

A Verified Algebra for Linked Data

Ross Horne and Vladimiro Sassone

Electronics and Computer Science, University of Southampton, United Kingdom
{rjh06r,vs}@ecs.soton.ac.uk

A foundation is investigated for the application of loosely structured data on the Web. This area is often referred to as Linked Data, due to the use of URIs in data to establish links. This work focuses on emerging W3C standards which specify query languages for Linked Data. The approach is to provide an abstract syntax to capture Linked Data structures and queries, which are then internalised in a process calculus. An operational semantics for the calculus specifies how queries, data and processes interact. A labelled transition system is shown to be sound with respect to the operational semantics. Bisimulation over the labelled transition system is used to verify an algebra over queries. The derived algebra is a contribution to the application domain. For instance, the algebra may be used to rewrite a query to optimise its distribution across a cluster of servers. The framework used to provide the operational semantics is powerful enough to model related calculi for the Web.

1 Introduction

The application of interest is a powerful emerging idea commonly referred to as the Web of Data [6]. The Web of Data marks a shift from publishing documents to publishing data. The Web is based on documents which contain links to other documents. The Web of Data is concerned with resources more general than documents. Data on the Web contains links to resources described in multiple data sources. In both the case of the Web and the Web of Data the links between documents and resources, respectively, are established by a standardised global naming system --- the URI. On the Web, URIs allow documents in distributed locations with distinct ownership to refer to each other. Similarly, in a Web of Data, URIs allow data in distributed locations with distinct ownership to refer to common resources.

Suppose that the URIs are not used as a standard naming system. In this case, each data source uses its own naming system. Typically, in this case each data source is disjoint, hence traditional database techniques may be applied. This is referred to as closed world system, since the boundaries of the data source are known. For instance, classical negation can be used to determine whether some data does not appear in a data source, and schemata can constrain the structure of data.

In contrast, the presence of URIs as a global naming system, enables an open world system. In an open world system a variety of protocols can be used to obtain data from multiple sources based on the URIs which appear. For instance, a request may be sent to a URI to directly obtain some data about that URI. Alternatively, services may be used to find data relevant to a URI. In this open world setting, there is no guarantee that mechanisms find all relevant data. There may always be data not known locally which refers to a resource; hence in general optimal query results cannot be obtained and classical negation cannot be applied. Another restriction in an open world system is that schemata which constrain data cannot be enforced globally.

A light semi-structured data format must be agreed for the Web of Data. The W3C recommends the Resource Description Framework (RDF) as a general format for presenting data [16]. RDF is based on triples which consist of a subject, predicate and object. The subject, predicate and object are all named by URIs. Each URI in a triple may represent resources in different locations, hence a triple links locations.

Other semi-structured data formats contain URIs, such as feeds. RDF is intended as a minimal data format to which other formats can be lifted.

Assuming that Linked Data can be gathered, observations about Linked Data can be made. The W3C recommendation is to use SPARQL Queries to make such observations [24]. In this work, to model this scenario, both RDF Data and SPARQL Queries are internalised in a process calculus. The operational semantics of the process calculus specifies how queries and data interact, to realise the W3C recommendations. The operational semantics are realistic since there is no guarantee of maximal responses, only that responses are correct.

Two SPARQL Queries may be indistinguishable with respect to their operational behaviour. Such operationally equivalent queries are bisimilar. In this work, bisimulation is used to derive an algebra over SPARQL Queries. The algebra agrees with expected equivalences analogous to those uncovered by relational algebra and exposes some new equivalences. The derived algebra can be used to rewrite a query to a normal form. Normal forms are useful for optimisation purposes. A query can be optimised before being distributed over multiple data sources. Distribution of queries is a key challenge for enabling a Web of Data [11].

Section 2 presents a syntax and semantics for RDF triples, SPARQL queries and processes which internalise both triples and queries. Section 3 provides an alternative operational semantics using a labelled transition system. The labelled transition system is proven to be sound with respect to the reduction system. Section 4 introduces two notions of equivalence over the calculus, which correspond to the two operational semantics. Bisimulation for the labelled transition system is proven to be complete with respect to contextual equivalence for the reduction system. An algebra for queries is verified using bisimulation.

2 A syntax and semantics for the syndication calculus

The concrete syntax for both RDF and SPARQL Query are specified in W3C recommendations [16, 24]. Here an abstract syntax is presented to model the core features of the concrete syntax. This abstract syntax is easier to define than the concrete syntax, which is sugared to make programming easier.

The operational semantics of the calculus is specified as a reduction system. The syntax and rules of the reduction system borrow from a fragment of Linear Logic, extended with a continuation. Related work has investigated other approaches to using Linear Logic for both query languages and process calculi [17, 13, 3].

Note that the description of the syntax and reduction system is brief. A similar syntax and reduction system are extensively discussed in the thesis of the first author [14]. The main contribution of this paper is the bisimulation results for queries.

2.1 A syntax for RDF triples

An abstract syntax for triples conveys the RDF data format. The atoms of the syntax are names and literals. Names represent occurrences of URIs, which are represented by identifiers in italics, such as *John* or *knows*. Literals are basic data values, such as the strings ‘Paul’ or ‘77-3426’. The definition of literals in the XML Schema Datatypes specification [4] is assumed. Variables a, b, \dots and x, y, \dots represent place holders for names and literals respectively.

A triple consists of three components: the subject, the predicate and the object, which is written (*subject predicate object*). The subject is related by the predicate to the object, similarly to simple sen-

	$U ::= C$		
$\phi ::= I$	ϕ	asked triple	
0	$U \oplus U$	filter	
$\phi \vee \phi$	$U \otimes U$	choice	$P ::= \perp$
$\phi \wedge \phi$	$\bigvee a.U$	tensor	$P \wp P$
$\neg \phi$	$\bigvee x.U$	select name	$\bigwedge a.P$
...	$*U$	select literal	U
true	$U;P$	iteration	\underline{C}
false		then	nothing
or			par
and			blank node
not			query
etc.			stored triple

Figure 1: The syntax of constraints (ϕ), queries (U) and processes (P), over triples (C).

tences in English of form subject-verb-object, where URIs and literals are used instead of words. The syntax ensures that literals can only appear as the object of a triple. The example below presents two RDF triples.

(b_4 *home* *starr.uk*) (b_4 *give_name* 'Ringo')

Predicates are names such as *home*. For instance, the first triple above means that a subject b_4 is related by predicate *home* to object *starr.uk*. The second triple above indicates that subject b_4 is related by predicate *given_name* to the literal 'Ringo'.

2.2 A syntax for SPARQL queries

In this section an abstract syntax for queries, Fig. 1, represents the core features of SPARQL Query [24]. SPARQL Queries are used to read from RDF triples. Synchronisation constructs allow substantial queries to be expressed. The syntax of processes, also in Fig. 1, demonstrates how both queries and content can be internalised in a process calculus, which suggests a high level language for Linked Data, which uses query results. In this model, persistently stored triples are used to answer queries. A stored triple is indicated by an underscore.

Ask queries and multiplicative operators. The simplest 'ask' query provides a triple to be matched. There are three multiplicative operators: a tensor product (\otimes) for synchronously joining queries, a par operator (\wp) for composing processes in parallel and the operator then ($;$) for guarding a process with a query. The difference between tensor and par is that queries composed using tensor must happen simultaneously (in the same atomic step), whereas processes composed in parallel may be used in different atomic steps. Tensor is the implicit join of queries used in SPARQL. Then and par are part of a higher level language, where query results are immediately used. These operators are multiplicative since they control the sharing of resources.

The additive operators and select queries. There are three additive operators: choose (\oplus), select (\bigvee) and the blank node quantifier (\bigwedge). The choose operator presents a choice between two queries, hence models the SPARQL keyword UNION. The select operator is a quantifier which binds a variable. Select is used to model SELECT queries in SPARQL, which discover names and literals. The names and literals discovered can also be bound in a continuation process, hence value passing is modelled at a high level. Blank node quantifiers provide a model for blank nodes in RDF [16]. A blank node is a local name

$$\begin{aligned}
P \wp \perp &\equiv P & P \wp Q &\equiv Q \wp P & P \wp (Q \wp R) &\equiv (P \wp Q) \wp R \\
\bigwedge a. \perp &\equiv \perp & \bigwedge a. \bigwedge b. P &\equiv \bigwedge b. \bigwedge a. P & \bigwedge a. P \wp Q &\equiv \bigwedge a. (P \wp Q) \quad a \notin \text{fn}(P)
\end{aligned}$$

Figure 2: The structural congruence over processes.

where the scope of the blank node is indicated by the scope of the quantifier. Blank nodes allow further data structures to be represented in RDF, including XML.

Constraints and optional queries. A constraint may be used in a query. Constraints form a Boolean algebra of basic predicates, such as inequalities and regular expressions. The specification of constraints can be found under the keyword `FILTER` in the recommendation [24]. A choice between a query and true models an optional query in SPARQL, so the keyword `OPTIONAL` is defined as follows: $\text{OPTIONAL } U \triangleq U \oplus \mathbf{I}$.

Repeated queries and iteration. A common requirement of a query language is that more than one result can be obtained. Bounded multiple copies of queries can be synchronously posed, using queries with natural number exponents and finite sums. Exponents and sums are just abbreviations defined as follows.

$$U^0 \triangleq \mathbf{I} \quad U^{n+1} \triangleq U \otimes U^n \quad \sum_{n=0}^k U^n \triangleq \mathbf{I} \quad \sum_{n=0}^{k+1} U^n \triangleq \sum_{n=0}^k U^n \oplus U^{k+1}$$

A natural number exponent n repeatedly applies the tensor product, so the query must be answered exactly n times. The sum with bound n allows the query to be answered between 0 and n times. Sums model the keyword `LIMIT`, such that $U \text{ LIMIT } k \triangleq \sum_{n=0}^k U^n$.

Unbounded iteration of queries is indicated by an explicit operator $(*)$, which allows zero or more copies of a query to be answered. Note that iteration differs from replication in common process calculi. All copies of an iterated query must be answered simultaneously using disjoint resources.

2.3 A reduction system for the calculus

The reduction system presents a concise operational semantics for the calculus. The reduction system is defined by a structural congruence and a relation over processes called the commitment relation. A further preorder over triples formalises key features of RDF Schema (RDFS [7]). RDFS is a light extension to RDF, which improves interoperability by resolving aliases between URIs.

The structural congruence (\equiv in Fig. 2) is defined such that (P, \wp, \perp) forms a commutative monoid. Alpha conversion can also be applied to blank node quantifiers. Furthermore, blank node quantifiers can be eliminated in the presence of nothing, commute and distribute over par. All reductions are considered up to structural congruence --- as standard in process calculi.

The commitment relation (\triangleright in Fig. 3) specifies atomic operational steps. The process on the left of the commitment relation, becomes the process on the right. A commitment is performed atomically.

Working with aliases for URIs is a key problem in Linked Data [2]. Aliases arise since different data sources use different URIs for similar purposes. For instance, in the context of a song, predicate *lyricist* may be more specific than predicate *creator* (see `subPropertyOf` in RDFS [7]). Similarly, *song₀* and *song₁* may be URIs for the same song (see `sameAs` in OWL [2]). Hence the aliases $\textit{lyricist} \sqsubseteq \textit{creator}$

$$\begin{array}{c}
\frac{C \sqsubseteq D}{\underline{C} \wp D \triangleright \underline{C}} \quad \frac{\models \phi}{\phi \triangleright \perp} \quad \frac{P \wp U \triangleright Q}{P \wp (U \oplus V) \triangleright Q} \quad \frac{P \wp V \triangleright Q}{P \wp (U \oplus V) \triangleright Q} \quad \frac{P \wp U \triangleright P' \quad Q \wp V \triangleright Q'}{P \wp Q \wp (U \otimes V) \triangleright P' \wp Q'} \\
*U \triangleright \perp \quad \frac{P \wp U \triangleright Q}{P \wp *U \triangleright Q} \quad \frac{P \wp (*U \otimes *U) \triangleright Q}{P \wp *U \triangleright Q} \quad \frac{P \wp U \{^b/a\} \triangleright Q}{P \wp \sqrt{a}.U \triangleright Q} \quad \frac{P \wp U \{^v/x\} \triangleright Q}{P \wp \sqrt{x}.U \triangleright Q} \\
\frac{P \wp U \triangleright Q}{P \wp (U ; R) \triangleright Q \wp R} \quad \frac{P \triangleright P'}{P \wp Q \triangleright P' \wp Q} \quad \frac{P \wp Q \triangleright P' \wp Q'}{P \wp \wedge a.Q \triangleright P' \wp \wedge a.Q'} \quad a \notin \text{fn}(P, P', \beta)
\end{array}$$

Figure 3: Commitment rules: ask, filter, choose left, choose right, tensor, weakening, dereliction, contraction, select name, select literal, guard, context, and blank node (fn indicates the free names).

and $\text{song}_0 \sqsubseteq \text{song}_1$ may be assumed. The application specific set of alias assumptions is referred to as β . The transitive reflexive closure of β gives rise to a preorder (\sqsubseteq) over URIs.

The ask axiom, guard rule and alias assumptions. The following example demonstrates the interaction of an ask query with a continuation and a stored triple. The axiom ‘ask’ allows a query triple and a stored triple to interact. The stored triple remains available after the commitment. The axiom ‘guard’ makes the continuation process available after the commitment.

$$(\underline{\text{song}_0 \text{ lyricist } b_4}) \wp ((\text{song}_1 \text{ creator } b_4); P) \triangleright (\underline{\text{song}_0 \text{ lyricist } b_4}) \wp P$$

Above, the conditions for a match are relaxed by the preorder over triples (\sqsubseteq). The preorder over triples is the point-wise extension of the preorder over URIs introduced above.

The tensor and select rules. The following example demonstrates two synchronised queries, in the presence of two stored triples. The first query poses a pattern to match, while the second query selects a name with respect to a pattern.

$$\frac{((b_2 \text{ role singer}) \otimes \sqrt{b}.((b \text{ role guitarist}); P)) \wp (\underline{b_2 \text{ role singer}}) \wp (\underline{b_3 \text{ role guitarist}})}{(\underline{b_2 \text{ role singer}}) \wp (\underline{b_3 \text{ role guitarist}})} \triangleright \frac{(\underline{b_2 \text{ role singer}}) \wp (\underline{b_3 \text{ role guitarist}}) \wp P \{^{b_3/b}\}}{(\underline{b_3 \text{ role guitarist}}) \wp P \{^{b_3/b}\}}$$

In the above example, the ‘tensor’ rule divides the stored triples between the two parts of the query. On the left the ‘select’ rule is applied. The ‘select’ rule substitutes a suitable URI for the quantified name. The result is that a URI is passed to the continuation.

The choose rule. The following example demonstrates a choice between queries. The ‘choose left’ rule is used in this case.

$$\sqrt{a}.(((a \text{ knows } b_2); P) \oplus ((b_2 \text{ knows } a); Q)) \wp (\underline{b_1 \text{ knows } b_2}) \triangleright (\underline{b_1 \text{ knows } b_2}) \wp P \{^{b_1/a}\}$$

The query result determines the continuation triggered.

Constraints in queries. The example query below selects a literal. The data literal appears in a triple and a constraint. The rules ensure that both a suitable triple appears and the constraint imposed holds.

$$\bigvee x.((|x| \leq 5) \otimes (b_1 \text{ name } x); P) \wp \underline{(b_1 \text{ name 'John'})} \triangleright \underline{(b_1 \text{ name 'John'})} \wp P\{^{\text{John}}/x\}$$

The satisfaction relation for evaluating constraints \models , is left to the W3C recommendation [24]. Satisfaction is assumed to define a Boolean algebra of constraints.

The rules for iteration of queries. The example below demonstrates iteration used to answer two copies of the same query. Two iterated queries are answered using ‘dereliction’, which are combined using the conventional tensor rule. The ‘contraction’ rule then reduces the combined queries to a single query.

$$*\bigvee c.((c \text{ is busy}); P) \wp \underline{(b_2 \text{ is busy})} \wp \underline{(b_3 \text{ is busy})} \triangleright \underline{(b_2 \text{ is busy})} \wp \underline{(b_3 \text{ is busy})} \wp P\{^{b_2}/c\} \wp P\{^{b_3}/c\}$$

A continuation for each result is triggered. Note the ‘weakening’ rule could be used to allow the query to be answered zero times.

Blank nodes as quantifiers. The example below demonstrates a query which discovers a blank node. The ‘blank node’ rule uses a temporary name to represent the blank node. The result is that the scope of the blank node quantifier is extended to include the continuation, which receives the blank node.

$$\bigvee c.((c \text{ creator } b_2); U) \wp \bigwedge a.(\underline{(a \text{ author } b_2)} \wp \underline{(a \text{ status open})}) \triangleright \bigwedge a. \left(\underline{U\{^a/c\}} \wp \underline{(a \text{ author } b_2)} \wp \underline{(a \text{ status open})} \right)$$

The alias $\text{author} \sqsubseteq \text{creator}$ is assumed above. The temporary name must not appear in the alias assumptions (β). The unused stored triple is idled.

Rules for an additive disjunction, tensor product, existential quantification, universal quantification and iteration, are borrowed from Linear Logic [9]. The sequent calculus is extended to indicate a continuation process, constraints extend the basic units with a Boolean algebra, and a preorder accommodates aliases over names.

3 A labelled transition system for the operational semantics

The operational semantics can be expressed as a labelled transition system. This provides an alternative operational semantics to the reduction system. This alternative semantics allows the behaviour of queries and data to be evaluated separately and then composed. Lemma 2 verifies that the labelled transition system and reduction system describe the same behaviour.

3.1 The purpose of labels

A labelled transition consists of two processes and a label. The first process is the process before the transition. The label is a constraint on the context in which a transition can take place. The second process is the resulting process after the transition.

The labels are formed from a commutative monoid over triples (E, \otimes, I) . A label indicates the inputs and outputs of a process. An input indicates that a process can proceed if it can receive the triples on the

$$\begin{array}{c}
\frac{C \sqsubseteq D}{D \xrightarrow{C} \perp} \quad \frac{U \xrightarrow{E} Q}{U; P \xrightarrow{E} Q \wp P} \quad \frac{U \xrightarrow{E} P \quad V \xrightarrow{F} Q}{U \otimes V \xrightarrow{E \otimes F} P \wp Q} \quad \frac{U \xrightarrow{E} P}{U \oplus V \xrightarrow{E} P} \quad \frac{V \xrightarrow{E} Q}{U \oplus V \xrightarrow{E} Q} \\
\frac{\varepsilon \phi}{\phi \xrightarrow{I} \perp} \quad \frac{U\{b/a\} \xrightarrow{E} Q}{\forall a. U \xrightarrow{E} Q} \quad \frac{U\{v/x\} \xrightarrow{E} Q}{\forall x. U \xrightarrow{E} Q} \quad *U \xrightarrow{I} \perp \quad \frac{U \xrightarrow{E} P}{*U \xrightarrow{E} P} \quad \frac{*U \otimes *U \xrightarrow{E} P}{*U \xrightarrow{E} P}
\end{array}$$

Figure 4: Labelled transitions for queries: input triple, trigger guard, tensor, choose left, choose right, filter, select name, select literal, weakening, dereliction and contraction.

label from its context. An output indicates that a process outputs the triple on the label to its context. For instance, the query below inputs a triple; while the stored triple below outputs a triple.

$$(b_4 \text{ knows } b_3); P \xrightarrow{(b_4 \text{ knows } b_3)} P \quad \underline{(b_4 \text{ knows } b_3)} \xrightarrow{\overline{(b_4 \text{ knows } b_3)}} \underline{(b_4 \text{ knows } b_3)}$$

A relevant interpretation is that the first transition above is an action from the perspective of a client which resolves a query; whereas the second is an action from the perspective of a server that provides a triple. Two processes composed in parallel with matching inputs and outputs may interact. For instance, the above processes can be composed, resulting in the following transition. The unit label indicates an operational step without side effects.

$$(b_4 \text{ knows } b_3); P \wp \underline{(b_4 \text{ knows } b_3)} \xrightarrow{I} P \wp \underline{(b_4 \text{ knows } b_3)}$$

Output labels can also indicate extruded names. For instance, the example below extrudes the name a . The extruded names represent blank nodes where the scope of the blank node quantifier may be extended. This is similar to extrusion of new names in the π -calculus [22].

$$\bigwedge a. \underline{(a \text{ has paper})} \wp \underline{(b_2 \text{ has stone})} \xrightarrow{\overline{a(a \text{ has paper})}} \underline{(a \text{ has paper})} \wp \underline{(b_2 \text{ has stone})}$$

The commutative monoid rules can always be applied to reorder labels.

3.2 Labelled transitions for queries

The input transitions allow the behaviour of a query to be modelled independently. The rules for queries are presented in Fig. 4. The rules accumulate RDF triples on an input label, which represents contexts in which a query may be answered.

The ‘input triple’ rule poses the triple as an input on the label. The triple on the label may be strengthened by the preorder over triples. The ‘trigger guard’ rule allows a continuation process to be triggered exposing the continuation. The following example demonstrates a query consisting of a single triple and a continuation process, where the preorder $colleague \sqsubseteq knows$ is assumed.

$$(b_4 \text{ knows } b_3); P \xrightarrow{(b_4 \text{ colleague } b_3)} P$$

Select quantifiers are resolved by anticipating the name or literal to input. For instance, the following labelled transition indicates that the query can be answered in a context where a name is chosen. The same name is passed to the continuation process.

$$\bigvee a. ((b_4 \text{ knows } a); P) \xrightarrow{(b_4 \text{ knows } b_3)} P\{b_3/a\}$$

$$\begin{array}{c}
\frac{C \sqsubseteq D}{\underline{C} \xrightarrow{\underline{D}} \underline{C}} \quad \frac{P \xrightarrow{\alpha \bar{E}} Q}{\wedge a.P \xrightarrow{\alpha + a \bar{E}} Q} \quad a \notin \text{fn}(\beta) \quad \frac{P \xrightarrow{\alpha \bar{E}} Q}{\wedge a.P \xrightarrow{\alpha \bar{E}} \wedge a.Q} \quad a \notin \alpha \cup \text{fn}(E) \\
\\
\frac{P \xrightarrow{\alpha \bar{E}} P'}{P \wp Q \xrightarrow{\alpha \bar{E}} P' \wp Q} \quad \alpha \cap \text{fn}(Q) = \emptyset \quad \frac{P \xrightarrow{\alpha_0 \bar{E}} P' \quad Q \xrightarrow{\alpha_1 \bar{F}} Q'}{P \wp Q \xrightarrow{\alpha_0 + \alpha_1 \overline{E \otimes F}} P' \wp Q'} \quad \alpha_0 \cap \text{fn}(Q) = \emptyset \\
\alpha_1 \cap \text{fn}(P) = \emptyset \\
\\
\frac{P \xrightarrow{E \otimes F} P' \quad Q \xrightarrow{\alpha \bar{F}} Q'}{P \wp Q \xrightarrow{E} \wedge \alpha.(P' \wp Q')} \quad \alpha \cap (\text{fn}(P) \cup \text{fn}(E)) = \emptyset
\end{array}$$

Figure 5: Process rules: output triple, open, blank node context, par context, parallel outputs and close. The symmetric versions of the par context and close rule are also assumed.

Choices are resolved by anticipating the left or right branch. For instance, the following transition indicates the label and continuation which results from choosing the left branch.

$$((b_4 \text{ knows } b_2); P) \oplus ((b_4 \text{ knows } b_3); Q) \xrightarrow{(b_4 \text{ knows } b_2)} P$$

Tensor synchronises two queries, by composing their respective labels and continuations. For instance, the following query simultaneously inputs two triples. The continuations of both queries are triggered in parallel, with the appropriate substitutions.

$$\forall a.(((b_4 \text{ knows } a); P) \otimes (\forall x.(a \text{ name } x); Q)) \xrightarrow{(b_4 \text{ knows } b_2) \otimes (b_2 \text{ name 'John'})} P\{b_2/a\} \wp Q\{b_2, \text{'John'}/a, x\}$$

A constraint is disposed when it is satisfied. For instance, in the following query the length of a selected literal is constrained, but satisfied by the substitution.

$$\forall x.((b_2 \text{ name } x) \otimes (|x| \leq 5); P) \xrightarrow{(b_2 \text{ name 'John'})} P\{\text{'John'}/x\}$$

Iteration anticipates the number of copies of a query to pose using weakening, dereliction and contraction. For instance, two copies of the following query are posed using contraction and dereliction. The label indicates the two separate triples which are to be answered simultaneously. Both continuations are composed in parallel.

$$*\forall a.((b_4 \text{ knows } a); P) \xrightarrow{(b_4 \text{ knows } b_2) \otimes (b_4 \text{ knows } b_3)} P\{b_2/a\} \wp P\{b_3/a\}$$

The rules of the labelled transition system are sufficient to model queries.

3.3 Labelled transitions for an RDF store

The behaviour of stored RDF triples can be modelled using output labels. The rules of output labels are presented in Fig. 5. The names extruded on the label are indicated by α , where $+$ indicates disjoint union of names. The abbreviation $\wedge \alpha.P$ is used to indicate the quantification of all names in α .

Stored triples can output the triple on the label. The same triple appears in the continuation unchanged. The preorder over names may be used to weaken the output triple. Names are extruded on

the label using the ‘open scope’ rule. For instance, the following triple outputs a triple and extrudes the blank node, using the assumption $colleague \sqsubseteq knows$.

$$\bigwedge b_4.(\underline{b_4 \text{ colleague } b_3}) \xrightarrow{b_4 | (b_4 \text{ knows } b_3)} \underline{b_4 \text{ colleague } b_3}$$

Output labels composed in parallel can be combined. Extruded names on both labels must be disjoint to preserve the scope of blank nodes. For instance, the following transition simultaneously outputs two triples and extrudes three names.

$$\bigwedge b_4.(\bigwedge b_2.(\underline{b_4 \text{ knows } b_2}) \wp \bigwedge b_3.(\underline{b_4 \text{ knows } b_3})) \xrightarrow{b_2, b_3, b_4 | (b_4 \text{ knows } b_2) \otimes (b_4 \text{ knows } b_3)} \underline{b_4 \text{ knows } b_2} \wp \underline{b_4 \text{ knows } b_3}$$

Two parallel processes may interact using the close rule. Close allows complementary inputs and outputs to be matched. Names extruded on the output label are introduced as quantifiers in the continuation. Any inputs not answered remain on the resulting label, to be answered later. For instance, the following iterated query is answered twice. One copy is answered by the available process and the other copy must be answered by the context for the transition to occur. In the continuation, the scope of the blank node is extended.

$$*\sqrt{a.((b_4 \text{ knows } a); P)} \wp \bigwedge b_3.(\underline{b_4 \text{ knows } b_3}) \xrightarrow{(b_4 \text{ knows } b_2)} \bigwedge b_3.(P\{b_2/a\} \wp P\{b_3/a\}) \wp \underline{b_4 \text{ knows } b_3}$$

The context rule for parallel composition allows a process which does not contribute to an interaction to idle. Similarly, the context rule for blank node quantifiers allows a blank node to be ignored in a transition if it does not appear on the label.

3.4 Comparison of the two operational semantics

To justify the labelled transition system, the labelled transitions are compared to the reductions of the reduction system. If a unit labelled transition can be derived then the corresponding reduction can also be derived. The significance is that, given the independent perspectives of the query and the store in terms of labelled transitions, their combination satisfies the global perspective specified by the reduction system.

Scope extrusion presents technical difficulties. The following technical lemma reduces these difficulties, by eliminating scope extrusion. The proof demonstrates that combinations of opening names and closing names can be eliminated from a proof tree which uses an extruded name.

Lemma 1 (Elimination of extrusion). *Suppose that a labelled transition proof uses name extrusion, but not in the conclusion. The same labelled transition, up to structural congruence, holds without any name extrusion.*

Note that full proofs for all theorems are provided in the thesis of the first author [14].

Every completed labelled transition can also be expressed as a reduction, Lemma 2. The proof works by transforming proof trees so that labels used in interactions are eliminated.

Lemma 2 (Elimination of labels). $P \xrightarrow{1} Q$ if and only if $P \triangleright Q$.

Thus the local perspective of the labelled transition system and the global perspective of the reduction system specify the same operational capabilities.

4 An algebra for the syndication calculus

In this section bisimulation is introduced as the natural notion of equivalence over the labelled transition system. Bisimulation is demonstrated to be sound with respect to equivalence in the reduction system. Thus every pair of bisimilar processes are equivalent with respect to the natural notion of equivalence over the reduction system. Bisimulation is then used to verify an algebra over queries and processes.

4.1 Bisimulation

Processes which are capable of the same observable behaviour can be regarded as equivalent. The observable behaviour of a process is given by the labels of the labelled transition system. Observational equivalence of processes is established using the technique of (strong) bisimulation, as follows.

Definition 1 (Bisimulation). *Bisimulation, written \sim , is the greatest symmetric relation such that the following holds, for any label l . If $P \sim Q$ and $P \xrightarrow{l} P'$ then there exists some Q' such that $Q \xrightarrow{l} Q'$ and $P' \sim Q'$.*

The following verifies that bisimulation is a congruence --- a relation which holds in any context. It is necessary that bisimulation is a congruence for it to be used as an algebra. A context is a process with a place holder for some syntax.

Lemma 3 (Bisimulation is a congruence). *If $P \sim Q$ and C is a context, then $CP \sim CQ$.*

An alternative notion of equivalence is defined using the reduction system. Contextual equivalence is used in related work to justify notions of bisimulation on the π -calculus and ambient calculus [15, 21].

Definition 2 (Contextual equivalence). *Contextual equivalence, written \simeq , is the greatest symmetric, reduction closed, context closed relation. A relation \mathcal{R} is reduction closed iff $P \mathcal{R} Q$ and $P \triangleright P'$ then there exists some Q' such that $Q \triangleright Q'$ and $P' \mathcal{R} Q'$. A relation \mathcal{R} is context closed iff $P \mathcal{R} Q$ yields that $CP \mathcal{R} CQ$, for all contexts C .*

Bisimulation is sound with respect to contextual equivalence. Soundness is essential to justify the chosen notion of bisimulation.

Theorem 1 (Bisimulation is a contextual equivalence). *If $P \sim Q$ then $P \simeq Q$.*

Proof. Reduction closure follows from Lemma 2 and context closure follows from Lemma 3. \square

Soundness of bisimulation ensures that algebraic properties proven using bisimulation also hold for contextual equivalence. Bisimulation simplifies proofs in the following section. Note that completeness (contextual equivalence is a bisimulation) is not required for this work. Completeness can only be achieved in an extended version of the calculus.

4.2 Algebraic properties of queries

Using bisimulation as an equivalence, key properties of queries are established. This section amounts to a soundness proof of the algebraic properties established. Thus if any two process are equivalent according to the algebraic properties then they are bisimilar; and furthermore, by Theorem 1, they are contextually equivalent.

For the labelled transition system, structural congruence is not assumed, hence verified here. The proof for the distributivity of blank node quantifiers over par requires extensive case analysis. The case of associativity of par follows from distributivity of blank node quantifiers. Proofs are similar to the analogous bisimulations in the π -calculus [22].

Proposition 2. *The structural congruence (Fig.2) is a bisimulation. So, (P, \wp, \perp) forms a commutative monoid. Blank node quantifiers annihilate with \perp , commute, and distribute over \wp .*

Bisimulation reveals some canonical algebraic properties of queries. Firstly, queries form an idempotent semiring. Semirings are ubiquitous in computer science. A notable feature of semirings is that the ideals of a semiring form a semiring.

Proposition 3. *$(U, \otimes, \oplus, I, 0)$ is a commutative idempotent semiring. That is, (U, \otimes, I) is a commutative monoid, $(U, \oplus, 0)$ is idempotent commutative monoid. \otimes distributes over \oplus and 0 annihilates with \otimes .*

Idempotent semirings have a natural preorder, given by $U \leq V$ iff $U \oplus V \sim V$. Hence queries have this natural preorder. An immediate consequence is that choice is a colimit, i.e. least upper bound, of two queries.

Proposition 4. *Choice is a colimit of its branches. That is, $V \leq W$ and $U \leq W$, if and only if $V \oplus U \leq W$.*

The preorder over queries can be used to optimise queries. If a query offers a choice between a query and a weaker query, with respect to the preorder, the stronger branch may be eliminated. For instance, in related work [23], it is claimed that $U \text{ OPTIONAL } (V \text{ OPTIONAL } W)$ is not the same as $(U \text{ OPTIONAL } V) \text{ OPTIONAL } W$. Under the interpretation of **OPTIONAL** in the calculus it holds that $U \otimes ((V \otimes (W \oplus I)) \oplus I) \leq U \otimes ((V \oplus I) \otimes (W \oplus I))$, by distributivity, commutativity and idempotency. So the first is a stronger query.

A single rule is sufficient to capture the algebra of the select quantifier. From this algebra common equalities can be derived. The derived rules are suitable for the optimisation technique of flattening nested selects used in relational algebra [8]. The proof of commutativity of quantifiers requires capture avoiding substitution to be assumed. The presence of the tensor in the rule is required to prove that $\forall a. U \otimes V \leq \forall a.(U \otimes V)$, when $a \notin \text{fn}(V)$.

Proposition 5. *Selects are colimits of substitutions. So, $U \left\{ \frac{b}{a} \right\} \otimes V \leq W$ for all b , if and only if $\forall a. U \otimes V \leq W$. Immediate consequences are that, select commutes, distributes over choice, is annihilated by true and distributes over tensor. Furthermore, alpha conversion of bound variables is verified.*

The following rules of regular algebra hold. The first of the rules is sufficient to demonstrate that $*V \otimes U$ is a fixed point of the (monotone) map $W \mapsto U \oplus (V \otimes W)$. The second rule demonstrates that $*V \otimes U$ is the least such fixed point. Historically, Redko demonstrated that no finite collection of equations could axiomatise iteration [25]. The formulation below, was proven to be complete by Kozen [19].

Proposition 6. *An iterated query expands as follows $*U \sim I \oplus (U \otimes *U)$. Furthermore, if $U \oplus (V \otimes W) \leq W$ then $*V \otimes U \leq W$.*

A classic consequence of the above is that queries without select can always be denested to a single iteration [20]. However, select breaks denesting since iteration and select do not commute. For instance the following query requires two iterations. The result is that for each of the first continuation triggered, zero or more instances of the second continuation are triggered. This query can be expressed using sub-queries in the current SPARQL Query working draft [10].

$$*\sqrt{a}.\sqrt{n}.\left(\left(\left(a \text{ name } n\right); P\right) \otimes *\sqrt{e}.\left(\left(a \text{ email } e\right); Q\right)\right)$$

Iteration can be expressed as a colimit of repeated queries. This is a strictly more general property than Proposition 6 [18]. Since all constructs are colimits which distribute over tensor, the ideals generated by queries form a (commutative) quantale, as exploited by Montanari, Hoare and others [5, 12]. Quantales are related to spectral theory, which is related to information retrieval techniques used by search engines. Clarification of this connection is future work.

Proposition 7. *Iteration is a colimit of powers of queries. So, $U^n \otimes V \leq W$ for all n , if and only if $*U \otimes V \leq W$.*

Kozen demonstrates that Boolean algebras can be embedded in Kleene algebras [20]. The ‘tests’ of Kozen correspond to ‘constraints’ in SPARQL. Bisimulation verifies that the Boolean algebra of constraints embeds in the Kleene algebra, in the same manner, with similar consequences.

Proposition 8. *The Boolean algebra of constraints embeds in queries. Using standard classical implication, $\phi \Rightarrow \psi$ if and only if $\phi \leq \psi$. Or is choice, and is tensor, exists is select and an iterated constraint is always true.*

As with classical implication, the preorder over triples can be embedded in the partial order over processes. However, since alias assumptions are only a preorder, if $C \sim D$ then it holds that $C \sqsubseteq D$ and $D \sqsubseteq C$, which is weaker than equality. Maintaining distinction of names is important for applications where β is not fixed over time.

Proposition 9. *$C \sqsubseteq D$ if and only if $C \leq D$.*

The multiplicatives then, par and times and the units are related in the following manner. Combined with the previous rules the properties of then are established. The second rule shows that ‘then’ can be replaced by the unit delay (as in [1]).

Proposition 10. *An empty continuation can be removed, a continuation can be decomposed into the guard and a unit delayed process, and two continuations can be combined in a single par continuation, as follows.*

$$I; \perp \sim I \quad U \otimes (I; P) \sim U; P \quad (U; P); Q \sim U; (P \wp Q)$$

The algebra can be applied to optimise queries for distribution. In the example below the first query is rewritten as the tensor product of two queries.

$$*\sqrt{a}(((a \text{ knows } b_2); P) \oplus ((a \text{ knows } b_3); Q)) \sim *\sqrt{a}((a \text{ knows } b_2); P) \otimes *\sqrt{a}((a \text{ knows } b_3); Q)$$

The second query above is better for distribution. The tensor product allows two smaller queries to be immediately evaluated in parallel. The tighter scope of the select quantifiers reduces the branching when potential values to select are considered. The distribution of queries across clusters of servers is a major problem for processing Linked Data [11].

5 Conclusion

The calculus introduced provides the first operational semantics for SPARQL Query -- a W3C recommendation for querying Linked Data. The calculus has a concise logical semantics defined by a reduction system. The power of the calculus lies in the synchronisation primitives for queries. The synchronisation primitives are required to match the expressiveness of the core of SPARQL Query. Queries are internalised in a high-level process calculus, where query results determine continuation processes.

An alternative labelled transition system is shown to match the expressive power of the reduction system. Furthermore, the notion of bisimulation in the labelled transition system is sound with respect to equivalence in the reduction system. Bisimulation is used to verify an algebra over queries, which extends existing notions of an algebra for SPARQL Query. An algebra of queries is useful when tackling problems associated with Linked Data, such as distributed query planning.

The operational semantics combines several formalisms, as expected for a real language. The queries form a semiring, which provides a natural partial order. This partial order is used to characterise choice, selects and iteration as colimits. Also, iteration is the least fixed point of a monotonic map over queries, hence queries form a Kleene algebra. A preorder over URIs allows small permissible mismatches between content and queries to be resolved, capturing key features of the RDFS standard. Also, a Boolean algebra of constraints is naturally embedded in queries, to provide further control. The calculus demonstrates that key features of SPARQL and related standards for Linked Data can be tightly integrated in one framework.

References

- [1] Samson Abramsky, Simon Gay & Rajagopal Nagarajan (1995): *Interaction Categories and the Foundations of Typed Concurrent Programming*. In Manfred Broy, editor: *Proceedings of the NATO Advanced Study Institute on Deductive Program Design, Marktobendorf, Germany*, Springer, pp. 35--114.
- [2] Harith Alani et al. (2002): *Managing Reference: Ensuring Referential Integrity of Ontologies for the Semantic Web*. In Gómez-Pérez et al., editors: *Knowledge Engineering and Knowledge Management: Ontologies and the Semantic Web*, 2473, Springer, pp. 235--246, doi:10.1007/3-540-45810-7_29.
- [3] G. Bellin & P. J. Scott (1994): *On the π -Calculus and Linear Logic*. *Theoretical Computer Science* 135, pp. 11--65, doi:10.1016/0304-3975(94)00104-9.
- [4] Paul V. Biron & Ashok Malhotra (2004): *XML Schema part 2: Datatypes Second Edition*. W3C, MIT, Cambridge, MA. REC-xmlschema-2-20041028.
- [5] Stefano Bistarelli, Ugo Montanari & Francesca Rossi (1997): *Semiring-Based Constraint Satisfaction and Optimization*. *Journal of the ACM* 44(2), pp. 201--236, doi:10.1145/256303.256306.
- [6] Christian Bizer (2009): *The Emerging Web of Linked Data*. *IEEE Intelligent Systems* 24, pp. 87--92, doi:10.1109/MIS.2009.102.
- [7] Dan Brickley & R.V. Guha (2004): *RDF Vocabulary Description Language 1.0: RDF Schema*. W3C, MIT, Cambridge, MA. REC-rdf-schema-20040210.
- [8] Richard Cyganiak (2005): *A relational algebra for SPARQL*. Hewlett Packard Labs, Bristol. HPL-2005-170.
- [9] Jean-Yves Girard (1987): *Linear Logic*. *Theoretical Computer Science* 50(1), pp. 1--112, doi:10.1016/0304-3975(87)90045-4.
- [10] Steve Harris, Andy Seaborne & Eric Prud'hommeaux (2010): *SPARQL 1.1 Query Language*. W3C, MIT, Cambridge, MA. WD-sparql11-query-20101014.
- [11] Olaf Hartig et al. (2009): *Executing SPARQL Queries over the Web of Linked Data*. In A. Bernstein et al., editors: *The Semantic Web -- ISWC 2009, Chantilly, VA*, 5823, Springer, pp. 293--309, doi:10.1007/978-3-642-04930-9_19.
- [12] C. A. R. Tony Hoare, Bernhard Möller, Georg Struth & Ian Wehrman (2009): *Concurrent Kleene Algebra*. In Mario Bravetti & Gianluigi Zavattaro, editors: *CONCUR 2009, Bologna, Italy*, 5710, Springer, pp. 399--414, doi:10.1007/978-3-642-04081-8_27.
- [13] Joshua S. Hods & Dale Miller (1994): *Logic Programming in a Fragment of Intuitionistic Linear Logic*. *Information and Computation* 110(2), pp. 327--365.
- [14] Ross Horne (2011): *Programming Languages and Principles for Read-Write Linked Data*. Ph.D. thesis, Electronics and Computer Science, University of Southampton.
- [15] Alan Jeffrey & Julian Rathke (2005): *Contextual equivalence for higher-order π -calculus revisited*. *Logical Methods in Computer Science* 1(4), pp. 1--22, doi:10.2168/LMCS-1(1:4)2005.
- [16] Graham Klyne & Jeremy Carroll (2004): *Resource Description Framework: Concepts and Abstract Syntax*. W3C, MIT, Cambridge, MA. REC-rdf-concepts-20040210.

- [17] Naoki Kobayashi & Akinori Yonezawa (1993): *ACL -- A Concurrent Linear Logic Programming Paradigm*. In: *Proceedings of the 1993 International Logic Programming Symposium*, MIT Press, pp. 279--294.
- [18] Dexter Kozen (1990): *On Kleene algebras and Closed Semirings*. In Rovan, editor: *Proceedings on Mathematical Foundations of Computer Science*, 452, Springer-Verlag, pp. 26--47.
- [19] Dexter Kozen (1994): *A Completeness Theorem for Kleene Algebras and the Algebra of Regular Events*. *Information and Computation* 110, pp. 366--390, doi:10.1006/inco.1994.1037.
- [20] Dexter Kozen (1997): *Kleene algebra with tests*. *ACM Transactions on Programming Languages and Systems* 19, pp. 427--443, doi:10.1145/256167.256195.
- [21] Massimo Merro & Matthew Hennessy (2002): *Bisimulation congruences in safe ambients*. In: *Principles of programming languages*, ACM, pp. 71--80, doi:10.1145/503272.503280.
- [22] Robin Milner, Joachim Parrow & David Walker (1992): *A calculus of mobile processes, part I and II*. *Information and Computation* 100(1), pp. 1--40, doi:10.1016/0890-5401(92)90008-4.
- [23] Jorge Pérez, Marcelo Arenas & Claudio Gutierrez (2009): *Semantics and Complexity of SPARQL*. *ACM Transactions on Database Systems* 34(3), pp. 1--45, doi:10.1145/1567274.1567278.
- [24] Eric Prud'hommeaux & Andy Seaborne (2008): *SPARQL Query Language for RDF*. W3C, MIT, Cambridge, MA. REC-rdf-sparql-query-20080115.
- [25] V. N. Redko (1964): *On defining relations for the algebra of regular events*. *Ukrainskii Matematicheskii Zhurnal*, pp. 120--126.