

A TYPED MODEL FOR LINKED DATA

ROSS HORNE AND VLADIMIRO SASSONE

Electronics and Computer Science, University of Southampton, United Kingdom
e-mail address: rjh06r@ecs.soton.ac.uk

Electronics and Computer Science, University of Southampton, United Kingdom
e-mail address: vs@ecs.soton.ac.uk

ABSTRACT. The term Linked Data is used to describe ubiquitous and emerging semi-structured data formats on the Web. URIs in Linked Data allow diverse data sources to link to each other, forming a Web of Data. A calculus which models concurrent queries and updates over Linked Data is presented. The calculus exhibits operations essential for declaring rich atomic actions. The operations recover emergent structure in the loosely structured Web of Data. The calculus is executable due to its operational semantics. A light type system ensures that URIs with a distinguished role are used consistently. The main theorem verifies that the light type system and operational semantics work at the same level of granularity, so are compatible.

Examples show that a range of existing and emerging standards are captured. Data formats include RDF, named graphs and feeds. The primitives of the calculus model SPARQL Query and the Atom Publishing Protocol. The subtype system is based on RDFS, which improves interoperability. Examples focus on the SPARQL Update proposal for which a fine grained operational semantics is developed. Further potential high level languages are outlined for exploiting Linked Data.

1. INTRODUCTION

The model presented is a contribution to understanding the principles of the architecture of modern Web applications, which has changed significantly due to recent developments in infrastructure. Web applications can now deliver user interfaces comparable to many traditional applications. Consequently mainstream application engineering and Web application engineering are increasingly interlinked. User interface concerns can be isolated in the View of an application. The problem of moving a View onto the Web is solved by presenting the View using Web standards. Web standards are a product of transparent negotiations between industry and standards bodies.

The View is one component in the Model-View-Controller architecture, which is widely adopted for application development. Another component, the Model, provides data which forms the subject of the application. The Controller coordinates interactions with the Model to achieve some objective. Having successfully moved the View onto the Web, standards bodies are tackling the problem of moving the Model onto the Web. The claim is that moving the Model onto the Web allows common subject matter to be shared between applications. Evidence of the potential of sharing data on the Web is the ubiquity of feeds, e.g., RSS and Atom [51], which deliver data on demand between news sources and consumers.

Key words and phrases: operational semantics, type systems, semi-structured data.

Making data available on the Web gives the potential for data to link across traditional boundaries. This is enabled by using the URI as a standardised naming system. By naming a resource with a URI the resource can be referred to from any other location. Efforts to exploit these links between data sources have resulted in several proposed standards. The common aim of proposed standards is often referred to as establishing a Web of Data [11]. Data which exploits the link structure of the Web is distinguished as Linked Data. Linked Data is supported by W3C recommendations and working drafts, which reflect a consensus on the aims of the initiative [34, 14, 48]. This work draws from key standards for Linked Data and presents an executable model in which the standards coexist.

At a low level, Linked Data is delivered as messages in a semi-structured data format. The Resource Description Format (RDF) is the standardised semi-structured data format for Linked Data [34]. At this level, an HTTP request to a URI produces some RDF which describes the resource represented by the URI. No requirements are enforced on how the RDF is produced or what is done with the RDF. Message passing on channels is modelled by many process calculi [40, 13, 1, 17].

At a higher level, Linked Data can be gathered in stores which are accessed using queries. A store responds to queries as prescribed by the SPARQL Query standard [53]. Rich data sources are now published as stores, notably the UK Government Data and DBpedia [46, 12, 35]. These examples gather data, from UK Government Databases and Wikipedia respectively, then prepare the data for queries. No requirements are placed on the method of preparation. SPARQL Query has been modelled as a graph query language and using relational algebra [47, 21].

The executable model presented here is tailored to problems introduced by an update mechanism. Challenges associated with updates are highlighted by Tim Berners-Lee in a note on Read-Write Linked Data. Updates are considered at several levels of granularity. At a coarse granularity of update the contents of a store are replaced periodically. Periodic updates are adequate when data changes infrequently. An intermediate granularity is achieved by dividing a store into regions, where each region is updated independently. This idea is captured by named graphs for RDF [18]. A protocol for updating named graphs is under development [44]. Feeds and standardised protocols for feeds also work at a similar level of granularity [43, 26].

The primary challenge is to model fine grained updates at the level of triples. Triples are the basic components of RDF which resemble simple sentences in natural language of the form subject-verb-object. Fine grained updates account for exactly the triples required to perform an update. Updates which use disjoint triples may occur concurrently. By using minimal resources, an update causes minimal interruption to a store. This approach avoids regions, which are difficult to design when the long term behaviour cannot be predicted. Fine grained updates are known to present conceptual difficulties, as highlighted by Reynolds in the traditional setting of shared memory [49]. The model is a contribution to the understanding of fine grained updates for Linked Data.

Implementing Read-Write Linked Data is necessary for using Linked Data in modern applications. For instance, in wikis or social media users increasingly write data. In contrast, existing Linked Data applications are limited to reading data. Furthermore, without an update mechanism for the Model, the Model and the Controller in the modern application architecture cannot be decoupled. A Controller instead requires lower level access to the Model to perform updates. This work therefore supports efforts towards a standardised approach to Read-Write Linked Data [52].

The structure of the paper. In Section 2, a core untyped calculus is presented. The core calculus motivates the basic unit of semi-structured data as resembling a simple sentence in natural language. A high level language for queries and updates recovers expressive power. Both the data format and the language are specified using an abstract syntax and operational semantics.

In Section 3, the untyped calculus is extended with key features. The motivation for the extensions are greater interoperability with existing semi-structured data formats and common use cases such as provenance and syndication. The extra features are specified by extending the syntax and semantics of the core calculus.

In Section 4, light types for URIs and literals are introduced. The role of each type is explained and a subtype system specifies a partial order over types. A distinction between static and dynamic classes is highlighted.

In Section 5, the types and the calculus are combined. The type system specifies a type checker for data, queries and updates. A theorem verifies that the subtype system and the type system are compatible.

In Section 6, the operational semantics and the type system are combined. The typed operational semantics enable a feasible type system, by minimising runtime type checks. The type preservation theorem verifies that the type system and typed operational semantics are compatible.

In Section 7, further contributions enabled by the calculus are outlined. The type system forms a foundation for type inference. The operational semantics forms a foundation for an algebra of updates. Further features suggest high level languages for Linked Data. The demand is a more subtle notion of logic than first envisioned for the Web of Data.

2. THE CORE SYNTAX AND SEMANTICS

This section focuses the core of the key standards for Linked Data. In particular, the core of the semi-structured data format RDF and the core of the SPARQL Update language are captured. The approach is that of structural operational semantics. An abstract syntax is defined, then the operational semantics are defined by relations over the abstract syntax. Operational semantics specifies the runtime behaviour of a language.

The role of the abstract syntax differs from the role of common concrete syntaxes for RDF [18]. A concrete syntax is intended for human readability or message exchange. In contrast, the abstract syntax is for the purpose of compiler engineering. It captures the essence of the language, without redundancies. Connectives are chosen to highlight close connections to connectives in constructive logics. Brief examples of the concrete syntax are provided, then the abstract syntax is fully defined.

2.1. A Syntax for the Resource Description Framework. The Web is based on documents, represented by one URI, which link to other documents, represented by another URI. The link structure of the Web can therefore be represented by pairs of URIs. This link structure has been exploited by organisations such as Google [33]. The source and target URIs are the subject and the object of the link, respectively.

RDF extends the link structure of the Web to the power of simple sentences. In natural languages, a verb indicates how a subject is related to an object. In English for instance the structure of a simple sentence is subject–verb–object e.g. “Kleinberg writes Authoritative Sources in a Hyperlinked Environment.” RDF extends the link structure of the Web to include a predicate. The predicate serves the same role as a verb, by indicating the nature of the connection between a subject and an object.

Like the subject and the object, the predicate is also a URI. A URI is a standardised global identifier for any resource, so need not identify a document. Thus the URI of a predicate is just a global identifier from some vocabulary. Similarly, the URI of the subject and the object need not refer to a document. Instead the URI could provide a global reference to some resource which, in a

$$\begin{array}{l}
 o ::= v \quad \text{literal} \\
 | x \quad \text{variable} \quad C ::= (a a o) \quad \text{triple} \\
 | a \quad \text{name}
 \end{array}$$

Figure 1: A syntax for RDF triples (C).

traditional setting, would normally be a local identifier in a database. The following is an example of two triples.

```
soton:24123 foaf:knows soton:10511. soton:doc1 dc:creator soton:10511.
```

Note that `soton:`, `foaf:` and `dc:` represent URI prefixes `http://ecs.soton.ac.uk/rdf/`, `http://xmlns.com/foaf/0.1/` and `http://purl.org/dc/elements/1.1/` respectively. The first is a namespace used by Southampton University. The second and third are namespaces used for terminology in the Friend-of-a-Friend and Dublin Core metadata vocabularies.

Another notion generalised by RDF is the idea that a URI is associated with a document on the Web. RDF allows several pieces of traditional data to be associated with a URI. As with links between resources a predicate indicates how a URI is related to a piece of traditional data. Again this resembles simple sentences in natural language. The following is an example of two triples.

```
soton:doc1 dc:title "Tae a Link". soton:doc1 dc:description "Some poem".
```

The examples above are written using the Turtle syntax for RDF. Turtle is one of several concrete languages for presenting RDF. Here an abstract syntax captures the essence of these formats, without redundancy.

2.1.1. An Abstract Syntax for RDF. For the purpose of defining operational semantics, an abstract syntax for RDF is defined. The abstract syntax captures the idea that RDF is a collection of triples of the form subject–predicate–object, as presented in Fig. 1.

The atoms of the syntax are names, variables and literals. Names represent URIs. For simplicity of examples, names are any identifier in italics, e.g. *24123*, *knows*. Names bound by quantifiers represent place holders for URIs. Literals represent traditional data such as a string of characters or an integer, such as ‘Authoritative Sources’ or ‘7’. Traditional data is well understood so, the technicalities of literals are left to the XML datatypes specification [10]. Variables are place holders for literals, which are used for pattern matching in queries.

A triple consists of a subject, a predicate and an object. The subject and predicate are names. The object is either a name, a literal or a variable. A URI as an object generalises the notion of a link between Web pages. Similarly, the use of a literal as an object generalises the notion of the document associated with a link. The following demonstrates two triples.

$$(\textit{doc1} \textit{ creator} \textit{ Burns}) \quad (\textit{doc1} \textit{ title} \textit{ 'Links'})$$

Note that the W3C recommendation describes how to obtain labelled directed graphs from the syntax of RDF [34]. This provides a denotational semantics, which is used by graph query languages [47]. In contrast, this work remains strictly at the level of syntax. Denotational semantics for concurrency are notoriously difficult [3].

2.2. A Syntax and Semantics for Queries and Updates. When data is published openly it is rarely possible to predict how it might be used. It is therefore difficult to decide a suitable format in which to publish the data. Preferably, the application which consumes the data should decide. For this reason RDF is a simple semi-structured data format. Power is regained from this minimal structure by an expressive query language. The query language enables the consumption of emergent structures conveyed in RDF. In this way the power is shifted from the producer to the consumer and lowers barriers to publishing Linked Data.

The language SPARQL Query is the agreed standard for the purpose of querying RDF. The first SPARQL Query standard has been widely deployed [48, 12]. A second draft for SPARQL Query learns from the experiences of the first [27]. A SPARQL end point is used to observe an RDF store. The observer declares the link patterns of interest using SPARQL Query. The query language also determines the format in which results are presented.

For instance an application may be interested the question, “Obtain names for either products related to the show or products related to an exhibitor at the show.” The example scenario can be specified as follows in the SPARQL Query concrete syntax, where `eg:show2011` identifies the show and `eg:exhibitor` and `eg:product` identify predicates from some vocabulary.

```
SELECT ?product WHERE {
  {
    eg:show2011 eg:exhibitor ?exhibitor .
    ?exhibitor eg:product ?product .
  }
  UNION
  eg:show2011 eg:product ?product .
}
```

An analogy is that queries support compound sentences although RDF only supports simple sentences. Several simple sentences can verify the truth of a compound sentence. The returned result is witness to the veracity of the compound sentence in the given context. While the truth of RDF is subjective, the truth represented by a successful SPARQL Query is intersubjective [25]. For intersubjective truth there is a subjective agreement between multiple parties. The parties involved are the client that poses the query (compound sentence) and the providers of the triples (simple sentences).

Current SPARQL recommendations have no constructs for maintenance. However, there is a proposal by Hewlett-Packard Labs and a working draft for a language called SPARQL Update [53, 52]. The proposals allow RDF to be inserted and deleted at the level of triples. Update operations reuse the operations of SPARQL Query for powerful updates.

The example above can be extended to specify the update, “For either products related to the show or products related to an exhibitor at the show, insert a link from another show to that product.” The extension is to add one clause to the query. The clause inserts a triple which relates the show to the discovered product.

A model for SPARQL Update is sufficient to model SPARQL Query. However, a model for SPARQL Update is more subtle than a model for SPARQL Query. Not only is the truth represented by a successful intersubjective interaction, it is also dependent on time.

2.2.1. An abstract syntax for updates and processes. An updates is a declarative specification of the intention of the programmer. The meaning of an update is independent of a particular implementation. An application expects some behaviour and an implementation provides a behaviour

$$P \wp \perp \equiv P \quad P \wp Q \equiv Q \wp P \quad P \wp (Q \wp R) \equiv (P \wp Q) \wp R$$

Figure 2: The structural congruence for processes.

		$U ::=$	$ C$	ask	$P ::=$	U	update
			$ C^d$	delete		$ P \parallel P$	interleave
$\phi ::=$	I	true	$ \tau C^\perp$	insert		$ \perp$	nothing
	0	false	$ \phi$	filter			
	$\phi \wedge \phi$	and	$ U \oplus U$	choose	$S ::=$	P	process
	$\phi \vee \phi$	or	$ U \otimes U$	join		$ C^\perp$	store
	$\neg \phi$	not	$ *U$	iteration		$ S \wp S$	par
	\dots	etc.	$ \bigvee a.U$	select name		$ S \parallel S$	interleave
			$ \bigvee x.U$	select literal		$ \perp$	nothing
			$ \tau P$	then			

Figure 3: A syntax for constraints, updates, processes and systems

within the bounds of expectation. The common interface between the application and the store is the syntax of the language. An abstract syntax for Updates is provided in Fig. 3.

!!!!

Par is associative and commutative, so brackets and the order of triples can be ignored. The unit of par is nothing, which represents the absence of data. Content therefore forms a commutative monoid, as defined by the structural congruence in Fig. 2. !!!!!

Basic queries are formed by embedding the syntax of RDF. The embedding ‘ask’ is used to demand that some RDF is matched. The following is an example of asking a query which is satisfied by the example RDF in Section 2.1.1.

$$|(doc1 creator Burns)| \otimes |(doc1 title 'Links')|$$

Basic updates are also formed by embedding the syntax of RDF in two ways. The embedding ‘delete’ demands that some persistently stored RDF should be removed. The embedding ‘insert’ stores some RDF persistently. Unlike queries, both delete and insert have side effects. For instance the above RDF could be inserted into a store then removed from the store.

Update can be formed using two binary operations, ‘join’ and ‘choose’. Join is the operation which combines two updates to ensure that they happen in the same atomic commitment. For instance, a query and an insert can be joined to ensure that an insert occurs if and only if the query is satisfied. Choose presents an option where either the left update or the right update occurs. For instance, a choice can be presented between two possible query patterns. The iteration operation indicates that zero, one, two or more copies of an update are simultaneously applied. The constructs choose, join, iterate, true and false form a Kleene algebra.

Syntactic conventions for Kleene algebras are adopted for examples. The operator \otimes binds stronger than \oplus and the operator \otimes can be omitted; hence $(U \otimes V) \oplus W$ is abbreviated $UV \oplus W$. The following example presents a syntax for the update described in the previous section. The update is an extension of the query given in concrete syntax. Here the query is translated into abstract syntax

and an insert is joined to the query.

$$*\bigvee a. \bigvee b. \left(\begin{array}{l} |(show2011 exhibitor a)| \\ |(a product b)| \\ \oplus \\ |(show2011 product b)| \\ (show2011 product b)^\perp \end{array} \right)$$

Updates extend Kleene algebras with quantifiers. The select quantifier binds occurrences of a name not known in advance. For instance, in the example above the name of the product is not known. The name is bound in both the query and the insert, so the name discovered by the query is also the name inserted. Names and literals are disjoint, so a separate quantifier is provided for variables. Quantifiers highlight the logical content of updates.

The syntax of constraints is embedded in the syntax of updates. A constraint imposes a condition on the update taking place. Typically variables which occur as the object of a triple are constrained. For instance a variable may represent a string of characters which satisfies a regular expression, or a numeral within a range of values. Like literals, constraints are well understood, so technicalities are left to the SPARQL Query standards [48, 27]. It is sufficient to note that constraints form a Boolean algebra.

For the purpose of defining operational semantics, a syntax for processes is introduced. Processes allow updates to be considered simultaneously. The commutative monoid using par and nothing defined for RDF is extended to processes, as in Fig. 2. The convention, common to sequent style deductive systems, is that the symbol \wp is abbreviated with a comma in examples.

2.2.2. An operational semantics for atomic updates. The behaviour of an update at the level of the syntax is captured by operational semantics. A preliminary draft of an operational semantics for SPARQL Update was published in October 2010 [52]. The operational semantics presented elaborates the draft. A fine grained operational semantics for updates are specified using atomic actions.

Atomic actions are specified as a relation over processes called the commitment relation. The process on the left of the relation is exactly the processes used by the action. The process on the right of the relation is exactly the processes after the action. Thus an commitment relation describes only the local behaviour of an update. A similar approach is given by commitment relations in the π -calculus [40]. In the π -calculus there is one type of commitment – the passing of a name on a channel. Coordination of Web Services motivates extending the commitments to joins [13]. SPARQL provides a compelling reason to extend commitments to all updates. The commitment relation \triangleright is defined in Fig. 4.

The delete axioms. A simple interaction is when an update deletes a triple and the triple is available to delete. For instance, the delete and triple below are expected to interact. The result of the interaction is that the delete and the matching triple are consumed. The following is an instance of the delete axiom.

$$(doc1 creator Burns)^d, (doc1 creator Burns)^\perp \triangleright \perp$$

Notice that the axioms of Linear Logic and the atomic commitments of CCS are of a similar form [24, 39]. Note the syntactic convention of using a comma for \wp .

$$\begin{array}{c}
C^\perp \wp C^d \triangleright \perp \quad \tau C^\perp \triangleright C^\perp \quad \frac{\vDash \phi}{\phi \triangleright \perp} \quad \frac{P \wp U \triangleright P' \quad Q \wp V \triangleright Q'}{P \wp Q \wp (U \otimes V) \triangleright P' \parallel Q'} \\
\\
\frac{P \wp U \triangleright Q}{P \wp (U \oplus V) \triangleright Q} \quad \frac{P \wp V \triangleright Q}{P \wp (U \oplus V) \triangleright Q} \quad *U \triangleright \perp \quad \frac{P \wp U \triangleright Q}{P \wp *U \triangleright Q} \quad \frac{P \wp (*U \otimes *U) \triangleright Q}{P \wp *U \triangleright Q} \\
\\
\frac{P \wp U\{^b/a\} \triangleright Q}{P \wp \forall a.U \triangleright Q} \quad \frac{P \wp U\{^v/x\} \triangleright Q}{P \wp \forall x.U \triangleright Q} \quad \frac{P \wp Q \wp R \triangleright S}{P \wp (Q \parallel R) \triangleright S} \quad \frac{P \wp Q \triangleright S}{P \wp (Q \parallel R) \triangleright S \parallel R} \quad \frac{P \wp R \triangleright S}{P \wp (Q \parallel R) \triangleright S \parallel Q}
\end{array}$$

Figure 4: The axioms and rules form atomic commitments: delete axiom, insert axiom, query axiom, filter axiom, join rule, choose left rule, choose right rule, weakening axiom, dereliction rule, contraction rule, select name rule, select literal rule.

The insert axiom and query axiom. RDF to be stored after an update appears on the right of a commitment relation. There are two ways in which RDF can appear on the right. The first scenario is that a triple is inserted into a store. This is captured by the insert axiom. The insert axiom states that some RDF intended to be stored is stored by a successful update.

The second scenario is that some stored RDF can be used to answer a query. The stored RDF then returned to the store unaltered. For instance, the following example consists of a stored triple and a query asking for that triple. The query is answered and the triple remains stored.

$$\begin{array}{c}
|(doc1 creator Burns)|, \\
(doc1 creator Burns)^\perp \triangleright (doc1 creator Burns)^\perp
\end{array}$$

The syntax for an insert is the same as the syntax for some stored RDF. Therefore a trivial update which inserts some RDF is used to model stored RDF. Other SPARQL Query results may be modelled similarly to inserts, by indicating the results on the right of the commitment relation. Related calculi investigate updates and queries over inserted data as concurrent constraint satisfaction problems for Web Services [16, 50].

The join rule. The join rule forces two updates to occur in the same commitment. Join meets a requirement of SPARQL Update that a delete and insert can occur atomically. Join avoids the issue of reversing an insert, when a delete fails.

Another requirement met by join is that updates can be dependent on queries. The following example demonstrates an insert joined with a query. The available triple is adequate for the query, so the insert takes place. Both the stored triple and the inserted triple persist, so are composed after the transaction.

$$\left(\begin{array}{c} |(doc1 title 'Links')| \\ \tau C^\perp \\ (doc1 title 'Links')^\perp \end{array} \right), \triangleright (doc1 creator Burns)^\perp \parallel (doc1 title 'Links')^\perp$$

Join splits a query into two updates which can be resolved in separate locations. For instance, in the above query the two joined parts can be resolved on different machines in a cluster of servers. Thus, join serves the same purpose as join in relational algebra [28, 21]. The join rule also appears as the rule for multiplicative conjunction (times) in linear logic and as atomic commitments in process calculi for Web Services [24, 13].

The choose rules. Choice allows the programmer to specify several possible updates. The example below asks for a triple where the predicate is one of two options. The branch with the query which matches the available data is chosen. This is an external choice dependent on the available data.

$$\left((doc1 \text{ creator } Burns)^d \oplus (Burns \text{ author } doc1)^d \right), (Burns \text{ author } doc1)^\perp \triangleright \perp$$

If both branches of a choice can be enabled, one is chosen non-deterministically. The choose rules correspond to the rules for additive disjunction (plus) of linear logic and an external choice in process calculi for Web Services [24, 17].

The select rules. Most constructs work at the level of triples. Quantifiers are required to access names within triples. The select name rule works by substituting a name for the quantified name. For instance, in the example below the bound name a is replaced by $person$. This particular substitution allows the query to be answered and determines the name in the triple inserted.

$$\forall a. \left(|(doc1 \text{ creator } a)| \tau C^\perp \right), (doc1 \text{ creator } Burns)^\perp \triangleright (Hamish \text{ knows } Burns)^\perp$$

The effect above is that the substituted name is passed from the triple to the update. This is also the effect of the atomic commitments of the π -calculus [40]. The commitments of the π -calculus are decomposed into: a select which inputs the name; join which composes a continuation; and ask which poses a guard. By replacing triples with channel-value pairs and inserts with processes, the π -calculus can be recovered as investigated by Miller [30].

The select literal rule substitutes literals for variables. As above, this captures the passing of literals from triples to updates. Value passing is achieved by atomic commitments in the applied π -calculus [1]. The select rules match the rule for first-order existential quantification (some) in linear logic [24].

The constraint satisfaction relation. In general constraints form a Boolean algebra. True formulae are indicated by \vDash the constraint satisfaction relation. The definition of the constraint satisfaction relation is left to the SPARQL Query standards [48, 27]. For instance, the constraint below is satisfied when x is at least 20 years before the current year. Select substitutes x for '1987', enabling the following commitment.

$$\forall a. \forall x. \left(\begin{array}{l} |(a \text{ year } x)| \\ (\text{year-now} - x > 20) \\ \tau C^\perp \end{array} \right), \triangleright \begin{array}{l} (\text{paper year } 1987)^\perp, \\ (\text{paper copyright open})^\perp \end{array}$$

An equality comparison over names is another form of constraint. An equality comparison joined to an update, captures match found in common process calculi [40, 1]. The constraints true and false are the top and bottom elements of the Boolean algebra. True always holds, so true is embedded as the multiplicative unit in linear logic. False never holds, so like the additive zero in linear logic, no rule can be applied [24]. The embedding of Boolean algebras in Kleene algebras is elaborated by Kozen [36].

$$\frac{P \triangleright Q}{P \rightarrow Q} \quad \frac{P \rightarrow Q \quad Q \rightarrow R}{P \rightarrow R} \quad \frac{P \rightarrow P' \quad Q \rightarrow Q'}{P \wp Q \rightarrow P' \wp Q'}$$

Figure 5: Reduction relations: action, transitivity and mix.

The rules for iteration. Without iteration updates are only applied once. This enables a protocol where the programmer requests an update. The user then observes the commitment relation. If the update was not as the user intended, the update can be revoked. Then the next update is observed until the user is satisfied. Caution is exercised when the exact update is difficult to express or the content to update is not certain. When a user is certain that the update is intended, the update can be applied iteratively. The replace tool in the reader's text editor probably has similar functionality.

To apply an iterated update zero times, the weakening axiom is used. To apply an iterated update once, the dereliction rule is used. To apply an iterated update twice, the contraction rule creates two joined copies of the update. Join ensures that both copies occur in the same commitment. The example below demonstrates two nested iterations. The outermost applies twice, the innermost applies both once and twice.

$$\left(\begin{array}{l} * \forall a. \\ |(a \text{ status hidden})|, \\ * \forall b.(a \text{ knows } b)^d \end{array} \right), \begin{array}{l} (Alice \text{ status hidden})^\perp, \\ (Bob \text{ status hidden})^\perp, \\ (Alice \text{ knows Bob})^\perp, \\ (Alice \text{ knows Chris})^\perp, \\ (Bob \text{ knows Chris})^\perp \end{array} \triangleright \begin{array}{l} (Alice \text{ status hidden})^\perp \parallel \\ (Bob \text{ status hidden})^\perp \end{array}$$

Iteration is the Kleene star in regular algebra. A classic result is that nested iteration can be represented by a single iteration [19, 36]. However, quantifiers ensure that the example above cannot be expressed without nested iteration. The SPARQL Query recommendation does not have nested iteration, so cannot express the corresponding query [48].

Iteration is not replication in process calculi. Iteration defines a single commitment of an unbounded size; whereas replication persists a process across an unbounded number of commitments [40]. The use of contraction, dereliction and weakening is similar to the exponentials in linear logic, but does not correspond to either. Iteration has been used by Hoare to specify unbounded behaviour [29].

2.2.3. *A transition system for SPARQL processes.* Atomic commitment relations describe the local behaviour of an update. Only processes consumed and inserted are accounted for in a commitment relation. Another relation between processes, called the reduction relation \rightarrow , captures the evolution of a store, defined in Fig. 5.

The action rule turns a local commitment relation into a reduction relation. The action rule can be applied repeatedly, to allow multiple updates to happen in parallel. Multiple updates and idled inserts are composed in parallel using the mix rule. The transitive closure of the transition relation permits a sequence of updates on a store. The following demonstrates two separate deletes which may be applied in parallel or in sequence with the same outcome.

!!!!change this !!!!

$$\begin{array}{l} (doc1 \text{ creator Burns})^d, \\ (doc1 \text{ creator Burns})^i, \\ (doc1 \text{ title 'Links'})^d, \\ (doc1 \text{ title 'Links'})^i \end{array} \rightarrow \perp$$

The reduction relation provides a concise model of the evolution of a store, with respect to updates. The commitment relation and the reduction relation cannot be combined, without breaking the atomicity of an update. The separation of the actions (given by \triangleright) and the representation of actions on the space on which they act (given by \rightarrow) is common across dynamic systems [5, 8]. The combination of mix and transitivity defines a truly concurrent process calculus [4].

3. FEATURES FOR SYNDICATION

The core calculus focusses on fine grained updates, where updates act at the level of individual triples. This section considers a coarser level of granularity. A coarser granularity of data divides a store into regions, each regions contains triples. Each region can be considered separately from other regions. Regions impact the querying of data by allowing queries to be directed at particular regions. Regions also enable a coarser level of update, where entire regions become atomic units.

This section argues that two key Web technologies work at the granularity of regions — feeds and named graphs [51, 18]. At a suitable level of abstraction, feeds and named graphs can be queried and updated in one model. The model demonstrates that several key standards for feeds and named graphs enable a programming language for Linked Data. Furthermore, prominent examples of feeds and named graphs suggest several useful scenarios, including syndication and provenance. Both syndication and provenance have been found to be essential for a Web of Data, where separate authorities contribute separate data.

3.1. Extensions of RDF for Named Graphs. Named graphs are introduced as a minimal extension to RDF such that a large monolithic knowledge base, consisting of a single RDF store can be divided into smaller stores, each individually named [18]. The name of the graph is a URI, which can be linked to like any other URI. The RDF associated with the name of a graph, can describe the nature of the knowledge represented by the graph. Applications including provenance and access control have lead to the widespread acceptance of named graphs. Named graphs are primitive in SPARQL Query and SPARQL Update [27, 52].

The following is an example of two named graphs. The example is expressed in the TriG syntax, an extension of the Turtle syntax [18]. The first graph contains some RDF data. The second graph contains some RDF data about the first graph. This enables the user make decisions based on the source of the RDF. A user may use RDF in *eg:G2* to decide how to use data in *eg:G1*.

```
eg:G1 {
  _:Monica eg:name "Monica Murphy" .
  _:Monica eg:email <monica@murphy.org> .
}
eg:G2 {
  eg:G1 eg:author eg:Chris .
  eg:G1 eg:date "2003-09-03"^^xsd:date
  eg:G1 eg:disallowedUsage eg:Marketing.
}
```

The above example also introduces a feature of RDF. Identifiers with prefix *eg:* are URIs in some example namespace. The identifier with prefix *_:* represents a blank node. A blank node is used in place of a URI when a URI is not explicitly assigned. A blank node is a local identifier which cannot be linked to directly. Blank nodes reduce the barrier between RDF and other data formats, by allowing common data structures to be encoded in RDF without introducing new URIs.

$$\begin{array}{l}
C ::= \perp \quad \text{nothing} \\
| (a \ a \ o) \quad \text{triple} \\
| C \wp C \quad \text{par} \\
| \nu a.C \quad \text{blank node} \\
\\
G ::= \mathcal{G}_a C \quad \text{named graph} \\
| C \quad \text{default graph} \\
| G \wp G \quad \text{par} \\
| \perp \quad \text{nothing}
\end{array}$$

Figure 6: An extended abstract syntax for RDF Content and Graphs.

A constraint placed on named RDF graphs is that blank nodes are local to each individual named graph. This preserves the integrity of data structures encoded in named graphs.

3.2. An abstract syntax for named graphs. The abstract syntax of RDF content (from Fig. 1) is extended with blank nodes (in Fig. 6). Blank nodes are indicated by a quantifier for names, similarly to N3 logic [9]. The scope of the quantifier indicates the RDF content in which a blank node is bound. Bound names represent blank nodes, whereas unbound names represent URIs.

The syntax of graphs indicates both named graphs and unnamed RDF content. Named graphs are represented by a prefix with a subscript indicating the name, called the naming operator. RDF content without a naming operator represent the default graph, which allows RDF to be published without the named graph mechanism. The example from the previous section is expressed below in the abstract syntax. The blank node quantifier binds the name *Monica* in the first graph. However, the name *G1* is not bound by the naming operator so can be linked to from the second graph.

$$\begin{array}{l}
\mathcal{G}_{G1} \left(\begin{array}{l} \nu Monica. \\ \left((Monica \ name \ 'Monica \ Murphy'), \right) \\ \left((Monica \ email \ monica@murphy.org) \right) \end{array} \right), \\
\mathcal{G}_{G2} \left(\begin{array}{l} (G1 \ author \ Chris), \\ (G1 \ date \ '2003-09-03'), \\ (G1 \ disallowedUsage \ Marketing) \end{array} \right)
\end{array}$$

The above example shows that par, abbreviated by comma, is used as before to compose triples and also graphs. The structural congruence over RDF content, in Fig. 2, ensures that par and nothing form a commutative monoid. The structural congruence is extended to blank node quantifiers in Fig. 7. The first rule allows blank nodes to be eliminated if they bind nothing. The second rule allows a blank node to be distributed across some RDF content where the name does not occur. The third rule allows the order of two quantifiers to be swapped. The blank node rules preserve the free URIs in RDF content.

As standard, bound names can be α -converted. This avoids name clashes between blank nodes. Content where only blank nodes differ are equivalent, by α -conversion. This is a syntactic approach to the graph isomorphisms defined in the RDF standards. Like α -conversion, the graph isomorphisms preserve the structure and URIs but allow the blank nodes to change [34, 7].

The structural congruence also extends to naming operators. A named graph can be split into two pieces each with the same name. This is for the purpose of fine grained updates, as only the part of a graph required for an update need be considered. The naming operator and the quantifiers do not commute, so blank nodes remain within their designated graph.

$$\begin{aligned}
 va.\perp &\equiv \perp & va.C \wp D &\equiv va.(C \wp D) \quad a \notin \text{fn}(D) \\
 va.vb.C &\equiv vb.va.C & \mathcal{G}_a(C \wp D) &\equiv \mathcal{G}_a C \wp \mathcal{G}_a D
 \end{aligned}$$

Figure 7: The structural congruence extended for blank node quantifiers and naming operators.

Related work constrains named graphs so that the boundaries of a named graphs are fixed [18]. For instance, knowing the boundaries of a named graph enable the named graph to be completely dropped from a store, as in the drop operation of SPARQL Update [52]. This structural constraint is a perpendicular concern to this work. Similar constraint on global structure are tackled, for instance, in dynamic epistemic logic as structure preserving maps [8]. Stronger preservation of structure is extensively researched in the context of the Web using ontologies [31]. This work focusses on Linked Data without such a global perspective on structure.

3.3. SPARQL Update over Named Graphs. Queries and updates also work in the named graph setting. The abstract syntax for updates, in Fig. 3, is extended to named graphs, by replacing RDF content (from Fig. 1) with named graphs (from Fig. 6). The same rules for atomic commitments (in Fig. 4) work for named graphs.

The SPARQL Update submission describes an update, where the title of a book is replaced by a new title [53]. This example is captured by the commitment relation below. The delete axiom removes a triple from a graph, the insert axiom inserts the new triple in to the graph and the join rule ensures the delete and the insert happen synchronously. The presence of the naming prefix makes no difference to the rules.

$$\begin{aligned}
 &\left(\begin{array}{l} \mathcal{G}_{store}(book_3 \text{ title 'Designs'})^d \\ \mathcal{G}_{store}(book_3 \text{ title 'Design'})^i \end{array} \right), \\
 &\mathcal{G}_{store}(book_3 \text{ title 'Designs'})^i \\
 &\triangleright \mathcal{G}_{store}(book_3 \text{ title 'Design'})^i
 \end{aligned}$$

The largest example given in initial SPARQL Update drafts can now be expressed [53]. The update combines a query which finds the date of a book, a filter which checks the date is in a certain range, and an iterated update on books in that range. The iterated update will only trigger if the query and filter are satisfied. The iterated update moves triples about a book across from one graph to another graph, by combining a delete and insert. The example can be expressed in the abstract syntax, as follows.

$$\begin{aligned}
 &* \forall d. \forall book. \\
 &\left(\begin{array}{l} (d \leq \text{'01-01-2000'}) \\ |\mathcal{G}_{store1}(book \text{ date } d)| \\ * \forall a. \left(\begin{array}{l} \mathcal{G}_{store1}(book \text{ note } a)^d \\ \mathcal{G}_{store2}(book \text{ note } a)^i \end{array} \right) \end{array} \right), \\
 &\mathcal{G}_{store1} \left(\begin{array}{l} (\text{Kidnapped date '01-05-1886'}), \\ (\text{Kidnapped note classic}) \end{array} \right)^i \\
 &\triangleright \begin{array}{l} \mathcal{G}_{store1}(\text{Kidnapped date '01-05-1886'}), \\ \mathcal{G}_{store2}(\text{Kidnapped note classic}) \end{array}^i
 \end{aligned}$$

$$\begin{array}{c}
\frac{P \triangleright P' \quad Q \triangleright Q'}{P \wp Q \triangleright P' \wp Q'} \quad \frac{P \wp Q \triangleright P' \wp Q'}{\nu a.P \wp Q \triangleright \nu a.P' \wp Q'} \quad a \notin \text{fn}(Q) \\
\quad \quad \quad a \notin \text{fn}(Q') \\
\\
\frac{\mathcal{G}_b P^i \wp Q \triangleright \mathcal{G}_b P'^i \wp Q'}{\mathcal{G}_b \nu a.P^i \wp Q \triangleright \mathcal{G}_b \nu a.P'^i \wp Q'} \quad a \notin \text{fn}(Q, Q', b)
\end{array}$$

Figure 8: Commitments extended for mix, blank nodes and blank nodes in named graphs.

The above example is compact compared to the example in the draft concrete syntax. The presentation here is enabled by the constructs in the abstract syntax, whereas a single compound construct is used in the concrete syntax.

3.4. Updates for RDF with Blank Nodes. To query and update blank nodes, the commitment relation is extended in Fig. 8. The trick is to replace a quantified name with a temporary free name. This allows the quantifier to be removed and the rules of the calculus to be applied as if there were no quantifiers.

The temporary name is chosen to be fresh in the context. By choosing a fresh name, the name can be tracked before and after the commitments. This ensures that the same name that was quantified before is quantified after, as expressed by the first blank node rule. The blank node rule is similar to universal quantification in Linear Logic and new name quantification in the π -calculus [24, 40]. Note $\text{fn}(P)$ represents the name which are free in P , where select and blank node quantifier bind names.

The second blank node rule ensures that a blank node which originates in an named graph is returned to the same named graph. The extra mix rule allows triples in the scope of a blank node which are not used in a commitment to idle.

The following example demonstrates an update which involves a blank node in a named graph. The name of the blank node in the graph is replaced by a temporary name. The temporary name is discovered by the select quantifier as normal. A new triple with the temporary name is inserted into the same graph as normal. The blank node rule ensures that the temporary is bound after the commitment.

$$\begin{array}{c}
\forall person. \forall x. \forall y. \\
\left(\begin{array}{l}
|\mathcal{G}_{graph1}(person \text{ nickname } x)| \\
\mathcal{G}_{graph1}(person \text{ email } y)^i \\
(y = \text{concat}(x, '@soton.ac.uk'))
\end{array} \right), \\
\mathcal{G}_{graph1} \nu a. (a \text{ nickname 'Rabbie'})^i \\
\\
\triangleright \mathcal{G}_{graph1} \nu a. \left(\begin{array}{l}
(a \text{ nickname 'Rabbie'}), \\
(a \text{ email 'Rabbie@soton.ac.uk'})
\end{array} \right)^i
\end{array}$$

In the example the blank node does not leave the graph. Suppose that instead the update inserts the new triple into a different named graph. In this case the update cannot be applied since the blank node would appear free in another graph. The side condition would be violated.

3.5. Feeds as a ubiquitous syndication format. RDF is the format standardised by the W3C, however feeds are ubiquitous on the Web. Like RDF, feeds are a semi-structured data format which identifies resources using URIs. The two ubiquitous feed formats are RSS and Atom. RSS was originally created by Netscape and comes in several varieties. Atom has the same purpose as RSS

but is standardised [51, 43]. Atom has been adopted by Google for its Google Data protocol, which shares data between applications.

Feeds are particularly suited to syndication. Syndication is the strategy of delivering data to the intended audience on demand. Feeds typically represent the view point of some authority. A BBC News feed on Africa contains data representing the viewpoint of the BBC on the topic of news in Africa. A user who is interested in that viewpoint can obtain that feed on demand. The user can answer questions such as, “According to BBC News on Africa, what are the headlines today?”

The following is an example of the Atom Syndication Format. Notice that the feed and the entry are identified by URIs, which are abbreviated here as *eg:feed_id* and *eg:entry_id*. The tags such as *title* and *updated* are also URIs indicated by the XML namespace.

```
<feed xmlns="http://www.w3.org/2005/Atom">
  <title>Example Feed</title>
  <link href="http://example.org/" />
  <updated>2003-12-13</updated>
  <author>
    <name>John Doe</name>
  </author>
  <id>eg:feed_id</id>
  <entry>
    <title>Example Entry</title>
    <link href="http://example.org/03" />
    <id>eg:entry_id</id>
    <updated>2003-12-13</updated>
    <summary>Some text.</summary>
  </entry>
</feed>
```

The above example can be represented using named graphs and blank nodes as follows. The entries are translated into triples and form the content of a named graph. The triples associate with the feed are part of the default graph. The XML style above does not indicate a URI for the author of the feed. Below the implicit author is represented by introducing a blank node.

$$\begin{aligned}
 & \nu_{Doe} \left(\begin{array}{l} (Doe \text{ name } 'John \text{ Doe}'), \\ (feed_id \text{ author } Doe) \end{array} \right), \\
 & (feed_id \text{ title } 'Example \text{ Feed}'), \\
 & (feed_id \text{ updated } '2003-12-13'), \\
 & (feed_id \text{ link } http://example.org), \\
 & \mathcal{G}_{feed_id} \left(\begin{array}{l} (entry_id \text{ title } 'Example \text{ Entry}'), \\ (entry_id \text{ link } http://example.org/03), \\ (entry_id \text{ updated } '2003-12-13'), \\ (entry_id \text{ summary } 'Some \text{ text.}') \end{array} \right)
 \end{aligned}$$

The above syntax demonstrates one representation of Atom in RDF, however there is no standard representation. Some varieties of RSS encode feeds using triples. However, named graphs are primitive in SPARQL, so are suggested here as a representation of the content of a feed.

3.6. The Atom Publishing Protocol. For RSS an application implements its own update mechanism. In contrast, the Atom Publishing Protocol extends Atom with a standard update mechanism [26]. The publishing protocol allows new resources to be published and existing resources

to be edited. The protocol works at the low level of passing messages using an HTTP protocol. However, feeds can still be updated at the high level offered by SPARQL. This section demonstrates a high level update of a feed and outlines the corresponding low level operations which realise the high level update.

The Atom publishing protocol specification allows variations on the basic protocol. The example in this section features a main feed of articles and comment feed linked to each entry of an article. Firstly, a feed is declared such that initially it contains no resources. The data associated with the feed indicates the author of the feed and a title for the feed.

$$\begin{aligned} &(\textit{feed author Hamish}), \\ &(\textit{feed title 'Caucasus reported'}), \\ &\mathcal{G}_{\textit{feed}} \perp \end{aligned}$$

The first update, defined below, creates a new article in the feed and an empty comment feed to go with the article. The comment feed is linked to the new article. The triple associated with *entry* indicates a title and a modification date.

$$\begin{aligned} &\mathcal{G}_{\textit{feed}} \left(\begin{array}{l} (\textit{entry title 'Invaded'}) \\ (\textit{entry updated '01-02-2008'}) \\ (\textit{entry comments discussion}) \end{array} \right) \\ &(\textit{discussion subject entry}) \\ &\mathcal{G}_{\textit{discussion}} \perp \end{aligned}$$

A second update, defined below, changes the title and the date the feed was updated. The update first discovers the old title and old date using select quantifiers. The update then deletes the old triples and inserts the new triples.

$$\begin{aligned} &\forall s, d. \mathcal{G}_{\textit{feed}} \left(\begin{array}{l} (\textit{entry title } s), \\ (\textit{entry updated } d) \end{array} \right)^d \\ &\mathcal{G}_{\textit{feed}} \left(\begin{array}{l} (\textit{entry title 'Ossetia invaded'}), \\ (\textit{entry updated '02-04-2008'}) \end{array} \right)^i \end{aligned}$$

A third update creates a new comment in the comment feed associate with the entry. A query discovers the relevant comment feed and a new entry is inserted in that comment feed. The new comment is identified by a blank node rather than a URI.

$$\begin{aligned} &\forall \textit{discussion}. \left[\mathcal{G}_{\textit{feed}}(\textit{entry comments discussion}) \right]_i \\ &\quad \textit{vreaction}. \\ &\mathcal{G}_{\textit{discussion}} \left(\begin{array}{l} (\textit{reaction content 'Why?'}), \\ (\textit{reaction author Dmitri}), \\ (\textit{reaction updated '05-04-2008'}) \end{array} \right) \end{aligned}$$

The updates can be applied to the initial configuration. By applying the operational semantics the results is the configuration bellow.

$$\begin{aligned} & (\textit{feed author Hamish}), \\ & (\textit{feed title 'Caucasus reported'}), \\ & \mathcal{G}_{\textit{feed}} \left(\begin{array}{l} (\textit{entry title 'Ossetia invaded'}), \\ (\textit{entry updated '02-04-2008'}), \\ (\textit{entry comments discussion}), \end{array} \right), \\ & \textit{vreaction}. \\ & \mathcal{G}_{\textit{discussion}} \left(\begin{array}{l} (\textit{reaction content 'Why?'}), \\ (\textit{reaction author Dmitri}), \\ (\textit{reaction updated '05-04-2008'}) \end{array} \right) \end{aligned}$$

This example demonstrates that the same language for updating named graphs can be used to update feeds. The underlying operations of a publishing protocol can realise these updates. Operations of the protocol are described by the verbs post, put and get as found in a REST protocol [23].

The underlying REST operations can be outlined as follows. The first update corresponds to posting an entry to the feed and posting a new feed to the store. The second update corresponds to getting the entry and putting it back in its updated form. The third entry corresponds to getting the entry, evaluating a query and posting a new entry in the comment feed.

This example demonstrates that other semi-structured data formats, such as Atom, are compatible with RDF. It also demonstrates that SPARQL Update can be realised by the operations of a lower level protocol. The details of the low level protocol are hidden from the programmer. Related work demonstrates high level operations encoded using low level operations in the setting of Web Services [13].

4. LIGHT TYPES FOR URIs AND LITERALS

Many structural constraints, often expressed using an ontology [31], are not tackled in this work. The issue is that many invariants on structures which are imposed by an ontology require a global perspective on data. Apparently simple invariants such as, “Resources with a surname have a nick name,” are difficult to impose in an open environment. If delete removes one nickname, is there another nickname that maintains the invariant? This cannot be confirmed without knowing the extent of the entire store. Furthermore, a query which checks for a triple indicating a nickname might be unsuccessful. An unsuccessful query does not mean that the triple does not exist, only that the query was unsuccessful.

A compromise between ontologies and pure data exists. A light type system which only deals with URIs and literals in triples, rather than structures across several triples, is proposed. Given one triple it is easier to tell whether the subject and object are of the correct type for a predicate. For instance, a type system may allow the assumption that a predicate *surname* relates a person to a string. Thus for any triple in which *surname* is observed as the predicate, the subject of the triple is a person and the object is a string. No knowledge of other triples is required.

It is still naive to assume that triples can be typed. Given a literal, say ‘3’, it is reasonable to assume that ‘3’ is an integer. However given a URI, say `http://eprints.ecs.soton.ac.uk/15017/`, what is the type of the URI? Is it a person? Is it a predicate? The only sure answer is that it is a URI. Inside knowledge may say that the prefix of the URI indicates a paper, however in general such policies are not available. In general on the Web few type assumptions can be made about URIs.

Intrinsic problems associated with semi-structured data on the Web are well known. In isolating the essential aspects of semi-structured data, Abiteboul highlights prevailing challenges [2]. Abiteboul argues for a light exchange model with an a-posteriori data guide, as opposed to an a-priori schema. The model should address challenges including, eclectic types and a blurring of the distinction between schema and data. A light flexible type system for Linked Data addresses issues highlighted by Abiteboul.

For flexibility the type system works at three levels. Firstly, the XML Schema Datatype standard is reused to form a solid basic type system for RDF, where only literals are typed. Secondly, some types for URIs are moved from the data to the type system as propositional types. Thirdly, the standard inference system from the W3C standard RDFS is adapted to form a subtype system [14]. The subtype system offers flexibility and interoperability between the different strengths of typing. The three perspectives offer a compromise between static typing and dynamic data. In the presence of updates, static types are preserved while dynamic data changes.

4.1. A Standardised Type System for Literals. A type system for literals can detect basic programming errors in queries. Literals can appear in constraints in which only literals of a certain type make sense. For instance a regular expression only makes sense over a string. A comparison between literals only makes sense if the two literals are of the same type. An inequality between literals only makes sense if the two literals are of the same type and there is a natural order over that type of data.

The SPARQL Query recommendation defines the operations which appear in filters [48]. The XML Schema Datatypes recommendation is reused to define the types for literals [10]. Literals are well understood, so it is assumed that a type system for literals exists. From these standards a basic type system for updates can be defined. The basic type system annotates variables with data types as follows.

$$\forall x : \text{DATE.} \left(\begin{array}{l} ('01-01-1960' \leq x) \\ (x < '01-01-1970') \\ |(document1 \text{ created } x)| \\ (document1 \text{ note candidate})^i \end{array} \right)$$

The data types can be used in a type system to check that the constraints are correctly typed. In the example above, the constraints are inequalities between dates and a variable which is presumed to be a date. A type system accepts this update. If a constraint that checks the variable for a regular expression is also added then a type error is triggered.

Due to the data type standards, this level of typing can always be applied. Furthermore, a type inference algorithm allows the programmer to specify an update without types and still take advantage of the type system. A suitable type system and inference algorithm for literals is assumed [15].

4.2. Light Propositional Types for RDF. By typing predicates, basic errors can be detected in the structure of triples. When the choice of verb does not match the choice of subject and object, no information needs to be known about the context of a sentence to reject a sentence. A simple sentence such as, “The mountain writes the fish,” will always be nonsense.

Without type information it is less obvious that the above sentence is nonsense. Naming the mountain ‘Ararat’ and the fish ‘Nemo’ might give, “Ararat writes Nemo.” By supposing that Ararat is someone’s name and Nemo is a story, the nonsense appears to make sense. However, cultural experience suggests that Ararat refers to Mount Ararat, rather than some person. A mountain cannot be the subject of the verb to write, so the sentence remains nonsense.

DATA ::=	STRING	string type
	DATE	date type
	...	etc.
τ ::=	A	atomic proposition
	$p(\tau, \tau)$	predicate type
	$p(\tau, \text{DATA})$	data predicate type
	\top	top type
	CLASS	dynamic class type
	$\tau \cup \tau$	union type
	$\#\tau$	container type

Figure 9: The syntax of types and type environments.

The type of a URI is harder to establish than in natural language. What is the type of the URI `dbpr:Mt_Ararat`? According to DBpedia the type is another URI `dbpo:Place`. The relationship between the URI and its type can be represented by the following triple.

`dbpr:Mt_Ararat rdf:type dbpo:Place .`

Note the namespace prefixes `dbpr:` and `dbpo:` abbreviate `http://dbpedia.org/resource/` and `http://dbpedia.org/ontology/`, which are used by DBpedia for resources and terminology respectively [12]. The namespace prefix `rdf:` corresponds to the `http://www.w3.org/1999/02/22-rdf-syntax-ns#`, which is used for standardised terminology for RDF [34]. The classification of the subject is indicated by the object of the predicate `rdf:type` in a triple, as above. The URIs used for classification are referred to as classes.

The type system can treat classes in two different ways. In the first interpretation, classes are treated like any other URI. A class treated as a URI can be used in data, so can be linked to and updated as normal. In the second interpretation, a class is distinguished as an atomic type. A class distinguished as an atomic type indicates static properties of URIs. The type system is designed so that both approaches may be used. This allows an interplay between applications which use classes in data and use classes as static types. This design decision is a blurring of the distinction between schema and data, highlighted by Abiteboul [2].

4.2.1. *The syntax of propositional types.* The syntax of types which defines propositions which can be assigned to URIs, is presented in Fig. 9. The definition of types uses atomic propositions, which depend on the application.

Atomic propositional types. For a type system, a number of atomic propositional types are fixed. Atomic types are indicated by small capitals, such as `ARTICLE`, `PERSON`. Atomic types are application specific. They indicate static assumptions about a URI. For instance, an application which plots resources on a map deals with URIs typed by proposition `PLACE`. A different application which maintains a calendar of concerts may use the proposition `EVENT`.

Unlike datatypes, which restrict the structure of literals, atomic propositional types do not impose structure on a URI. Atomic propositional types are just syntax which guides how a URI is used.

Predicates between URIs. Atomic types can be used to construct predicates types. A predicate type indicates the type of a subject and object. For instance, the predicate *writes* may relate a person to a document. This is indicated by the predicate type $p(\text{PERSON}, \text{DOCUMENT})$. The predicate *knows* may relate a person to a person, indicated by type $p(\text{PERSON}, \text{PERSON})$.

Predicate types can be used to catch basic errors in triples, where the subject or the object does not match the expected type. For instance, a subject of type `LOCATION` and object of type `ANIMAL` are not valid for the predicate *writes*, under the assumptions above.

Datatype predicates. Predicates which allow literals as objects, are indicated using datatype predicates. For instance, the predicate *created* relates a document to a date literal, indicated by type $p(\text{DOCUMENT}, \text{DATE})$. This allows both literals and variables of type `DATE` to be used as the object of *created*.

In the example in Sec. 4.1, the variable appears as the object of a triple with predicate *created*. The type assumption for *created* can be used to check that the type of the variable, matches how the variable is used in the triple. In the example, the variable also appears in a constraint. A type error occurs if the type of the variable in the triple and the type of the same variable in the constraint do not match.

The top type. A URIs can be assigned the top type \top , which represents anything. The top type is used to indicate URIs where no static type information is known. This allows a light type system, by allowing any URI as the subject or object of a predicate. For instance, the atomic proposition `DOCUMENT` may be too stringent for the predicates *writes* and *created*. Instead, the types $p(\text{PERSON}, \top)$ and $p(\top, \text{DATE})$ can be used where any resource can be written and any resource can be created.

The top type can only be applied to URIs. Literals have their own top type defined in the datatype standard [10]. By keeping these two top types distinct the type system for URIs and for literals do not interfere with each other.

The union type. The union type offers a compromise between atomic propositions and top. For instance, the predicate *writes* may apply to two atomic propositions `ARTICLE` and `BOOK`. The object of the predicate becomes the union type of the two atoms, as follows $p(\text{PERSON}, \text{ARTICLE} \cup \text{BOOK})$.

The class type. For classes which appear in data (dynamic classes) the class type is introduced. The class type is just another propositional type. Just as the propositional type `PERSON` is used to type URIs of type `person`, the propositional type `CLASS` is used to type URIs which are used as classes. The predicate *rdf:type* may be assigned the type $p(\top, \text{CLASS})$.

A common misunderstanding here arises from two approaches to using using classes, illustrated as follows. A URI *person* can be assigned the type `CLASS`. This represents a dynamic class in data which may be linked to and updated like any other URI. However, the URI *Hamish* may be assigned the type `PERSON`. In this case the class `PERSON` has been lifted to the type system, so can no longer be treated like any other URI.

The mistake at this point is to link the dynamic class *person* and the static type `PERSON`. By making this mistake the static type would be of type `CLASS` and the dynamic URI would be a type. The distinction between data and types is lost. The RDFS standard unfortunately makes this mistake by not distinguishing between data and types. The result is that serious paradoxes are breached, which may be partially resolved by a higher-order type system [20]. Higher-order type systems are technically complicated and add almost nothing to this application domain. By treating `CLASS` as a simple proposition these problems disappear.

$$\begin{array}{c}
\frac{\tau_0 \leq \tau_1}{\tau_0 \leq \tau_1 \cup \tau_2} \quad \frac{\tau_0 \leq \tau_2}{\tau_0 \leq \tau_1 \cup \tau_2} \quad \frac{\tau_0 \leq \tau_2 \quad \tau_1 \leq \tau_2}{\tau_0 \cup \tau_1 \leq \tau_2} \\
\tau \leq \top \quad \frac{\tau'_0 \leq \tau_0 \quad \tau'_1 \leq \tau_1}{\mathbf{p}(\tau_0, \tau_1) \leq \mathbf{p}(\tau'_0, \tau'_1)} \quad \frac{\tau_0 \leq \tau_1 \quad \mathbf{D}_0 \leq \mathbf{D}_1}{\mathbf{p}(\tau_1, \mathbf{D}_1) \leq \mathbf{p}(\tau_0, \mathbf{D}_0)} \\
\frac{\tau_0 \leq \tau_1}{\#\tau_1 \leq \#\tau_0} \quad \frac{\tau_0 \leq \tau_1 \quad \tau_1 \leq \tau_2}{\tau_0 \leq \tau_2} \quad \tau \leq \tau
\end{array}$$

Figure 10: Axioms and rules of the subtype system: left injection, right injection, least upper bound, top, predicate, data predicate, feed, transitivity and reflexivity.

Types for feeds and named graphs. To type named graphs and feeds the container type is introduced. Typically, the subject of a triple is what is described by the triple. By analogy, the subject of a simple sentence is what is described by the sentence. A feed can therefore indicate what type of resources its triples describe. For instance, a feed of articles contains triples with subjects of type `article`. A feed of articles is indicated by the type `#ARTICLE`.

Container types are well suited to feeds. For instance, BBC News delivers feeds of articles, Flickr delivers feeds of photos and Google Calendar delivers feeds of events. However, named graphs are intended to contain diverse triples. The most general container type `#T` indicates a named graph with no restrictions on content.

Container types allow novel features. For instance, the *seeAlso* predicate indicates where to find more information about a resource. A suitable type would be `p(T, #T)`, which suggests that more information about the subject can be found in a named graph indicated by the object. Another novel type is `#(p(T, T) ∪ CLASS)` which suggests a feed of predicates and dynamic classes. This type of container is useful for presenting meta vocabularies using feeds.

Note that types for atomic propositions, predicates and union types are implicit in the RDFS standard. However, there is no standard type for named graphs, since named graphs and RDFS were introduced independently [14, 18]. The named graph type suggests one light approach to typing named graphs.

4.2.2. *A subtype system based on RDFS.* Subtypes are essential for enabling some basic scenarios. The subtype system, presented in Fig. 10, defines a preorder over types. The subtype system enables interoperability by enabling different strengths of type system to coexist. For instance, data which is heavily typed can still be used if very little type information is required. This light approach to interoperability avoids typical data integration problems, such as the integration of schema with conflicting constraints [22].

The subtype axioms. The subtype relation $\tau_0 \leq \tau_1$ can be read τ_0 is stronger than or equivalent to τ_1 . So, if something is of type τ_0 , then it is also of type τ_1 . The basic axiom of the subtype system, reflexivity, states that every type is at least as strong as itself.

Interoperability of systems can be further enabled by application specific axioms. For instance an application may define types `IMAGE` and `MEDIA`. To indicate that an image can also be treated as media, the axiom `IMAGE ≤ MEDIA` can be included. Application specific subtype axioms always relate an atomic proposition to another atomic proposition. They correspond to the `subClass` predicate in RDFS, lifted to the type system [14].

Subtypes for union types. A union type indicates that a term is of one of two types. For instance, the type `ARTICLE` is a subtype of `ARTICLE ∪ IMAGE`. This is the left injection of a type into a union type. Similarly, the type `IMAGE` is a subtype of `ARTICLE ∪ IMAGE` by the symmetric right injection rule.

A union type is the least upper bound of two types. If `ARTICLE ≤ MEDIA` and `IMAGE ≤ MEDIA` are subtype axioms, then `ARTICLE ∪ IMAGE` is also bound above by `MEDIA`. A URI of type `ARTICLE ∪ IMAGE` means that the URI may vary between identifying either an article or an image.

Subtypes for top types. Every type is bound above by the top type. Thus both `PERSON` and `p(PERSON, PERSON)` are bound above by `top`. If a URI can be anything then a person or a predicate are both suitable. This eases the restriction in ontologies that URIs for resources and predicates are disjoint [31]. Similarly, the type `CLASS` is also bound above by the top type, so dynamic classes in the data are just URIs, like resources and predicates.

Subtypes for predicates. Predicate types are contravariant in both the subject and object. Contravariance switches the direction of the subtype relation. For instance, a predicate of type `p(PERSON, ⊤)` is also of type `p(PERSON, ARTICLE)`. Contravariance allows the top type to be strengthened to the article type. So a predicate which allows anything as the object can certainly have an article as the object.

Data predicates are also contravariant in both arguments. For instance, a predicate of type `p(⊤, STRING)` can be used as a predicate of type `p(PERSON, STRING)`. The subtype relation for datatypes is defined in the XML Datatypes standard [10]. For instance, a ‘normalised string’ from the standard can be used in place of a string. Because datatypes and types for URIs are separate the subtype systems do not interfere.

Subtypes for named graphs and feeds. With subtypes, the type of feeds containing more than one type of resource can be expressed. For instance, `#(ARTICLE ∪ PHOTO)` is a feed of articles and photos collectively have their union type. The feed constructor is contravariant meaning that the type of the content of the feed may be strengthened. For instance, a feed containing resources which are either articles or photos can be treated as a feed containing only articles by ignoring the photos. This is captured by the subtype relation `#(ARTICLE ∪ PHOTO) ≤ #ARTICLE`.

4.2.3. Cut Elimination for the Subtype System. Cut elimination allows transitivity to be eliminated from subtype proofs. Cut elimination exhibits the categorical content of the subtype system. Arrows are normalised proof trees and objects are types. The theorem is justification for the rules of the subtype system.

Theorem 1 (Cut elimination). Given a proof of a subtype relation $\tau_0 \leq \tau_1$, there exists a normalised proof with the same conclusion which does not use cut (i.e., the transitivity rule in Fig. 10). Furthermore, proofs form a category such that union types are coproducts and top is terminal.

The routine proof appears in the forthcoming thesis of the first author. It relies on using the completion of the subtype axioms. For instance, consider the atomic types, `REPORTER`, `JOURNALIST`, `PERSON`, and the subtype axioms `REPORTER ≤ JOURNALIST` and `JOURNALIST ≤ PERSON`. Cut is used to determine the subtype relation `REPORTER ≤ PERSON`, which must also be a primitive subtype axiom.

The Dedekind-MacNeille completion can be used to determine all the subtype axioms for atomic types [37]. The completion associates each atomic type with a set of atomic types which are less than all types greater than the type. This set is the principle ideal of the atomic type. The ideals form a complete join semi-lattice from which the subtype axioms can be efficiently recovered.

4.2.4. *Interoperability of Subtype Systems.* A store may include Linked Data from more than one source with static type information. The subtype system enables interoperability between different subtype type systems. For instance, suppose there exist three stores, for distinct applications. Suppose that one store uses atomic types `MUSICIAN` and `VENUE`, while another store uses atomic types `PERSON` and `LOCATION`. A third store uses content from both servers, so must handle all four atomic types. Furthermore, the third server is given the subtype assumptions `MUSICIAN ≤ PERSON` and `VENUE ≤ LOCATION`, which improves interoperability of content from both servers.

In the example involving three stores, the subtype systems of the first two stores can be extended to the subtype system of the third store. A subtype system δ_0 is defined to extend to a subtype system δ_1 if and only if the completion of subtype axioms in δ_0 is contained in the completion of subtype axioms in δ_1 . Valid extensions can be checked efficiently using the Dedekind-MacNeille completion [37]. The following corollary justifies the definition of an extension of a subtype system.

Corollary 2. If δ_0 extends to δ_1 , then all subtype assumptions with respect to δ_0 are subtype assumptions with respect to δ_1 .

Subtypes ease restrictions imposed by types. Linked Data involves data from stores with different subtype systems. Corollary 2 emphasises the role of subtypes for lightweight interoperability across diverse Linked Data systems.

5. THE TYPED SYNDICATION CALCULUS

This section introduces the typed calculus. The typed calculus builds on the calculus described in previous sections. The extra type information assigns types to URIs and literals. The type system investigates the feasibility of lifting a small amount of the data to a type system. In particular, the typed calculus provides a model to evaluate the effectiveness of RDFS as a type system for Linked Data [14].

The type system allows data and updates to be statically type checked. This type system ensures that URIs assigned a distinguished role are always used consistently. Notice that the data formats and query system exists in major deployments [12]. Also, a preliminary update system is under development, so is considered a requirement for Linked Data. The type system for literals has certain benefits for catching basic programming errors. However, the type system for URIs is a design decision, rather than a requirement. It depends on the application whether a type system for URIs should be used.

Although, the type system requires a design decision, the barriers imposed by the type system are less than those imposed by traditional database schema. For instance, an application may decide that a URI refers to an article, but the data associated with that article may change. An entirely new vocabulary might be used to replace the data about an article. However the URI of type `article` remains a URI of type `article`. As long as the new vocabulary allows articles to be described, then the article can still be described.

As expected from an application dependant type system, care should be taken with what data is part of the type system. For instance, a URI may be of type `person`. It is reasonable to assume that a person will not morph into a bat, so `person` is a good choice of static type. However in an application, if a person is a banker, that person may become a bar tender. In this case role `banker` is too strong to be a static type, so should remain part of the data. For flexibility, the RDFS standard can also be considered at the level of data [42].

$$\begin{array}{lcl}
\alpha ::= & a : \tau & \text{name assignment} \\
| & x : \text{DATA} & \text{variable assignment} \\
| & \epsilon & \text{empty environment} \\
| & \alpha, \alpha & \text{environment composition} \\
\\
\alpha, \epsilon \equiv \alpha & \alpha_0, (\alpha_1, \alpha_2) \equiv (\alpha_0, \alpha_1), \alpha_2 & \\
\\
\alpha_0, \alpha_1 \equiv \alpha_1, \alpha_0 & \alpha, \alpha \equiv \alpha &
\end{array}$$

Figure 11: A syntax for type environments and structural rules over type environments: unit, associativity, exchange, contraction and weakening.

5.1. Type Rules for Linked Data and Updates. RDF content and updates are typed to ensure that URIs used in RDF content are consistent with type assumptions. This section presents type rules for both RDF content and updates. The type rules for updates ensure that an update is only well typed if it updates well typed RDF content. A type rule for each construct of content is provided in Fig. 12. In a type judgement, the turnstile \vdash separates the context on the left, represented by a type environment, from the well typed term on the right.

5.1.1. Type Environments for names and literals. Type environments are finite partial functions from names to types. Syntactically, a type assignment is a name–type pair. If the pair *Alice*: PERSON occurs, the URI *Alice* is said to be assigned type PERSON. Similarly, type assignments allow variables to be assigned datatypes. The type environment is built from comma separated type assignments. Type environment composition is associative, with the empty environment as a unit, as indicated by the congruence over type environments in Fig. 11.

The type system uses the standard structural rules exchange and contraction. Exchange allows the order of type assignments to be changed. Contraction allows two identical assumptions can be reduced to a single assumption. For instance *Alice*: PERSON, *Alice*: PERSON is equivalent to *Alice*: PERSON. Exchange and contraction are captured by the congruence over type environments in Fig. 11. These structural rules are standard for type systems. The congruence can always be applied to the environment on the left of the turnstile in a type judgement.

To ensure that environments are partial functions from URIs and variables to types and datatypes respectively, type environments must satisfy the following condition. If a URI or variable occurs in two type assignments within one type environment, then in each case the URI must be assigned the same type. Two type environments are compatible if and only if their composition still satisfies this constraint. For instance, the type environment *Alice*: PERSON and the type environment *Alice*: BOOK are incompatible. It is useful to denote the domain of a type environment α , by $\text{dom}(\alpha)$.

5.1.2. Axioms, weakening, subsumption and literals. The type system uses the standard axioms, which states that, assuming that a URI is of a particular type, the URI is of the given type. Thus if *Ossetia* is assumed to be an article, then *Ossetia* is an article. The same shape of axiom applies to variables, but types and datatypes do not overlap so the axioms are separate.

The weaken environment rule allows unused type assignments to be added to the context. So weakening can be applied to an axiom to give, if *Ossetia* is an article and *Exchange* is an article, then *Ossetia* is an article. The subsumption rule allows the subtype system to be applied at any point.

$$\begin{array}{c}
a: \tau \vdash a: \tau \quad x: \text{DATA} \vdash x: \text{DATA} \\
\\
\frac{\alpha \vdash a: \tau_0 \quad \alpha \vdash p: \mathbf{p}(\tau_0, \tau_1) \quad \alpha \vdash b: \tau_1}{\alpha \vdash (a \ p \ b): \tau_0} \\
\\
\frac{\alpha \vdash a: \tau \quad \alpha \vdash p: \mathbf{p}(\tau, \text{DATA}) \quad \alpha \vdash e: \text{DATA}}{\alpha \vdash (a \ p \ e): \tau} \\
\\
\frac{\alpha \vdash a: \#\tau \quad \alpha \vdash C: \tau}{\alpha \vdash \mathcal{G}_a C: \#\tau} \quad \frac{\alpha, a: \tau_0 \vdash P: \tau_1}{\alpha \vdash \nu a: \tau_0.P: \tau_1} \\
\\
\vdash \perp: \tau \quad \frac{\alpha \vdash P: \tau \quad \alpha \vdash P: \tau}{\alpha \vdash P \wp P: \tau} \\
\\
\frac{\alpha_0 \vdash P: \tau}{\alpha_0, \alpha_1 \vdash P: \tau} \quad \frac{\alpha \vdash P: \tau_0 \quad \tau_0 \leq \tau_1}{\alpha \vdash P: \tau_1}
\end{array}$$

Figure 12: Type rules for RDF content and named graphs: name assignment, variable assignment, type triple, type triple with literal object, type named graph, type blank node, type nothing, type par, weakening and subsumption.

So, if *Ossetia* is an article, then *Ossetia* is anything, by the top axiom from the subtype system. Weakening and subsumption enable the intuitive presentation of the type system in Fig. 12.

Data literals are defined independently from the calculus. For the purpose of examples, intuitive type judgements are assumed to hold, such as $\vdash \text{'09-09-2008'}: \text{DATE}$ or $\vdash \text{'Hamish'}: \text{STRING}$. Technical details are left to the standards [10].

5.1.3. Type rules for triples and simple RDF content. Predicate types indicate the subject and object of predicate. The type rules for triples ensure that the subject and object are of the correct type. For instance, suppose that *author* is of type $\mathbf{p}(\text{ARTICLE}, \text{PERSON})$. If *Ossetia* is of type `ARTICLE` and *Hamish* is type `PERSON` then the triple (*Ossetia author Hamish*) is well typed.

Data predicates are typed using a similar rule. The data predicate type indicates the type of the subject and the datatype of the object. For instance, suppose that the predicate *name* is of type $\mathbf{p}(\text{PERSON}, \text{STRING})$. Given that the name *Hamish* is of type `PERSON` and the literal 'Hamish' is of data type `STRING`, then the triple (*Hamish name 'Hamish'*) is well typed.

In both cases a well typed triple takes on the type of the subject. So the first triple above is of type `ARTICLE` and the second triple is of type `PERSON`. Only the type of the subject of the triple is indicated. This allows collections of triples with the same type of subject to be identified. For instance, some content may consist of triples with subjects which are articles. As noted in Sec. 4, typing triples according to the type of the subject is an application specific choice. The top type can be used to indicate that the type of the subject is irrelevant.

Triples and processes are composed using par. The type rule for par allows two triples of the same type to be composed. For instance, the two triples in this section can be composed. Subtyping

is applied to weaken the types of both triples to the appropriate union type.

$$\begin{aligned}
 & \textit{Hamish} : \text{PERSON}, \textit{Ossetia} : \text{ARTICLE}, \\
 & \textit{author} : \mathbf{p}(\text{ARTICLE}, \text{PERSON}), \textit{name} : \mathbf{p}(\text{PERSON}, \text{STRING}) \\
 \vdash & \left(\textit{Hamish give_name 'Hamish'} \right), \\
 & \left(\textit{Ossetia author Hamish} \right) : \text{PERSON} \cup \text{ARTICLE}
 \end{aligned}$$

The type judgement indicates that the locality contains triples which describe either people or articles.

5.1.4. *Type rules for blank nodes.* The blank node quantifier binds names which represent blank nodes. In the typed calculus, bound names are annotated with a type information. The rule for blank nodes first types some RDF assuming that the blank node is a normal URI. The rule then internalises the type information as a quantifier.

The following example internalises three type assumptions, which represent blank nodes. Two blank nodes indicate that they are two separate events. The third blank nodes is of the top type. The scope of the quantifier indicates that the same resource judged both events but no information is known about that resource.

$$\begin{aligned}
 & \textit{judge} : \mathbf{p}(\top, \top), \textit{date} : \mathbf{p}(\top, \text{DATE}) \\
 \vdash & \textit{va} : \top. \left(\begin{array}{l} \textit{vevent1} : \text{EVENT}. \\ \left(\begin{array}{l} (\textit{event1 judge } a), \\ (\textit{event1 date '13-01-2011'}) \end{array} \right), \\ \textit{vevent2} : \text{EVENT}. \\ \left(\begin{array}{l} (\textit{event2 judge } a), \\ (\textit{event2 date '14-01-2011'}) \end{array} \right) \end{array} \right) : \text{EVENT}
 \end{aligned}$$

Subjects bound by typed blank nodes help determine the type of the content. In the above example, since the subject of all triples are events, the whole resource is of type event.

5.1.5. *Type rules for named graphs.* The type of a named graph indicates the type of content that may be contained in the named graph. For instance, a named graph of type #ARTICLE has content of type ARTICLE. The following named graph models a feed named *Caucuses*, appearing below. The four triples in the named graph have subjects which are articles, so the feed is well typed.

$$\begin{aligned}
 & \textit{title} : \mathbf{p}(\text{ARTICLE}, \text{STRING}), \textit{published} : \mathbf{p}(\text{ARTICLE}, \text{DATE}), \\
 & \textit{editor} : \mathbf{p}(\#\text{ARTICLE}, \text{PERSON}), \textit{Caucuses} : \#\text{ARTICLE}, \\
 & \textit{Hamish} : \text{PERSON}, \textit{Ossetia} : \text{ARTICLE}, \textit{exchange} : \text{ARTICLE} \\
 \vdash & \left(\textit{Caucuses editor Hamish} \right), \\
 & \mathcal{G}_{\textit{Caucuses}} \left(\begin{array}{l} (\textit{Ossetia title 'Ossetia invaded'}), \\ (\textit{Ossetia published '09-09-2008'}), \\ (\textit{exchange title 'Stock collapse'}), \\ (\textit{exchange published '08-10-2008'}) \end{array} \right) : \top
 \end{aligned}$$

A URIs for a named graph are be treated like any other URI. Thus triples can be assigned to named graphs, such as the triple which indicates the editor of the named graph in the example above.

$$\begin{array}{c}
 \frac{\alpha \vdash G : \tau}{\alpha \vdash |G| : \tau} \quad \frac{\alpha \vdash G : \tau}{\alpha \vdash G^d : \tau} \quad \frac{\alpha \vdash \phi}{\alpha \vdash \phi : \tau} \\
 \\
 \frac{\alpha \vdash S : \tau \quad \alpha \vdash T : \tau}{\alpha \vdash S \oplus T : \tau} \quad \frac{\alpha \vdash S : \tau \quad \alpha \vdash T : \tau}{\alpha \vdash S \otimes T : \tau} \\
 \\
 \frac{\alpha, a : \tau_0 \vdash S : \tau_1}{\alpha \vdash \forall a : \tau_0. S : \tau_1} \quad \frac{\alpha, x : D \vdash S : \tau}{\alpha \vdash \forall x : D. S : \tau} \quad \frac{\alpha \vdash S : \tau}{\alpha \vdash *S : \tau}
 \end{array}$$

Figure 13: Type rules for updates: type ask, type delete, type filter, type choice, type join, type select name, type select literal, type exponential.

5.1.6. *Type rules for updates and queries.* A delete, insert or a query have the same type as the RDF content that they act on. This ensures that only RDF content which makes sense can be updated or queried. For instance, the following type judgement holds, which indicates a resource and an update which intends to replace *in* with *out*.

$$\begin{array}{l}
 \text{Dmitri} : \text{PERSON}, \quad (\text{Dmitri status in})^i, \\
 \text{status} : \text{p}(\top, \top), \quad \vdash \left(\begin{array}{l} (\text{Dmitri status in})^d \\ (\text{Dmitri status out})^i \end{array} \right) : \text{PERSON} \\
 \text{in} : \top, \text{out} : \top
 \end{array}$$

In the example above deleted and inserted data has a subject of type person, so the update maintains the type of the context. A type checker can detect malformed triples in a delete or insert before the update is applied.

5.1.7. *Type rules for select quantifiers.* Select quantifiers consist of a type assignment and an update. The type environment constrains the type of name to select. The example below selects a URI of type person. The type information permits the assumption that a selected URI will be of type person. The object of the triple in both the query and the insert are expected to be of type person.

$$\begin{array}{l}
 \text{article} : \text{ARTICLE}, \text{editor} : \text{p}(\text{ARTICLE}, \text{PERSON}), \\
 \text{club} : \top, \text{member} : \text{p}(\top, \text{PERSON}) \\
 \vdash \forall p : \text{PERSON}. \left(\begin{array}{l} |(\text{article editor } p)| \\ (\text{club member } p)^i \end{array} \right), \quad \text{vHamish} : \text{PERSON}. (\text{article editor Hamish}) : \top
 \end{array}$$

The above update is in the presence of some data where a blank node appears. The type assignment for the blank node is the same as the type assignment for the select quantifier.

5.1.8. *Type rules for literals in filters and selects.* A constraint which contains variables or names can be typed. For instance, under the assumption that $x : \text{DATE}$, constraint $x \leq \text{'01-01-1950'}$ is well typed. A date literal substituted for x results in a constraint such as $\text{'01-05-1886'} \leq \text{'01-01-1950'}$, which is also well typed. In the example below, the select quantifier introduces the assumption that x is a date. This type assumption allows the filter and triple in the query to be typed. The update

satisfies the following type judgement.

$$\begin{array}{l}
\textit{Kidnapped} : \text{BOOK}, \textit{published} : \text{p}(\text{BOOK}, \text{DATE}), \\
\textit{note} : \text{p}(\top, \top), \textit{classic} : \top \\
\hline
\forall x : \text{DATE}. \forall \textit{book} : \text{BOOK}. \\
\vdash \left(\begin{array}{l} (x \leq \text{'01-01-1950'}) \\ |(book \textit{published} x)| \\ (book \textit{note} \textit{classic})^i \end{array} \right), \quad : \text{BOOK} \\
(\textit{Kidnapped} \textit{published} \text{'01-05-1886'})
\end{array}$$

Typing literals is the minimal type system for updates. Literals can still be typed without application specific type information for URIs.

5.1.9. *Type rules for joins, choice and iteration.* Joins ensure that two well typed updates are applied atomically. A choice between two well typed updates is presented. A join or choice assumes a type that both components can assume. Iteration does not affect the type of a process. In the following example all components are of type person so the whole update is of type person.

$$\begin{array}{l}
\textit{guard} : \text{CLASS}, \textit{attendant} : \text{CLASS}, \textit{porter} : \text{CLASS}, \\
\textit{type} : \text{p}(\top, \text{CLASS}) \\
\hline
\vdash * \forall a : \text{PERSON}. \left(\begin{array}{l} \left((a \textit{type} \textit{attendant})^d \right) \\ \oplus \\ \left((a \textit{type} \textit{guard})^d \right) \\ (a \textit{type} \textit{porter})^i \end{array} \right), \quad : \text{PERSON} \\
\forall b : \text{PERSON}. (b \textit{type} \textit{attendant}), \\
\forall c : \text{PERSON}. (c \textit{type} \textit{guard})
\end{array}$$

The above example demonstrates a mix of static classes as types and dynamic classes as data. The people are always people, but their role changes.

5.2. **Algorithmic Typing for the Calculus.** In the type system in the previous section, the subsumption and weakening rules can be applied at any point. An algorithmic type system controls the use of subsumption and weakening. Subsumption can be applied as early as possible. Weakening can be applied as late as possible. The algorithmic type system can be less intuitive but is syntax directed, so easier to work with for proofs and type inference algorithms [41].

Key differences between the rules of the type system and the algorithmic type system are presented in Fig. 14. The first variation is that the axioms immediately weaken the type to the correct type required. The second variations is that, when two terms are composed, the type environments are merged whenever they are compatible. This is characterised by the type rule for join. Merging environments avoids weakening both environments before updates are combined. The third variation is that the blank node and select rules, which internalise the type environment permit weakening of the environment. This permitted weakening is expressed using the congruence over environments. Exchange and contraction still apply.

The soundness and completeness of the algorithmic type system with respect to the intuitive type system ensures that results carry from one system to the other. The proof begins with a technical lemma. The lemma demonstrates that, for the algorithmic type system, the environment on the left of the turnstile covers exactly the URIs that occur free in the term.

Lemma 3. If $\alpha \Vdash P : \tau$ then $\text{fn}(P) = \text{dom}(\alpha)$.

$$\begin{array}{c}
 \frac{\tau_0 \leq \tau_1}{a: \tau_0 \Vdash a: \tau_1} \quad \frac{\alpha, a: \tau_1 \Vdash P: \tau_0}{\alpha \Vdash \nu a: \tau_1.P: \tau_0} \quad a \in \text{fn}(P) \\
 \\
 \frac{\alpha \Vdash P: \tau_0}{\alpha \Vdash \nu a: \tau_1.P: \tau_0} \quad a \notin \text{fn}(P) \quad \frac{\alpha_0 \Vdash U: \tau \quad \alpha_1 \Vdash V: \tau}{\alpha_0, \alpha_1 \Vdash U \otimes V: \tau} \\
 \\
 \frac{\alpha_0 \Vdash a: \tau_0 \quad \alpha_1 \Vdash p: \mathfrak{p}(\tau_0, \tau_1) \quad \alpha_2 \Vdash b: \tau_1 \quad \tau_0 \leq \tau_2}{\alpha_0, \alpha_1, \alpha_2 \Vdash (a \ p \ b): \tau_2}
 \end{array}$$

Figure 14: Variations in rules for the algorithmic type system.

Soundness of the algorithmic type system is established by a straightforward rewrite from an algorithmic type tree to a normal type tree. The effect of the typing is preserved by the rewrite.

Theorem 4 (Soundness of algorithmic typing). If $\alpha \Vdash P: \tau$ then $\alpha \vdash P: \tau$.

Proof. Soundness is established by a straight forward translation of proof trees. Each algorithmic type rule which involves subtypes can be replaced by their equivalent type rule followed by a subsumption link.

For blank nodes, if $a \notin \text{fn}(P)$ then the algorithmic type rule is transformed into the type rule, preceded by application of weakening, as follows.

$$\frac{\frac{\pi}{\alpha \Vdash P: \tau_0}}{\alpha \Vdash \nu a: \tau_1.P: \tau_0} \quad \text{yields} \quad \frac{\frac{\pi}{\alpha \vdash P: \tau_0}}{\alpha_0, a: \tau_1 \vdash P: \tau_0}}{\alpha_0 \vdash \nu a: \tau_1.P: \tau_0}$$

Hence, each algorithmic type tree corresponds to a type tree with the same conclusion. \square

The proof of completeness of the algorithmic type system is a transformation of proof trees which pushes subsumption towards the leaves and weakening towards the root of a type tree.

Theorem 5 (Completeness of algorithmic subtyping). If $\alpha \vdash P: \tau$, then there exists some α_0, α_1 such that $\alpha_0, \alpha_1 \equiv \alpha$ and $\alpha_0 \Vdash P: \tau$.

Proof. The transformation $\llbracket \cdot \rrbracket$ pushes subsumption rules as deep as possible into the proof tree and suspends weakening.

There are two special cases. If two subsumption links appear consecutively, then they can be performed in a single subsumption link using cut in the subtype system. Also, weakening rules can be deleted, since weakening is controlled by the induction hypothesis.

For axioms, subsumption is absorbed by the algorithmic rule.

$$\left\| \frac{\frac{\pi}{a: \tau_0 \vdash a: \tau_0} \quad \vdash \tau_0 \leq \tau_1}{a: \tau_0 \vdash a: \tau_1} \right\| \quad \text{yields} \quad \frac{\pi}{a: \tau_0 \Vdash a: \tau_1}$$

For join, the subsumption link is pushed up each branch. By structural induction, α_0, α_1 are type environments, such that $\alpha \equiv \alpha_0, \alpha'_0$ and $\alpha \equiv \alpha_1, \alpha'_1$ and also each forms the premises of the

conclusions of the respective branches of the resulting tree.

$$\left\| \frac{\frac{\pi_0 \quad \pi_1}{\alpha \vdash P: \tau_0 \quad \alpha \vdash Q: \tau_0} \quad \pi_2}{\alpha \vdash P \otimes Q: \tau_0} \quad \vdash \tau_0 \leq \tau_1}{\alpha \vdash P \otimes Q: \tau_1} \right\|$$

yields

$$\frac{\alpha_0 \Vdash P: \tau_1 \quad \alpha_1 \Vdash Q: \tau_1}{\alpha_0, \alpha_1 \Vdash P \otimes Q: \tau_1}$$

where

$$\alpha_0 \Vdash P: \tau_1 = \left\| \frac{\pi_0 \quad \pi_2}{\alpha \vdash P: \tau_0 \quad \vdash \tau_0 \leq \tau_1} \right\|$$

and

$$\alpha_1 \Vdash Q: \tau_1 = \left\| \frac{\pi_1 \quad \pi_2}{\alpha \vdash Q: \tau_0 \quad \vdash \tau_0 \leq \tau_1} \right\|$$

The triple rule absorbs a subsumption link. By induction, there exists type environments $\alpha_0, \alpha_1, \alpha_2$ such that $\alpha \equiv \alpha_0, \alpha_1, \alpha_2, \alpha'$ and the following transformation holds.

$$\left\| \frac{\frac{\pi_0 \quad \pi_1 \quad \pi_2}{\alpha \vdash a: \tau_0 \quad \alpha \vdash p: \mathfrak{p}(\tau_0, \tau_1) \quad \alpha \vdash b: \tau_1} \quad \pi_3}{\alpha \vdash (a p b): \tau_0} \quad \vdash \tau_0 \leq \tau_2}{\alpha \vdash (a p b): \tau_2} \right\|$$

yields

$$\frac{\alpha_0 \Vdash a: \tau_0 \quad \alpha_1 \Vdash p: \mathfrak{p}(\tau_0, \tau_1) \quad \alpha_2 \Vdash b: \tau_1 \quad \vdash \tau_0 \leq \tau_2}{\alpha_0, \alpha_1, \alpha_2 \Vdash (a p b): \tau_2} \quad \pi_3$$

For blank nodes, subsumption is pushed straight up the tree. Consider the following proof tree.

$$\left\| \frac{\frac{\pi_0}{\alpha, a: \tau_2 \vdash P: \tau_0} \quad \pi_1}{\alpha \vdash \nu a: \tau_2.P: \tau_0} \quad \vdash \tau_0 \leq \tau_1}{\alpha \vdash \nu a: \tau_2.P: \tau_1} \right\|$$

By induction, there is some α_1 and α' such that the following transformation holds, where $\alpha_1, \alpha' \equiv \alpha, a: \tau_2$.

$$\left\| \frac{\pi_0 \quad \pi_1}{\alpha, a: \tau_2 \vdash P: \tau_0 \quad \vdash \tau_0 \leq \tau_1} \right\| = \alpha_1 \Vdash P: \tau_1 \quad \pi_3$$

If $a \notin \text{fn}(P)$, then the following proof tree holds.

$$\frac{\alpha_1 \Vdash P: \tau_1}{\alpha_1 \Vdash \nu a: \tau_2.P: \tau_1} \quad \pi_3$$

If $a \in \text{fn}(P)$ then, by Lemma 3 there exists α'_1 such that $\alpha_1 \equiv \alpha'_1, a: \tau_2$, and the following holds.

$$\frac{\pi_3 \quad \alpha'_1, a: \tau_2 \Vdash P: \tau_1}{\alpha'_1 \Vdash \nu a: \tau_2. P: \tau_1}$$

Further cases are similar to the above. Thus by induction over the proof trees a transformation from any type tree to an algorithmic type tree exists. \square

The algorithmic type system demonstrates that type checking is syntax directed. Even for a light type system, type checking a store at each operational step is costly. A feasible approach is to type check updates.

6. THE TYPED OPERATIONAL SEMANTICS

Given a well typed update, the expectation is that Linked Data need only be typed once. Well typed updates applied to well typed Linked Data should result in well typed Linked Data, without the need to recheck the Linked Data. The light type system works locally, in the sense that the correctness of one triple is not affected by other triples. Similarly, the commitment relation over updates describes the local behaviour of updates, since unused triples are ignored. The type system and commitment relations therefore work at the same level of granularity, so are compatible.

This section demonstrates that the specification of atomic commitments can be extended to ensure that the type system and the commitment relation are compatible. The typed commitment rules introduce minimal assumptions about the context. The assumptions about the context are that names selected in an update are of the correct type. This amounts to a minimal dynamic type check on selected names. Under minimal assumptions about the context, type judgements are preserved by the dynamically typed commitment relation, as verified by Theorem 8.

The typed operational semantics are defined by combining the following components. α -conversion of bound names, the structural congruence in Fig. 15 and a typed commitment relation in Fig. 16. Examples throughout this section illustrate the operational behaviour of typed updates.

6.1. The Structural Congruence for Typed Linked Data. The structural congruence for typed content is gathered in Fig. 15. The structural congruence captures the commutative monoid formed by par and nothing, which is used for both RDF content and processes. The structural congruence allows blank node quantifiers to distribute over tensor and be eliminated in the presence of nothing. Compatible blank node quantifiers may be swapped. The side condition for swapping type assignments ensures that the composition of the assignments form a partial function. Type environments must be partial functions. The split named graph rule allows named graphs to be decomposed for fine grained updates.

The first type preservation result verifies that the structural congruence preserves types. Thus given a well typed process, processes structurally congruent to the process are well typed. Lemma 6 verifies this compatibility between the structural congruence and the type system. The proof makes use of algorithmic typing to simplify proofs.

Lemma 6 (Structural congruence preserves types). Assuming that $P \equiv Q$, $\alpha \vdash P: \tau$ if and only if $\alpha \vdash Q: \tau$.

$$\begin{aligned}
P \wp \perp &\equiv P & P \wp (Q \wp R) &\equiv (P \wp Q) \wp R \\
P \wp Q &\equiv Q \wp P & \mathcal{G}_a(C \wp D) &\equiv \mathcal{G}_a C \wp \mathcal{G}_a D \\
va: \tau.(P \wp Q) &\equiv va: \tau.P \wp Q & a &\notin \text{fn}(Q) \\
va: \tau.\perp &\equiv \perp & va: \tau_0.vb: \tau_1.P & \\
& & \equiv vb: \tau_1.va: \tau_0.P & \quad a \neq b \text{ or } \tau_0 = \tau_1
\end{aligned}$$

Figure 15: The structural congruence over content and processes: unit, associativity, commutativity, split named graph, distribute blank node, eliminate blank node and commute blank node.

Proof. For each algorithmic type tree and rule of structural congruence, an algorithmic type tree of the equivalent process can be constructed, by Theorems 4 and 5.

For distributivity of blank nodes over par, assume that $a \notin \text{fn}(Q)$. Also assume that $a \in \text{fn}(P)$ and the following proof tree holds.

$$\frac{\frac{\alpha_0, a: \tau_0 \Vdash P: \tau_1}{\alpha_0 \Vdash va: \tau_0.P: \tau_1} \quad \alpha_1 \Vdash Q: \tau_1}{\alpha_0, \alpha_1 \Vdash va: \tau_0.P \wp Q: \tau_1}$$

Now $a \notin \text{dom}(\alpha_1)$, by Lemma 3, and $a \notin \text{dom}(\alpha_0)$ thus $a \notin \text{dom}(\alpha_0, \alpha_1)$, so the following proof tree holds. The converse is immediate.

$$\frac{\frac{\alpha_0, a: \tau_0 \Vdash P: \tau_1 \quad \alpha_1 \Vdash Q: \tau_1}{(\alpha_0, \alpha_1), a: \tau_0 \Vdash P \wp Q: \tau_1}}{\alpha_0, \alpha_1 \Vdash va: \tau_0.(P \wp Q): \tau_1}$$

Now, assume that $a \notin \text{fn}(P)$ and the following proof tree holds. Clearly $a \notin \text{fn}(P \wp Q)$ so the following proof trees can be interchanged.

$$\frac{\frac{\alpha_0 \Vdash P: \tau_1}{\alpha_0 \Vdash va: \tau_0.P: \tau_1} \quad \alpha_1 \Vdash Q: \tau_1}{\alpha_0, \alpha_1 \Vdash va: \tau_0.P \wp Q: \tau_1}$$

iff

$$\frac{\frac{\alpha_0 \Vdash P: \tau_1 \quad \alpha_1 \Vdash Q: \tau_1}{\alpha_0, \alpha_1 \Vdash P \wp Q: \tau_1}}{\alpha_0, \alpha_1 \Vdash va: \tau_0.(P \wp Q): \tau_1}$$

Remaining cases are straight forward. The result follows by induction over the derivation of an equivalence. \square

The structural congruence covers the reorganisation of content and processes. The structural congruence is always reversible. In contrast, the effect of updates are generally irreversible, so are captured by a commitment relation.

6.2. Typed Atomic Commitments. Atomic commitments were introduced in Sec. 2 to specify an operational semantics for queries and updates over Linked Data. In Sec. 3, atomic commitments were extended to cover key features for syndication. In this section, atomic commitments are extended with a type environment, called the context. Otherwise, the role of atomic commitments remains the same. The process on the left indicates exactly the processes consumed. The processes on the right indicates the exact processes which replace the processes consumed.

The context for typed atomic commitments indicates a minimal dynamic type check required by a commitment. By minimising dynamic type checks, a feasible type system is enabled. The context represents these minimal type checks as a type environment. Any processes in the vicinity of the commitment must agree on the assignments of names to types in the context.

For instance, a context for a commitment relation may indicate that the name *Burns* is of type `WRITER`. However, if it is assumed elsewhere that *Burns* is of type `PERSON`, then the commitment cannot be applied, since the required context indicates a stronger type. The rules for typed commitments are presented in Fig. 16. The higher-order π -calculus similarly constrains the context of a transition using type environments [32].

6.2.1. Fully type safe commitments. Assuming that a process is well typed, a commitment which only uses axioms requires no dynamic type checks. When there are no type checks most other rules behave like the untyped calculus. For instance, in the example below the deletes and inserts have an empty context, so their join has an empty context.

$$\vdash \left(\begin{array}{l} (\textit{studio status closed})^d \\ (\textit{studio status open})^i \\ \mathcal{G}_{\textit{studio}}(\textit{Dmitri status out})^d \\ \mathcal{G}_{\textit{studio}}(\textit{Dmitri status in})^i \end{array} \right),$$

$$\begin{array}{l} (\textit{studio status closed}), \\ \mathcal{G}_{\textit{studio}}(\textit{Dmitri status out}) \\ \triangleright (\textit{studio status open}), \\ \mathcal{G}_{\textit{studio}}(\textit{Dmitri status in}) \end{array}$$

Because the context above is empty, any environment which types the process before the commitment also types the process after the commitment.

6.2.2. The dynamically typed select quantifier. The select quantifier introduces the need for dynamic type checks. A select quantifier annotates a name with a type. The type annotation imposes an upper bound on the type of the selected name. For instance, the select quantifier below requires that the selected name is of type `PERSON`. The commitment selects the name *Dmitri* according to the triple to be deleted. However, the given process does not indicate that *Dmitri* is of type `PERSON`. The missing assumption is indicated by the context in front of the commitment.

$$\forall a : \text{PERSON}.$$

$$\textit{Dmitri} : \text{PERSON} \vdash (\textit{Hamish knows } a)^d, \quad \triangleright \perp$$

$$(\textit{Hamish knows } \textit{Dmitri})$$

The context above indicates that the commitment can only be applied safely when *Dmitri* is of type `PERSON`. Further information required to type the process, such as *knows* is of type `p(PERSON, PERSON)` and *Hamish* is of type `PERSON`, is not required for the commitment.

$$\begin{array}{c}
\vdash C^i \wp C^d \triangleright \perp \quad \vdash C^i \triangleright C^i \quad \vdash C^i \wp |C| \triangleright C^i \\
\\
\frac{\alpha \vdash P \wp U \triangleright Q}{\alpha \vdash P \wp (U \oplus V) \triangleright Q} \quad \frac{\alpha \vdash P \wp V \triangleright Q}{\alpha \vdash P \wp (U \oplus V) \triangleright Q} \\
\\
\frac{\alpha_0 \vdash P \wp U \triangleright P' \quad \alpha_1 \vdash Q \wp V \triangleright Q'}{\alpha_0, \alpha_1 \vdash P \wp Q \wp (U \otimes V) \triangleright P' \wp Q'} \quad \frac{\varepsilon \phi}{\vdash \phi \triangleright \perp} \\
\\
\frac{\alpha \vdash P \wp U \triangleright Q}{\alpha \vdash P \wp *U \triangleright Q} \quad \frac{\alpha \vdash P \wp (*U \otimes *U) \triangleright Q}{\alpha \vdash P \wp *U \triangleright Q} \\
\\
\frac{\alpha \vdash P \wp U\{^b/a\} \triangleright Q \quad \vdash \tau_0 \leq \tau_1}{\alpha, b: \tau_0 \vdash P \wp \bigvee a: \tau_1. U \triangleright Q} \\
\vdash *U \triangleright \perp \\
\\
\frac{\alpha \vdash P \wp U\{^v/x\} \triangleright Q \quad \vdash v: D}{\alpha \vdash P \wp \bigvee x: D. U \triangleright Q} \\
\\
\frac{\alpha_0 \vdash P \triangleright P' \quad \alpha_1 \vdash Q \triangleright Q'}{\alpha_0, \alpha_1 \vdash P \wp Q \triangleright P' \wp Q'} \\
\\
\frac{\alpha, a: \tau \vdash P^i \wp Q \triangleright P^i \wp Q'}{\alpha \vdash \nu a: \tau. P^i \wp Q \triangleright \nu a: \tau. P^i \wp Q'} \quad a \notin \text{fn}(Q, Q') \\
\\
\frac{\alpha, a: \tau \vdash \mathcal{G}_b P^i \wp Q \triangleright \mathcal{G}_b P^i \wp Q'}{\alpha \vdash \mathcal{G}_b \nu a: \tau. P^i \wp Q \triangleright \mathcal{G}_b \nu a: \tau. P^i \wp Q'} \quad a \notin \text{fn}(Q, Q', b)
\end{array}$$

Figure 16: The axioms and rules form atomic commitments: delete axiom, insert axiom, query axiom, choose left rule, choose right rule, join rule, filter axiom, dereliction rule, contraction rule, weakening axiom, select name rule, select literal rule, mix rule, blank node rule, named graph rule.

6.2.3. *Joined commitments with non-empty context.* Join is used to synchronise updates. If two updates each require a context, then the join of the updates composes the contexts. For instance, the following example consists of two commitments where each requires a URI to be of type person.

The commitments are joined so the context indicates that both URIs are of type `PERSON`.

$$\begin{array}{l}
 \text{user1: PERSON, user2: PERSON} \\
 \\
 (user1 \text{ status busy}), \\
 (user2 \text{ status ready}), \\
 \vdash \left(\begin{array}{l} \forall a: \text{PERSON.} \\ \left(\begin{array}{l} (a \text{ status busy})^d \\ (a \text{ status ready})^i \end{array} \right) \\ \forall b: \text{PERSON.} \\ \left(\begin{array}{l} (b \text{ status ready})^d \\ (b \text{ status busy})^i \end{array} \right) \end{array} \right) \triangleright \begin{array}{l} (user1 \text{ status ready}), \\ (user2 \text{ status busy}) \end{array}
 \end{array}$$

In the above example, the names in the context are distinct. The join rule forces combined contexts to be disjoint. By forcing disjoint contexts, two select quantifiers cannot discover the same name. Consequently, the more controlled ‘select distinct’ quantifier is modelled from SPARQL Query [48].

In contrast, the join rule in Sec. 2 models the normal select quantifier in SPARQL Query. The normal select allows different selects to discover the same name. The normal quantifier can be achieved here by removing the constraint that joined contexts are disjoint. Removing the constraint allows the same name to appear in the combined environment, hence contraction may be applied. Contraction allows two different select quantifiers to share the same resource. The two variations on the select quantifier may coexist by extending the type environment in the calculus. A more subtle type environment can control the use of contraction, as investigated in the logic of bunched implications [45].

6.2.4. *Dynamic type checks for selected literals.* The select literal quantifier annotates a variable with a data type. The annotation constrains the type of a literal discovered using the select literal rule. To enforce the constraint, the select literal rule dynamically type checks the selected literal. In the example below, the literal input by the select quantifier is successfully checked to be a date. The syntax of the literal is enough information to check the type.

$$\begin{array}{l}
 \text{Kidnapped: BOOK} \\
 \\
 \forall x: \text{DATE.} \forall book: \text{BOOK.} \\
 \vdash \left(\begin{array}{l} (x \leq \text{'01-01-1950'}) \\ |(book \text{ published } x)| \\ (book \text{ status classic})^i \end{array} \right), \\
 (\text{Kidnapped published '01-05-1886'}) \\
 \triangleright (\text{Kidnapped published '01-05-1886'}), \\
 (\text{Kidnapped status classic}),
 \end{array}$$

The select literal performs the dynamic type check immediately. No further information about the literal is required from the environment. In contrast, there is not enough information to check the name *Kidnapped* is a book. This minimum requirement placed on the context is indicated by the type environment.

6.2.5. *Typed Commitments involving Choice.* The branches in a choice may depend on different contexts. In the update below a person and a string are always selected. The string is immediately type checked and the check for the person is indicated by the context. The update features a third select which demands a name of type place, but alternatively offers the choice of the unit update.

In the commitment below the unit branch is chosen, so the third select does not contribute to the context.

$$\begin{array}{l}
\text{Burns: PERSON} \\
\vdash \forall x: \text{STRING}. \forall a: \text{PERSON}. \\
\left(\begin{array}{l} |(a \text{ email } x)| \\ \mathcal{G}_{\text{poets}}(a \text{ email } x)^i \\ \left(\begin{array}{l} \forall h: \text{PLACE}. \\ \left(\begin{array}{l} |(a \text{ home } h)| \\ \mathcal{G}_{\text{poets}}(a \text{ home } h)^i \end{array} \right) \oplus \mathbf{I} \end{array} \right) \end{array} \right), \\
(a \text{ email 'Rabbi@soton.ac.uk'}) \\
\triangleright (\text{Burns email 'Rabbi@soton.ac.uk'}), \\
\triangleright \mathcal{G}_{\text{poets}}(\text{Burns email 'Rabbi@soton.ac.uk'})
\end{array}$$

The choice between an update and the unit update models the operator ‘opt’ in SPARQL Query [48, 47]. Since the unit update is always enabled, the other branch may always be ignored so is optional. This demonstrates that, firstly, opt is not primitive and, secondly, opt works for updates. In related work, opt is borrowed from relational algebra for modelling queries [21].

6.2.6. *Iterated updates and dynamic types.* Iteration allows multiple copies of an update to be applied. For instance, the following iteration joins two copies of the inner update. Due to the use of join in the contraction rule, selected names are forced to be disjoint. In the following example, the context demands three disjoint names of type person.

$$\begin{array}{l}
\text{Alice: PERSON, Bob: PERSON, Chris: PERSON} \\
\forall a: \text{PERSON}. \\
\vdash \left(\begin{array}{l} |(a \text{ type journalist})| \\ * \forall b: \text{PERSON}. \\ \left(\begin{array}{l} |(b \text{ type photographer})| \\ (a \text{ knows } b)^i \end{array} \right) \end{array} \right), \\
(\text{Alice type journalist}), \\
(\text{Bob type photographer}), \\
(\text{Chris type photographer}) \\
\triangleright (\text{Alice type journalist}), (\text{Alice knows Bob}), \\
(\text{Alice knows Chris}), (\text{Bob type photographer}), \\
(\text{Chris type photographer})
\end{array}$$

Now suppose that the whole of the above update is also iterated. Due to the disjunction of environments forced by join, each journalist is assigned distinct photographers. To allow names to be shared the join rule can be relaxed, as discussed above.

6.2.7. *Commitments for typed blank nodes.* The blank node rule allows a blank node to be used in place of a URI. The typed blank node also indicates a lower bound on the type of URI the blank node can represent. For instance, the example below involves a blank node quantifier annotated with type person. The query demands a name of any type, so the assumption that the blank node is

of type person is strong enough for the following commitment.

$$\frac{\begin{array}{l} \forall b: \top. \\ \vdash \left(\begin{array}{l} |(b \text{ name 'Burns'})| \\ (Dmitri \text{ knows } b)^i \end{array} \right), \\ \nu a: \text{PERSON.} \\ (a \text{ name 'Burns'}) \end{array}}{\vdash \left(\begin{array}{l} \nu a: \text{PERSON.} \\ (a \text{ name 'Burns'}), \\ (Dmitri \text{ knows } a) \end{array} \right)} \triangleright$$

The context is used to ensure that type of the select quantifier and the blank node quantifier match. The select quantifier introduces to the context an assignment of a name to type person. The blank node rule eliminates that assignment from the context. In the above example, this leaves an empty context so no dynamic checks are required.

6.3. Type Preservation for Commitments. Type preservation verifies that given a well typed process the resulting process after a commitment is well typed with respect to the same environment. This means that the use of a URI after an update is consistent with the use of a URI before the update. For a commitment relation with a non-empty context, the context must agree with the type environment used to type the process. The following substitution lemma is required for selected names and literals.

Lemma 7 (Substitution preserves types). For names,

$$\text{if } \alpha, a: \tau_1 \vdash U: \tau \text{ and } \tau_0 \leq \tau_1, \text{ then } \alpha, b: \tau_0 \vdash U\{b/a\}: \tau.$$

Similarly for literals,

$$\text{if } \alpha, x: D \vdash U: \tau \text{ and } \vdash v: D, \text{ then } \alpha \vdash U\{v/x\}: \tau.$$

The proof of the lemma follows by structural induction. The type preservation theorem also uses type preservation of the structural congruence, Lemma 6. The soundness and completeness of the algorithmic type system eliminate the need to consider subsumption and weakening rules, Theorems 4 and 5.

Theorem 8 (Commitments preserve types). If $\alpha_0 \vdash P \triangleright Q$, then $\alpha_0, \alpha_1 \vdash P: \tau$ yields that $\alpha_0, \alpha_1 \vdash Q: \tau$.

Proof. The axioms are immediate. The structural induction proof for choose, join, select and blank nodes are demonstrated.

Consider the choose rule and assume that the following type tree holds.

$$\frac{\frac{\alpha_0 \vdash P: \tau \quad \frac{\alpha_1 \vdash U: \tau \quad \alpha_2 \vdash V: \tau}{\alpha_1, \alpha_2 \vdash U \oplus V: \tau}}{\alpha_0, \alpha_1, \alpha_2 \vdash P \wp (U \oplus V): \tau}}$$

Therefore the following type tree holds.

$$\frac{\alpha_0 \vdash P: \tau \quad \alpha_1 \vdash U: \tau}{\alpha_0, \alpha_1 \vdash P \wp U: \tau}$$

Now assume that the choose left rule is used to resolve a commitment, where $\alpha, \alpha' \equiv \alpha_0, \alpha_1, \alpha_2$.

$$\frac{\alpha \vdash P \wp U \triangleright Q}{\alpha \vdash P \wp (U \oplus V) \triangleright Q}$$

By induction, $\alpha \vdash P \wp U \triangleright Q$ and $\alpha_0, \alpha_1 \vdash P \wp U: \tau$ yields the following type judgement, as required.

$$\alpha_0, \alpha_1, \alpha_2 \vdash Q: \tau$$

Consider the join rule and suppose that the following type judgement holds.

$$\frac{\frac{\alpha_0 \vdash P : \tau \quad \alpha_1 \vdash Q : \tau}{\alpha_0, \alpha_1 \vdash P \wp Q : \tau} \quad \frac{\alpha_2 \vdash U : \tau \quad \alpha_3 \vdash V : \tau}{\alpha_2, \alpha_3 \vdash U \otimes V : \tau}}{\alpha_0, \alpha_1, \alpha_2, \alpha_3 \vdash P \wp Q \wp (U \otimes V) : \tau}$$

Hence the following two judgements hold.

$$\alpha_0, \alpha_2 \vdash P \wp U : \tau \quad \text{and} \quad \alpha_1, \alpha_3 \vdash Q \wp V : \tau$$

Now, assume that the following commitment holds, where $\alpha_0, \alpha_2 \equiv \beta_0, \beta'_0$ and $\alpha_1, \alpha_3 \equiv \beta_1, \beta'_1$.

$$\frac{\beta_0 \vdash P \wp U \triangleright P' \quad \beta_1 \vdash Q \wp V \triangleright Q'}{\beta_0, \beta_1 \vdash P \wp Q \wp (U \otimes V) \triangleright P' \wp Q'}$$

Hence by induction, the following type judgement holds, as required.

$$\frac{\alpha_0, \alpha_2 \vdash P' : \tau \quad \alpha_1, \alpha_3 \vdash Q' : \tau}{\alpha_0, \alpha_1, \alpha_2, \alpha_3 \vdash P' \wp Q' : \tau}$$

Consider the select rule and suppose that the following type tree holds.

$$\frac{\frac{\alpha_1, a : \tau_1 \Vdash U : \tau}{\alpha_1 \Vdash \bigvee a : \tau_1. U : \tau}}{\alpha_0, \alpha_1 \Vdash P \wp \bigvee a : \tau_1. U : \tau}$$

Assuming that $\tau_0 \leq \tau_1$, by the substitution lemma, $\alpha_1 + a : \tau_1 \vdash U : \tau$ yields $\alpha_1, b : \tau_0 \vdash U\{b/a\} : \tau$, so the following proof tree can be constructed.

$$\frac{\alpha_0 \vdash P : \tau \quad \alpha_1, b : \tau_0 \vdash U\{b/a\} : \tau}{\alpha_0, \alpha_1, b : \tau_0 \vdash P \wp U\{b/a\} : \tau}$$

Also, assume that the following commitment holds, where $\alpha_0, \alpha_1 \equiv \alpha, \alpha'$, for some α' .

$$\frac{\alpha \vdash P \wp U\{b/a\} \triangleright Q}{\alpha + b : \tau_0 \vdash P \wp \bigvee a : \tau_1. U \triangleright Q}$$

By the induction hypothesis, $\alpha \vdash P \wp U\{b/a\} \triangleright Q$ and $\alpha_0, \alpha_1, b : \tau_0 \vdash P \wp U\{b/a\} : \tau$ yield the following, as required.

$$\alpha_0, \alpha_1, b : \tau_0 \vdash Q : \tau$$

Consider the blank node rule and assume that the following type tree holds.

$$\frac{\alpha_0 \vdash \nu a : \tau_0. P : \tau \quad \alpha_1 \vdash Q : \tau}{\alpha_0, \alpha_1 \vdash \nu a : \tau_0. P \wp Q : \tau}$$

Hence, assuming that $a \notin \text{fn}(Q)$, the following type tree holds.

$$\frac{\alpha_0 + a : \tau_0 \vdash P' : \tau \quad \alpha_1 \vdash Q : \tau}{\alpha_0, \alpha_1 + a : \tau_0 \vdash P \wp Q : \tau}$$

Hence by induction, there exists α'_0, α'_1 such that $\alpha_0, \alpha_1 \equiv \alpha'_0, \alpha'_1$ and the following type tree holds.

$$\frac{\alpha'_0 + a : \tau_0 \vdash P' : \tau \quad \alpha'_1 \vdash Q' : \tau}{\alpha'_0, \alpha'_1 + a : \tau_0 \vdash P' \wp Q' : \tau}$$

Therefore, assuming that $a \notin \text{fn}(Q')$ the following proof tree holds, as required.

$$\frac{\frac{\alpha'_0 + a : \tau_0 \vdash P' : \tau}{\alpha'_0 \vdash va : \tau_0.P' : \tau} \quad \alpha'_1 \vdash Q' : \tau}{\alpha'_0, \alpha'_1 \vdash P' \wp Q' : \tau}$$

The remaining cases follow a similar pattern. Therefore, by induction on the structure of a commitment derivation, types are preserved by atomic commitments. \square

6.3.1. Monotonicity of contexts. The examples in the previous section indicate the weakest context for a transition. However, the example involving the blank node quantifier uses subtyping in the select name rule to select a stronger name. Similarly, in all examples subtyping allows a stronger type to be used in the environment.

For instance, a context which requires that *Burns* is of type `PERSON`, is satisfied by context which instead assigns the subtype `WRITER` to the same name. In general, Proposition 9 verifies that a stronger context can be used in place of a weaker context without breaking a commitment. The preorder extends subtyping point-wise to environments.

Proposition 9 (Monotonicity). If $\alpha_0 \leq \alpha_1$, then $\alpha_1 \vdash P \triangleright Q$ yields that $\alpha_0 \vdash P \triangleright Q$.

The proof pushes the strengthening of the context towards the select quantifiers, where it is eliminated. A similar proof shows the monotonicity of typing. Monotonicity facilitates integration by allowing processes to be moved to a stronger environment without further type checks.

6.3.2. Recovering the untyped calculus. The relationship between the typed and untyped calculus is acknowledged through erasure. Erasure removes all type annotations whilst retaining operational behaviour. In particular, the type annotations which appear in select and blank node quantifiers are removed, as defined by the transformation `erase`. As verified in Proposition 10, all transitions possible in the typed calculus are possible in the untyped calculus. For an exact match, the join rule is relaxed to remove requirement that joined context are disjoint.

Proposition 10. If $\alpha \vdash P \triangleright Q$, then `erase` $P \triangleright$ `erase` Q .

However, as expected, the converse does not hold. There exist transitions in the untyped calculus that are impossible in the typed calculus. For instance in the following example the blank node can only be selected after erasure, since `⊤` is not a subtype of `DOCUMENT`.

$$\text{erase} \left(\frac{va : \top.(a \text{ status official}),}{\vee b : \text{DOCUMENT}.(b \text{ status official})^d} \right) \triangleright \perp$$

The main differences between the typed commitments in Fig 16, and the untyped commitments in Figs. 4 and 8, are summarised as follows. The select name quantifier inputs a name of a given type, rather than any name. The select literal rule checks the datatype of the literal, rather than accepting any literal. Rules propagate resulting constraints on the context, whereas the untyped calculus does not constrain the context. The blank node quantifier rule simulates URIs of a given type, rather than any URI.

The combination of the subtype system, literal only typing, monotonicity and erasure allow different strengths of type system to be used in different applications.

7. FURTHER WORK

The calculus provides an abstract syntax and operational semantics which model a language where Linked Data is primitive. The focus of this work is to investigate the capabilities of a light type system for Linked Data programming languages. The rules of the type system provide a model, which can be used to implement type inference algorithms for Linked Data. The operational semantics can also be used to derive an algebra for optimisations of queries and updates. A range of extensions can enrich the calculus for particular applications. In particular, extensions for data interoperability and further higher level programming languages are outlined.

7.1. Type inference algorithms. Type inference reduces constraints imposed by the type system by inferring types from partial type information. An update can be provided untyped by a programmer. An algorithm then automatically infers the missing type annotations. Type inference makes programming easier and improves interoperability with Linked Data systems with different degrees of static type information.

For instance, in the example below the types of *Hamish* and *Dmitri* are unknown, so are assigned fresh type variables x and y . The constraints $x \leq \text{PERSON}$ and $y \leq \text{PERSON}$ are obtained by unfolding the tree of the following algorithmic type judgement.

$$\begin{array}{l} \text{knows} : \mathbf{p}(\text{PERSON}, \text{PERSON}), \\ \text{Hamish} : x, \text{Dmitri} : y \end{array} \Vdash (\text{Hamish knows Dmitri}) : \top$$

A type inference algorithm discovers the minimal unifier. The minimal unifier above is $x \mapsto \text{PERSON}, y \mapsto \text{PERSON}$. This unifier is a substitution, which gives an valid type judgement when applied to the above tree.

The general inference algorithm proceeds as follows. Firstly, apply algorithmic subtyping to obtain a proof tree indicating a set of subtype constraints over types involving fresh atomic types, as in the example above. Secondly, apply unification to the constraints, until either the constraints are rejected or the algorithm terminates successfully. If the constraints are rejected, there is no unifier. If the constraints are accepted, then the minimal unifier is generated [38].

Type inference appears implicitly in the RDFS standard. The rules of RDFS state that given a predicate, the type of the domain of the predicate bounds the type of any subject of the predicate. Similarly, the range of a predicate bounds the type of any object [14]. As demonstrated above, the same effect is achieved by type inference in this work. Type inference is performed at compile time, so incurs no cost to queries or updates.

7.2. An algebra for optimising updates. The calculus can be used as a basis for optimisation of queries. In particular, an algebra over updates can be derived. A suitable algebra allows updates to be rewritten whilst preserving the operational semantics of an update. An update before the algebra is applied should have the same operational behaviour as an update after the algebra is applied. For instance, the following two updates are expected to have the same operational behaviour.

$$\left(\begin{array}{c} \forall a. \\ \left(\begin{array}{c} \forall b. (b \text{ knows } a)^d \\ |(a \text{ type spy})| \\ \oplus \\ |(a \text{ type hermit})| \end{array} \right) \end{array} \right) \sim \left(\begin{array}{c} \forall a. \forall b. \\ \left(\begin{array}{c} (b \text{ knows } a)^d \\ |(a \text{ type spy})| \\ \oplus \\ (b \text{ knows } a)^d \\ |(a \text{ type hermit})| \end{array} \right) \end{array} \right)$$

The algebra can be applied to rewrite queries to normal forms. For instance, a suitable normal form can be used to optimise an update for distribution over a cluster of machines. The details of the algebra and proof that it preserves operational behaviour is worthy of detailed attention, so is the focus of a paper in preparation.

7.3. Extensions for interoperability of data. Interoperability problems similar to those for static types, tackled in Sec. 4, also apply to dynamic Linked Data. The intrinsic interoperability problem for URIs is known as co-reference [6]. Linked Data from different authorities may use different URIs to identify the same resource. For instance, many different URIs identify Robert Burns.

The problem results in an equivalence, or more generally a preorder, over URIs. The preorder, say \sqsubseteq , can be used to improve queries and updates, by permitting discrepancies between URIs. This light URI level reasoning can be built into the axioms. The following instance of an extended delete axiom allows a triple to be deleted by a weaker triple.

$$\frac{(Burns' \text{ author tribute}) \sqsubseteq (Burns \text{ creator tribute})}{(Burns' \text{ author tribute})^i \wp (Burns \text{ creator tribute})^d \triangleright \perp}$$

The above example works under the assumptions that $Burns = Burns'$ and $author \sqsubseteq creator$. This demonstrates one use of subProperty from the RDFS standard and sameAs from the OWL standard [31].

7.4. Extensions for future high level languages. Extensions can specify long term behaviour. A continuation process (indicated by then below) describes update which follows the current update. Recursion (indicated by ν below) allows a process to be unfolded. For instance, the following example unfolds a process twice to navigate a chain of knows predicates until it reaches a leader.

$$\begin{array}{l} \nu Y(Ali). \\ \left(\begin{array}{l} \vee b. \\ |(Ali \text{ knows } b)| \text{ then } Y[b] \\ \oplus \\ |(Ali \text{ type leader})| \text{ then } P \end{array} \right), \rightarrow \begin{array}{l} (Ali \text{ knows } Bob), \\ (Bob \text{ type leader}), \\ P\{^{Bob/Ali}\} \end{array} \\ (Ali \text{ knows } Bob), \\ (Bob \text{ type leader}) \end{array}$$

The above example, Upon reaching the leader performs some process P using the URI of the leader. Recursion is useful for describing long term behaviour of an automata. An automata can control navigation in a site dependant on Linked Data.

In another paradigm the right hand side of a commitment relation gathers a monoid of query results. Results are then manipulated in a functional language. The suggestion is that the standards are approaching a high level language for engaging with Linked Data.

8. CONCLUSIONS

The calculus presented in this paper has a distinguished role in the emerging Web of Linked Data. The calculus provides a level of abstraction where atomic actions can be executed using only local observations about data. The level of abstraction is above concerns such as the messages sent between machines, but below concerns such as the global structure of a store.

The calculus models substantial queries and updates which span diverse data sources. It is assumed that mechanisms can aggregate Linked Data in one environment — typically a store. The

contribution is an operational semantics for standards which operate at this level of abstraction. In particular, an operational semantics for SPARQL Query and SPARQL Update are specified. Examples focus on SPARQL Update which is crucial for Read–Write Linked Data.

The calculus works with a light type system, which can be checked locally. This contrasts to stronger schema, which require a global perspective on data. The calculus demonstrates that appropriate data can be lifted to a feasible type system. Literals in a query or update can always be type checked. By providing type assumptions, URIs can also be type checked. The contribution is an investigation into the RDFS standard as a type system.

The calculus simultaneously models key Linked Data standards. The challenge is that collectively the standards do not form an established logical system. Classical logic is found in constraints, linear logic is found in atomic actions and a truly concurrent process algebra combines actions, Sec. 2. A partial order over URIs and types enables interoperability, Sec. 4. The type system introduces two intuitionistic logics, Sec. 5, 6. Kleene algebras found optimisations, Sec. 7. Related work includes models based on graphs, epistemic logics and bunched implications [47, 8, 45]. The suggestion is that any model for Linked Data accommodates several notions of truth.

Acknowledgements. We appreciate our stimulating discussions with Mariangiola Dezani-Ciancaglini on this proposal.

REFERENCES

- [1] Martín Abadi and Cédric Fournet. Mobile values, new names, and secure communication. In *Proceedings of the 28th ACM SIGPLAN-SIGACT*, volume 36, pages 104–115, London, January 17-19 2001. ACM, NY.
- [2] Serge Abiteboul. Querying semi-structured data. In Foto Afrati and Phokion Kolaitis, editors, *Database Theory – International Conference on Database Theory*, pages 1–18, Delphi, Greece, January 8-10 1997. Springer Berlin/Heidelberg.
- [3] Samson Abramsky. What are the fundamental structures of concurrency?: We still don’t know! In *Proceedings of the Workshop Essays on Algebraic Process Calculi*, volume 162, pages 37–41. Elsevier, September 29 2006.
- [4] Samson Abramsky, Simon Gay, and Rajagopal Nagarajan. Interaction categories and the foundations of typed concurrent programming. In Manfred Broy, editor, *Proceedings of the NATO Advanced Study Institute on Deductive Program Design*, pages 35–114, Marktobendorf, Germany, July 26 - August 7 1995. Springer, Berlin/Heidelberg.
- [5] Samson Abramsky and Steven Vickers. Quantales, observational logic and process semantics. *Mathematical Structures in Computer Science*, 3(02):161–227, 1993.
- [6] Harith Alani et al. Managing reference: Ensuring referential integrity of ontologies for the Semantic Web. In Asunción Gómez-Pérez and V. Benjamins, editors, *Knowledge Engineering and Knowledge Management: Ontologies and the Semantic Web*, volume 2473, pages 235–246. Springer, Berlin/Heidelberg, 2002.
- [7] Jean-Franois Baget. RDF entailment as a graph homomorphism. In Yolanda Gil, Enrico Motta, V. Benjamins, and Mark Musen, editors, *The Semantic Web – ISWC 2005*, volume 3729, pages 82–96, Galway, Ireland, November 2005 2005. Springer Berlin/Heidelberg.
- [8] Alexandru Baltag, Bob Coecke, and Mehrnoosh Sadrzadeh. Epistemic Actions as Resources. *Journal of Logic and Computation*, 17(3):555–585, 2007.
- [9] Tim Berners-Lee, Dan Connolly, Lalana Kagal, Yosi Scharf, and Jim Hendler. N3logic: A logical framework for the world wide web. *Theory and Practice of Logic Programming*, 8(03):249–269, 2008.
- [10] Paul V. Biron and Ashok Malhotra. *XML Schema part 2: Datatypes Second Edition*. W3C, MIT, Cambridge, MA, 2004.
- [11] Christian Bizer. The emerging Web of Linked Data. *IEEE Intelligent Systems*, 24:87–92, 2009.
- [12] Christian Bizer et al. DBpedia: A crystallization point for the Web of Data. *Web Semantics: Science, Services and Agents on the World Wide Web*, 7(3):154–165, 2009.
- [13] Laura Bocchi and Roberto Lucchi. Atomic commit and negotiation in service oriented computing. In Paolo Ciancarini and Herbert Wiklicky, editors, *Coordination Models and Languages*. Springer, Berlin/Heidelberg, 2006.
- [14] Dan Brickley and R.V. Guha. *RDF Vocabulary Description Language 1.0: RDF Schema*. W3C, MIT, Cambridge, MA, 2004.

- [15] Allen Brown, Matthew Fuchs, Jonathan Robie, and Philip Wadler. MSL: A model for W3C XML Schema. *Computer Networks*, 39(5):507–521, 2002.
- [16] Maria Buscemi and Ugo Montanari. CC-Pi: A constraint-based language for specifying service level agreements. In Rocco De Nicola, editor, *Programming Languages and Systems*, volume 4421 of *Lecture Notes in Computer Science*, pages 18–32. Springer, Berlin/Heidelberg, 2007.
- [17] Marco Carbone, Kohei Honda, and Nobuko Yoshida. Structured communication-centred programming for Web Services. In Rocco De Nicola, editor, *Programming Languages and Systems*, volume 4421 of *Lecture Notes in Computer Science*, pages 2–17. Springer, Berlin/Heidelberg, 2007.
- [18] Jeremy J. Carroll, Christian Bizer, Pat Hayes, and Patrick Stickler. Named graphs. *Web Semantics: Science, Services and Agents on the World Wide Web*, 3(4):247–267, 2005.
- [19] John Horton Conway. *Regular Algebra and Finite Machines*. Chapman and Hall, London, 1971.
- [20] Thierry Coquand and Gerard Huet. The calculus of constructions. *Inf. Comput.*, 76:95–120, February 1988.
- [21] Richard Cyganiak. *A relational algebra for SPARQL*. Hewlett Packard Labs, Bristol, 2005.
- [22] Ronald Fagin, Phokion Kolaitis, Rene Miller, and Lucian Popa. Data exchange: Semantics and query answering. In Diego Calvanese, Maurizio Lenzerini, and Rajevev Motwani, editors, *Database Theory – International Conference on Database Theory*, volume 2572, pages 207–224, Siena, Italy, January 8-10 2003. Springer Berlin/Heidelberg.
- [23] Roy T. Fielding and Richard N. Taylor. Principled design of the modern Web architecture. *ACM Transactions on Internet Technology*, 2(2):115–150, 2002.
- [24] Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50(1):1–112, 1987.
- [25] Jean-Yves Girard. Truth, modality and intersubjectivity. *Mathematical Structures in Computer Science*, 17:1153–1167, December 2007.
- [26] J. Gregorio and B. de hOra. *The Atom Publishing Protocol*. Internet Engineering Task Force, Fremont, CA, October 2007.
- [27] Steve Harris, Andy Seaborne, and Eric Prud’hommeaux. *SPARQL 1.1 Query Language*. W3C, MIT, Cambridge, MA, October 2010.
- [28] Olaf Hartig et al. Executing SPARQL Queries over the Web of Linked Data. In A. Bernstein et al., editors, *The Semantic Web – ISWC 2009*, volume 5823, pages 293–309, Chantilly, VA, October 25-29 2009. Springer Berlin/Heidelberg.
- [29] C. Hoare, Bernhard Möller, Georg Struth, and Ian Wehrman. Concurrent Kleene algebra. In Mario Bravetti and Gianluigi Zavattaro, editors, *CONCUR 2009*, volume 5710, pages 399–414, Bologna, Italy, September 1-14 2009. Springer Berlin/Heidelberg.
- [30] J. S. Hodas and D. Miller. Logic programming in a fragment of intuitionistic linear logic. *Information and Computation*, 110(2):327–365, 1994.
- [31] Ian Horrocks and Peter Patel-Schneider. Reducing OWL entailment to description logic satisfiability. *Web Semantics: Science, Services and Agents on the World Wide Web*, 1(4):345–357, 2004.
- [32] Alan Jeffrey and Julian Rathke. Contextual equivalence for higher-order π -calculus revisited. *Logical Methods in Computer Science*, 1(1):paper 4, 2005.
- [33] Jon M. Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM*, 46(5):604–632, 1999.
- [34] Graham Klyne and Jeremy Carroll. *Resource Description Framework: Concepts and Abstract Syntax*. W3C, MIT, Cambridge, MA, 2004.
- [35] Georgi Kobilarov et al. Media meets Semantic Web: How the BBC uses DBpedia and Linked Data to make connections. In Lora Aroyo, editor, *The Semantic Web: Research and Applications. 6th European Semantic Web Conference*, pages 723–737, Heraklion, Greece, May 31 - June 4 2009. Springer, Berlin/Heidelberg.
- [36] Dexter Kozen. Kleene algebra with tests. *ACM Trans. Program. Lang. Syst.*, 19:427–443, May 1997.
- [37] Holbrook Mann MacNeille. Extensions of partially ordered sets. *Proceedings of the National Academy of Sciences of the United States of America*, 22(1):45–50, 1936.
- [38] Alberto Martelli and Ugo Montanari. An efficient unification algorithm. *ACM Transactions on Programming Languages and Systems*, 4(2):258–282, 1982.
- [39] Robin Milner. *A calculus of communicating systems*, volume 92. Springer, NJ, 1980.
- [40] Robin Milner. The polyadic π -calculus: A tutorial. In F.L. Bauer, W. Brauer, and H. Schwichttenburg, editors, *Logic and Algebra in Specification*. Springer, New York, 1993.
- [41] John C. Mitchell. Type inference with simple subtypes. *Journal of Functional Programming*, 1(03):245–285, 1991.
- [42] S. Muñoz, J. Pérez, and C. Gutierrez. Minimal deductive systems for RDF. In *The Semantic Web: Research and Applications, 4th European Semantic Web Conference*, pages 53–67, Innsbruck, Austria, June 3-7 2007. Springer.

- [43] Mark Nottingham and Robert Sayre. *The Atom Syndication Format*. Internet Engineering Task Force, Fremont, CA, December 2005.
- [44] Chimezie Ogbuji. *SPARQL 1.1 Uniform HTTP Protocol for Managing RDF Graphs*. W3C, MIT, Cambridge, MA, October 2010.
- [45] Peter W. O’Hearn and David J. Pym. The logic of bunched implications. *Bulletin of Symbolic Logic*, 5(2):215–244, 1999.
- [46] Tope Omitola et al. Put in your postcode, out comes the data: A case study. In Lora Aroyo et al., editors, *The Semantic Web: Research and Applications. 7th Extended Semantic Web Conference*, volume 6088, pages 318–332, Heraklion, Crete, May 30 - June 3 2010. Springer Berlin/Heidelberg.
- [47] Jorge Pérez, Marcelo Arenas, and Claudio Gutierrez. Semantics and complexity of SPARQL. *ACM Transactions on Database Systems*, 34(3):1–45, 2009.
- [48] Eric Prud’hommeaux and Andy Seaborne. *SPARQL Query Language for RDF*. W3C, MIT, Cambridge, MA, 2008.
- [49] John Reynolds. Toward a grainless semantics for shared-variable concurrency. In Kamal Lodaya and Meena Mahajan, editors, *Foundations of Software Technology and Theoretical Computer Science*, volume 3328, pages 11–38. Springer, Berlin/Heidelberg, 2005.
- [50] Vijay A. Saraswat, Martin Rinard, and Prakash Panangaden. The semantic foundations of concurrent constraint programming. In *Proceedings of the 18th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 333–352, Orlando, Florida, January 21-23 1991. ACM, NY.
- [51] Robert Sayre. Atom: the standard in syndication. *IEEE Internet computing*, 9(4):71–78, 2005.
- [52] Simon Schenk and Paul Gearon. *SPARQL 1.1 Update*. W3C, MIT, Cambridge, MA, October 2010.
- [53] Andy Seaborne and Geetha Manjunath. *SPARQL/Update: A language for updating RDF graphs*. Hewlit Packard Labs, Bristol, 2007.