.