

Metadata of the chapter that will be visualized in SpringerLink

Book Title	Theoretical Aspects of Software Engineering	
Series Title		
Chapter Title	Protocol Conformance of Collaborative SPARQL Using Multiparty Session Types	
Copyright Year	2024	
Copyright HolderName	The Author(s), under exclusive license to Springer Nature Switzerland AG	
Author	Family Name	Hernawan
	Particle	
	Given Name	Ari
	Prefix	
	Suffix	
	Role	
	Division	
	Organization	University of Oxford
	Address	Oxford, UK
	Email	ari.hernawan@cs.ox.ac.uk
	ORCID	http://orcid.org/0009-0001-2465-424X
Corresponding Author	Family Name	Yoshida
	Particle	
	Given Name	Nobuko
	Prefix	
	Suffix	
	Role	
	Division	
	Organization	University of Oxford
	Address	Oxford, UK
	Email	nobuko.yoshida@cs.ox.ac.uk
	ORCID	http://orcid.org/0000-0002-3925-8557
Abstract	Decentralised linked data gives users rights over their data while being accessible to other domains. The RDF (Resource Description Framework) and SPARQL have been the standard specifications for managing linked data for several years. Recent research and development introduce scalable, centralised and distributed RDF store engines with the SPARQL. However, writing SPARQL federated queries may grow more complex as the number of domain participants increases, presenting challenges such as source discovery, completeness and performance. This paper presents a SPARQL Query Template (SQT) that applies Multiparty Session Types (MPST) to determine the order of federated queries. We also guarantee protocol conformance between MPST and SPARQL relational algebra.	
Keywords (separated by '-')	Multiparty session types - SPARQL - Federated Query - Protocol Conformance	



Protocol Conformance of Collaborative SPARQL Using Multiparty Session Types

Ari Hernawan¹ and Nobuko Yoshida²

University of Oxford, Oxford, UK
{ari.hernawan,nobuko.yoshida}@cs.ox.ac.uk

Abstract. Decentralised linked data gives users rights over their data while being accessible to other domains. The RDF (Resource Description Framework) and SPARQL have been the standard specifications for managing linked data for several years. Recent research and development introduce scalable, centralised and distributed RDF store engines with the SPARQL. However, writing SPARQL federated queries may grow more complex as the number of domain participants increases, presenting challenges such as source discovery, completeness and performance. This paper presents a SPARQL Query Template (SQT) that applies Multiparty Session Types (MPST) to determine the order of federated queries. We also guarantee protocol conformance between MPST and SPARQL relational algebra.

AQ1

Keywords: Multiparty session types · SPARQL · Federated Query · Protocol Conformance

1 Introduction

The RDF (Resource Description Framework) and SPARQL 1.1¹ are the two W3C recommended standards for querying and manipulating linked data on the Web. Since SPARQL's introduction in 2008, significant research and technological advancements have supported RDF and SPARQL specifications. Many RDF repositories, such as DBpedia, UniProt (Universal Protein Resource) and Rhea, are publicly available to provide information from diverse domains [1–6]. DBpedia, for example, contains around 900 million triplets as of January 2021² and maintains a distributed infrastructure that extracts approximately 5500 triplets per second and 21 billion triples per release. In real-world applications, RDF datasets are often partitioned and built upon distributed environments [7].

There has been rigorous work on handling massive RDF datasets. One way to manage the datasets is by partitioning and distributing them under the same RDF engine or by decentralising them on domain-specific repositories. The latest distributed RDF engine implemented with a custom communication protocol

¹ <https://www.w3.org/TR/sparql11-overview/>.

² <https://www.dbpedia.org/resources/latest-core/>.

has shown superior performance compared to other centralised or MapReduce-based engines [7–16]. In addition, SPARQL 1.1 federated query allows merging data across decentralised repositories and heterogeneous RDF engines. Federated query processors have been improved in many aspects, including querying documents, query planner optimisation and sensitive data access policies [17–23]. A federated query execution framework depends on source selection as it determines which endpoints are relevant to evaluate a given query. The Source selection is a separate step before execution, aiming to reduce the number of requests to the RDF server, optimise resource allocations and minimise network overhead, improving overall query execution time.

Well-specified communication is crucial for distributed RDF engines and federated queries, leading to significant performance increases. For communication-focused programming, a framework of Multiparty Session Types (MPST) provides a typed process that abstracts communication between multiple participants into a global protocol (type) [24]. The theory formalise endpoint projection of the global description in the Web Services Choreography Description Language³, but it can accommodate other diverse specifications [25]. We aim to integrate this typed theory with SPARQL. This theory matches common issues in the software development process and is essential for a collaborative team of engineers writing SPARQL. A global type (protocol) gives engineers a global view of federated query communication between RDF repositories. The engineers can validate their SPARQL federated query following a global protocol.

We highlight current SPARQL issues as follows. (1) *Source discovery*: Writing a SPARQL query that involves many RDF repositories and complex query joins can be burdensome; and (2) *Performance and completeness*: Query execution time depends on the performance of SPARQL runtime and federated communication cost. More importantly, the query must return a complete result.

This paper introduces a novel SPARQL Query Template (SQT) generator that leverages the benefits of Multiparty Session Types [24, 26, 27]. The theory notably gives some clarity to complex processes. The SQT generates federated query body syntax that includes the RDF URL source address and guarantees conformance with the projection of a given global type. Moreover, by using a global type to describe communication patterns, we can identify and prevent overhead communications, speed up query execution and ensure the completeness of the results.

In the remainder, Sect. 2 gives preliminaries of RDF, MPST and SPARQL relational algebra through examples. Section 3 illustrates a use case that motivates this work. We explain in detail how SQT produces query templates from a global type. Section 4 explains how local types conform to relational algebra. Section 5 evaluates different network communication costs associated with different implementations. Finally, Sect. 6 discusses current limitations and future works of current work.

³ <https://www.w3.org/2002/ws/chor/>.

2 Preliminaries and Examples

This section introduces RDF notations, SPARQL relational algebra and MPST through simple examples.

2.1 RDF and SPARQL

Our basic definitions of RDF and SPARQL follow [28,29]. In addition, we add a representation of RDF graph [16] and conjunctive query [7].

Definition 1 (RDF Triple). *RDF terms (t) are consist of IRIs (**I**), literals (**L**) and blank nodes (**B**). These terms form a triple or 3-tuple, $\langle t_s, t_p, t_o \rangle \in (\mathbf{I} \cup \mathbf{B}) \times (\mathbf{I}) \times (\mathbf{I} \cup \mathbf{B} \cup \mathbf{L})$, where s , p and o denote subject, predicate and object of the triple, $\beta = \{s, p, o\}$.*

A set of RDF triples indicates a semantic relationship between terms and can be represented as a directed graph [16]. We can construct and label multiple graphs into a dataset. However, for convenience, this paper focuses on a single graph in a dataset.

Definition 2 (RDF Dataset). *Dataset (D) is a set of triples. The semantic relationship on triples can be described as a directed labelled graph. A graph consists of vertices (V) and directed edges (E), denoted by $D = \{V, E\}$. Vertices V denotes a set of resource nodes $t_s \cup t_o \in V$. Edges E is a set of directed edges connecting the nodes in V . To simplify, the directed edge $t_s \xrightarrow{t_p} t_o$ is expressed by $\langle t_s, t_p, t_o \rangle$.*




Fig. 1. RDF Dataset

Example 1. To illustrate, we define a dataset carrying a relation between Alice and Bob, presented as vertices and directed edges as shown in Fig. 1.

$$D = \{V, E\}$$

$$V = \{Alice, Bob, foaf:Person, p_0, p_1\}$$

$$E = \{\langle p_0, foaf:knows, p_1 \rangle, \langle p_0, rdf:type, foaf:Person \rangle,$$

$$\langle p_1, rdf:type, foaf:Person \rangle, \langle p_0, foaf:name, Alice \rangle, \langle p_1, foaf:name, Bob \rangle\}$$

Definition 3 (SPARQL Query Template). *The SPARQL Query Template (SQT) syntax is denoted as follows:*

$$\begin{aligned} Q, Q' &::= A && (\text{triple pattern or atom}) \\ &| \mathcal{S}.Q' && (\text{service federation}) \\ &| Q \wedge Q' && (\text{group of basic graph pattern}) \end{aligned}$$

The SPARQL relational algebra formally defines a query language semantics [7, 28, 29]. A query's result or solution is a tuple containing variables from queries matched against the RDF dataset. This paper recursively constructs queries from graph patterns, including triple patterns, basic graph patterns and service federation.

Definition 4 (Triple Pattern or Atom). *Triple pattern or atom (A) forms $\langle t_s, t_p, t_o \rangle \in (\mathbf{I} \cup \mathbf{B} \cup \mathbf{V}) \times (\mathbf{I} \cup \mathbf{V}) \times (\mathbf{I} \cup \mathbf{B} \cup \mathbf{L} \cup \mathbf{V})$ where $(\mathbf{V})$ refers to variables. The variables occurring in an atom are denoted by $\text{vars}(A)$. The triple pattern evaluates RDF triples in D by selection function (σ) with constant terms t_k and by mapping function (π) that maps terms t_j to the variables v_i :*

$$A ::= \pi_{v_i \leftarrow t_j}(\sigma_{t_k}(D)) \quad (j, k \in \beta \text{ and } i \in \mathbb{N})$$

$$\text{vars}(A) = \{t_\beta | t_\beta \cap \mathbf{V}\} \quad (t_k \in (\mathbf{I} \cup \mathbf{B} \cup \mathbf{L}) \text{ and } t_j \in (\mathbf{I} \cup \mathbf{B} \cup \mathbf{L} \cup \mathbf{V}))$$

Example 2. A query for finding a person's name on dataset D is defined as:

$$\{?x \text{ foaf:name } ?y\}$$

This query translates to the relational operation:

$$\langle x, \text{foaf:name}, y \rangle = \pi_{x, y \leftarrow t_s, t_o}(\sigma_{t_p=\text{foaf:name}}(D))$$

The process begins with searching any triples in a dataset that match the selection's predicate, $t_p = \text{foaf:name} \in \mathbf{I}$. Here, we have two variables $?x$ and $?y$ in subject and object respectively. Subject $t_s \in \mathbf{I} \cup \mathbf{B} \cup \mathbf{V}$ and object $t_o \in \mathbf{I} \cup \mathbf{B} \cup \mathbf{L} \cup \mathbf{V}$ intersect with $\mathbf{V}$ and will be included in $\text{vars}(A) = \{t_s, t_o\}$.

Definition 5 (Basic Graph Pattern). *The basic graph pattern is a group of triple patterns. It can be expressed syntactically as a conjunctive query Q from finite set atoms, where $\text{vars}(Q)$ denotes variables occurring in the basic graph pattern.*

$$Q ::= A_1 \wedge \cdots \wedge A_n \quad \text{and} \quad \text{vars}(Q) = \text{vars}(A_1) \cup \cdots \cup \text{vars}(A_n)$$

Multiple triple patterns stacked together will construct a basic graph pattern [7]. The basic graph pattern variables are the union of all triple pattern variables. These graph patterns can be evaluated in a local or remote dataset, formalised in Definition 6 [29]. The definition of set variables also applies equally to remote datasets.

Definition 6 (Service Federation). The SPARQL service ($\mathcal{S}$) is a query evaluation on a remote dataset. Query evaluation Q on remote service $\mathcal{S}$ is denoted as $(\mathcal{S}.Q)$ while local query evaluation is denoted as (Q) .

The SPARQL query Q is evaluated over dataset D and then returns solution mapping that represents matched variables on triples in a dataset [29]. The solution mapping function is defined in Definition 7. We assume all variables are bounded and no error returns from the graph pattern.

Definition 7 (Solution Mapping Function). A solution mapping (π) partially maps variables $\mathbf{V}$ from terms $\mathbf{I} \cup \mathbf{B} \cup \mathbf{L}$. Suppose $\{v_i, \dots, v_n\} \subseteq \mathbf{V}$ and $\{x_i, \dots, x_n\} \subseteq \text{vars}(Q)$ for $n \in \mathbb{N}$. Then we denote by $\{v_i/x_i, \dots, v_n/x_n\}$ mapping function π such that the solution contains a set of $\text{dom}(\pi) = \{v_i/x_i\}$.

$$\pi_{v_i} = x_i \quad (v_i \in \text{dom}(\pi)) \quad \text{and} \quad \pi_\emptyset = \emptyset \quad (\text{dom}(\pi) = \emptyset)$$

Example 3. Assume there are two datasets. The local dataset contains triples of persons. The remote dataset contains people's relations with each other. The first graph pattern is evaluated locally and the second pattern is service federation, which is evaluated on a remote dataset as shown in Fig. 2. The local graph pattern finds subject terms with type foaf:Person and map into variable $?x \in \mathbf{I}$. Then, a service federation pattern finds the person $?y \in \mathbf{I}$ that is acquainted with $?x$. A complete query syntax is presented in Fig. 3.




Fig. 2. SPARQL relational algebra

```

1  SELECT ?x ?y /* return variables */
2  WHERE {
3    ?x rdf:type foaf:Person . /* graph pattern is evaluated on local dataset */
4    SERVICE <https://remote>{
5      ?x foaf:knows ?y . /* graph pattern is evaluated on remote dataset */
6    }
7  }
```

Fig. 3. SPARQL 1.1 Query

2.2 Multiparty Session Types

Global types (protocol) G provides a blueprint describing communication between participants as a pair of a sender and a receiver [24, 26, 27]. In Definition 8, *message type* formalises protocol where p sends to q a value with type U . Similarly, *branching type* formalises participant p chooses a label l_i from q for some $i \in I$. The *termination type* **end** denotes communication is finished and no further process.

Definition 8 (Global and Local Types). The global type G and local type T are denoted as follows:

$$\begin{aligned}
 G &::= \text{end} \mid \mathbf{p} \rightarrow \mathbf{q} : [U]; G \mid \mathbf{p} \rightarrow \mathbf{q} : \{l_i : G_i\}_{i \in I} && (\text{termination, message, branching}) \\
 T &::= \text{end} \mid \mathbf{q}![U]; T \mid \mathbf{p}?[U]; T && (\text{termination, send, receive}) \\
 &\mid \mathbf{q} \oplus \{l_i : T_i\}_{i \in I} \mid \mathbf{p} \& \{l_i : T_i\}_{i \in I} && (\text{internal choice, external choice})
 \end{aligned}$$

We require that $\mathbf{p} \neq \mathbf{q}$, $I \neq \emptyset$, and $i \neq j$ whenever $i \neq j$, for all $i, j \in I$.

Local type T defines a process for a single participant following a global type G [24, 26, 27]. *Termination type* (**end**) denotes the end of a process. *Send type* $\mathbf{q}![U]; T$, sends a message to $\mathbf{q}$ with payload U . *Receive type* $\mathbf{p}?[U]; T$, waits for a message from $\mathbf{p}$ with payload U . *External choice* $\mathbf{p} \& \{l_i : T_i\}_{i \in I}$ waits for selection label l_i from local type to $\mathbf{p}$ to continues. *Internal choice* $\mathbf{q} \oplus \{l_i : T_i\}_{i \in I}$ formalises a label selection l_i at $\mathbf{q}$ and continues.

Definition 9 (Participants). We define the set of participants of a global type G as follows:

$$\begin{aligned}
 pt(\text{end}) &= \emptyset \\
 pt(\mathbf{p} \rightarrow \mathbf{q} : [U]; G) &= \{\mathbf{p}, \mathbf{q}\} \cup pt(G) \\
 pt(\mathbf{p} \rightarrow \mathbf{q} : \{l_i : G_i\}_{i \in I}; G) &= \{\mathbf{p}, \mathbf{q}\} \cup \bigcup_{i \in I} pt(G_i)
 \end{aligned}$$

A set of participants in *message type* is their sender $\mathbf{p}$ and receiver $\mathbf{q}$ which are defined in a global type. Similarly, for *branching type*, a set of participants is the label sender $\mathbf{p}$ and label receiver $\mathbf{q}$. A *branching type* has multiple continuations G_i and the set of participants is a union from all branches. A *termination type* does not have continuation and does not involve any participants.

Definition 10 (Projection). The projection of a global type G onto a participant $\mathbf{r}$ is defined as follows:

$$\begin{aligned}
 \text{end} \upharpoonright \mathbf{r} &= \text{end} \\
 \mathbf{p} \rightarrow \mathbf{q} : [U]; G \upharpoonright \mathbf{r} &= \begin{cases} \mathbf{q}![U]; G \upharpoonright \mathbf{r} & \mathbf{r} = \mathbf{p} \\ \mathbf{p}?[U]; G \upharpoonright \mathbf{r} & \mathbf{r} = \mathbf{q} \\ G \upharpoonright \mathbf{r} & \mathbf{r} \notin \{\mathbf{p}, \mathbf{q}\} \end{cases} \\
 \mathbf{p} \rightarrow \mathbf{q} : \{l_i : G_i\}_{i \in I}; G \upharpoonright \mathbf{r} &= \begin{cases} \mathbf{q} \oplus \{l_i : G_i \upharpoonright \mathbf{r}\}_{i \in I} & \mathbf{r} = \mathbf{p} \\ \mathbf{p} \& \{l_i : G_i \upharpoonright \mathbf{r}\}_{i \in I} & \mathbf{r} = \mathbf{q} \\ \bigcap_{i \in I} G_i \upharpoonright \mathbf{r} & \mathbf{r} \notin \{\mathbf{p}, \mathbf{q}\} \end{cases}
 \end{aligned}$$

A global type can be projected into a local type of participant $G \upharpoonright \mathbf{r} = T_{\mathbf{r}}$. The *termination type* **end** does not transform during the projection. The projection of *message type* becomes a *send type* where $\mathbf{r} = \mathbf{p}$, $\mathbf{r}$ is a sender and becomes a *receive type* where $\mathbf{r} = \mathbf{q}$, $\mathbf{r}$ is a receiver. The *branching type* is projected to an *external choice* where $\mathbf{r} = \mathbf{p}$ and into an *internal choice* where $\mathbf{r} = \mathbf{q}$. Projection of *branching type* to the third participant $\mathbf{r}$, neither $\mathbf{p}$ nor $\mathbf{q}$, must specified to have the same continuation for any branch label.

Definition 11 (Global Type Reduction). Evaluating $\mathbf{p} \rightarrow \mathbf{q}$ in global type G is defined as follows:

$$\begin{aligned}
 & \text{(message reduction)} \\
 & (\mathbf{p} \rightarrow \mathbf{q} : [U]; G) \setminus \mathbf{p} \xrightarrow{U} \mathbf{q} = G \\
 & \text{(pass message reduction)} \\
 & (\mathbf{r} \rightarrow \mathbf{s} : [U]; G) \setminus \mathbf{p} \xrightarrow{U'} \mathbf{q} = \mathbf{r} \rightarrow \mathbf{s} : [U]; (G \setminus \mathbf{p} \xrightarrow{U'} \mathbf{q}) \\
 & \quad \text{if } \{\mathbf{r}, \mathbf{s}\} \cap \{\mathbf{p}, \mathbf{q}\} = \emptyset \text{ and } \{\mathbf{p}, \mathbf{q}\} \subseteq G \\
 & \text{(branching reduction)} \\
 & (\mathbf{p} \rightarrow \mathbf{q} : \{l_i : G_i\}_{i \in I}) \setminus \mathbf{p} \xrightarrow{l} \mathbf{q} = G_i \quad \text{for } l = l_i \\
 & \text{(pass branching reduction)} \\
 & (\mathbf{r} \rightarrow \mathbf{s} : \{l_i : G_i\}_{i \in I}) \setminus \mathbf{p} \xrightarrow{l} \mathbf{q} = \mathbf{r} \rightarrow \mathbf{s} : \left\{ l_i : G_i \setminus \mathbf{p} \xrightarrow{l} \mathbf{q} \right\}_{i \in I} \\
 & \quad \text{if } \{\mathbf{r}, \mathbf{s}\} \cap \{\mathbf{p}, \mathbf{q}\} = \emptyset \text{ and } \forall i \in I : \{\mathbf{p}, \mathbf{q}\} \subseteq G_i
 \end{aligned}$$

Message reduction $\mathbf{p} \xrightarrow{U} \mathbf{q}$ reduces the global type G with *message type* $\mathbf{p} \rightarrow \mathbf{q} : [U]$ and continues to the next protocol G . *Branching reduction* $\mathbf{p} \xrightarrow{l} \mathbf{q}$ reduces multiple label selection $\mathbf{p} \rightarrow \mathbf{q} : \{l_i : G_i\}_{i \in I}$ by choosing one label l_i then continue with G_i . *Pass reduction* applies when participants of reduction $\{\mathbf{p}, \mathbf{q}\}$ do not meet $\{\mathbf{r}, \mathbf{s}\}$ such as reduction $\mathbf{p} \xrightarrow{U} \mathbf{q}$ on $\mathbf{r} \rightarrow \mathbf{s} : [U]$ or $\mathbf{p} \xrightarrow{l} \mathbf{q}$ on $\mathbf{r} \rightarrow \mathbf{s} : \{l_i : G_i\}_{i \in I}$. The reduction continues through the rest of the global type until it meets the same participant.

3 Overview

Many research institutes [1–6] provide open RDF datasets. These datasets are often used towards drug discovery. Researchers use them to work with genomic, metabolomic, molecular structures, reaction pathways and screening drug candidates. The information resides in different domains and locations. Data integration and retrieval are the keys to comprehensive analyses. The SPARQL federated query techniques allow engineers to query remote datasets like a single unified database.

We use UniProt query example number 45⁴ to cover all discussions in this paper. The use case describes a query of retrieving drug targets for human enzymes involved in sterol metabolism. The example uses three SPARQL endpoints, Rhea⁵, UniProt⁶ and Wikidata⁷. The query can be run directly at

⁴ <https://sparql.uniprot.org/.well-known/sparql-examples/?offset=40>.

⁵ <https://sparql.rhea-db.org/sparql>.

⁶ <https://sparql.uniprot.org/sparql>.

⁷ <https://query.wikidata.org/sparql>.




Fig. 4. The relational algebra of retrieving drug targets for human enzymes involved in sterol metabolism

UniProt query editor. Alternatively, it can be executed in Rhea or Wikidata but requires some rewriting at the federated query part. We aim to make it less tiring to write a federated query by modelling this in a global type.

We present a query relational algebra and its local type representation in Fig. 4. The query involves three participants: **rhea**, **uniprot** and **wikidata**. Query $Q_{uniprot}$ is evaluated at **uniprot**. Conjunctive query $Q_{uniprot}$ consists of two service federations and a local graph pattern. The query evaluation follows the sequence: Rhea, UniProt and Wikidata. First, **uniprot** evaluates federated query at **rhea** and retrieves reactions involving some molecules. Then, it joins the result with a local query at **uniprot** to find human enzymes that catalyse reactions. Finally, **uniprot** finds the interaction of enzymes from previous results with drug descriptions in **wikidata**. It is important to note that query $Q_{uniprot}$ and local type $T_{uniprot}$ are specifications for **uniprot** only. Other participants will need a different query algebra and a local type definition. The global type provides a global view and ensures accurate and consistent results across all participants during query execution.

4 Implementation of SQT Workflow

This section explains the workflow used to generate the SPARQL Query Template. The top-down diagram in Fig. 5 shows the process from global type to query template. We define global type G with four participants: **rhea**, **uniprot**, **wikidata** and **local**. The **local** participant physically does not exist. Instead, **local** acts as **rhea**, **uniprot** or **wikidata**.

We begin with building a global type G , in Example 4, using the Scribble script defined in Fig. 6. Suppose a script `Payload()` from **p** to **q**; the participant **q** expects a message return from **p**. This is defined as $p \rightarrow q : [U]$ where payload U contains free tuples of terms t_β . Here, **p** is a participant who owns a dataset and executes query Q either in a local or remote environment.




Fig. 5. Top-down view of global type for UniProt query number 45. The r stand for **rhea**, u for **uniprot**, w for **wikidata** and l stand for **local**

The global type is implemented in Scribble script in Fig. 6. Lines 4-8 represents G_{rhea} a choice of **local** acts as **rhea**, lines 18-22, G_{uniprot} where **local** as **uniprot** and lines 32-36, G_{wikidata} where **local** as **wikidata**. Implementation of **rhea** evaluates query locally shown in line 10 and federated query in line 11. The definition of global type G is required to be projectable. Otherwise, it cannot produce any local types T and queries Q .

We introduce a *broadcast type* with the global type definition. The *broadcast type* formed as a sequence of *branching type* and uses one label that same as the previously selected branch. In Definition 12, participant p broadcast a l_i to $\bar{q}$ is expressed by selecting a label l_i to a finite set number of participant $\bar{q} = \{q_1, q_2, \dots, q_n\}$.

Definition 12 (Broadcast Type). A *broadcast type* is sending or selecting the same label to multiple participants. A complete set of global type syntax with *broadcast type* is given as follows:

$$\begin{aligned}
 G &::= \text{end} \mid p \rightarrow \bar{q} : [U] ; G && (\text{termination, message}) \\
 &\mid p \rightarrow \bar{q} : \{l_i : G_i\}_{i \in I} && (\text{branching}) \\
 p \rightarrow \bar{q} : [U] ; G &\stackrel{\text{def}}{=} p \rightarrow q_1 : [U] ; p \rightarrow q_1 : [U] ; \dots \\
 &\quad p \rightarrow q_n : [U] ; G && (\text{broadcast message}) \\
 p \rightarrow \bar{q} : \{l_i : G_i\}_{i \in I} &\stackrel{\text{def}}{=} p \rightarrow q_1 : \{l_i : p \rightarrow q_2 : \{l_i : \dots \\
 &\quad p \rightarrow q_n : \{l_i : G_i\}\}_{i \in I} && (\text{broadcast branching}) \\
 \text{with } \bar{q} &= \{q_1, q_2, \dots, q_n\} \quad n \in \mathbb{N}
 \end{aligned}$$

The broadcast reduction evaluates communication between participants of a *broadcast type* in global type G or continues to go through the next types when participants involved in broadcast type differ from the reduction criteria. Reducing *branching type* $p \xrightarrow{l_i} q_1$ or $p \xrightarrow{l_i} q_2$ in global type G can happen individually in a different order, but a broadcast reduction must reduce all $\bar{q}$ together.

Definition 13 (Broadcast Reduction). *Evaluating a broadcast type in a global type G is defined as follows:*

(message reduction)

$$(p \rightarrow \bar{q} : [U]; G) \setminus p \xrightarrow{U} \bar{q} \cong G$$

(pass message reduction)

$$(r \rightarrow \bar{s} : [U]; G) \setminus p \xrightarrow{U'} \bar{q} \cong r \rightarrow \bar{s} : [U]; (G \setminus p \xrightarrow{U'} \bar{q})$$

if $\{r, \bar{s}\} \cap \{p, \bar{q}\} = \emptyset$ and $\{p, \bar{q}\} \subseteq G$

(branching reduction)

$$p \rightarrow \bar{q} : \{l_i : G_i\}_{i \in I} \setminus p \xrightarrow{l} \bar{q} \cong G_i \quad l = l_i$$

(pass branching reduction)

$$r \rightarrow \bar{s} : \{l_i : G_i\}_{i \in I} \setminus p \xrightarrow{l} \bar{q} \cong r \rightarrow \bar{s} : \left\{ l_i : G_i \setminus p \xrightarrow{l} \bar{q} \right\}_{i \in I}$$

if $\{r, \bar{s}\} \cap \{p, \bar{q}\} = \emptyset$ and $\forall i \in I : \{p, \bar{q}\} \subseteq G_i$

Theorem 1 (Broadcast). *If $G \setminus p \xrightarrow{l} \bar{q} \cong G'$ then $G \setminus p \xrightarrow{l} q_1 \setminus p \xrightarrow{l} q_2 \cdots \setminus p \xrightarrow{l} q_n = G'$ where $n \in \mathbb{N}$ and $\bar{q} = \{q_1, q_2, \dots, q_n\}$.*

Proof. By induction on n and G .

Example 4. We formalise a global type for UniProt query number 45 as follows:

$$I = \{\text{rhea}, \text{uniprot}, \text{wikidata}\}$$

$$G = \text{local} \rightarrow \text{rhea} \{l_i : \text{local} \rightarrow \text{uniprot} \{l_i : \text{local} \rightarrow \text{wikidata} \{l_i : G_i\}\}_{i \in I}$$

$$\begin{aligned} G_{\text{rhea}} &= \text{local} \rightarrow \text{rhea} : [U]; \text{uniprot} \rightarrow \text{rhea} : [U]; \\ &\quad \text{wikidata} \rightarrow \text{rhea} : [U]; \text{end} \\ G_{\text{uniprot}} &= \text{rhea} \rightarrow \text{uniprot} : [U]; \text{local} \rightarrow \text{uniprot} : [U]; \\ &\quad \text{wikidata} \rightarrow \text{uniprot} : [U]; \text{end} \\ G_{\text{wikidata}} &= \text{rhea} \rightarrow \text{wikidata} : [U]; \text{uniprot} \rightarrow \text{wikidata} : [U]; \\ &\quad \text{local} \rightarrow \text{wikidata} : [U]; \text{end} \end{aligned}$$

Below is a global type reduction for **local** by choosing label $l_i = l$:

$$G \setminus \text{local} \xrightarrow{l} \{\text{rhea}, \text{uniprot}, \text{wikidata}\} = G_i$$

Algorithm 1 translates a global type to a query template, following translation on Table 1. The algorithm has two inputs, the projection of global protocol $G \upharpoonright p$ and target participant q where a query will be executed. Suppose we use a global type G in Example 4 to build a query template for **uniprot**. This global type G has three branches for different continuations: G_{rhea} , G_{uniprot} and G_{wikidata} . Initially, we set $p = \text{uniprot}$ for $G \upharpoonright p$ and participant $q = \text{uniprot}$. Next, global type G reduced by selecting a branch for **uniprot**, $\setminus \text{local} \xrightarrow{1} \bar{s}$ where $l = q$ and $\bar{s} = \{\text{rhea}, \text{uniprot}, \text{wikidata}\}$. Selecting label $l_i = l$ continues

```

1  global protocol FederatedQuery(role Local, role Rhea, role UniProt,role Wikidata){
2      choice at Local{
3          /* Local tells everyone that it chooses to act as Rhea */
4          Rhea() from Local to Rhea;
5          choice at Local{
6              Rhea() from Local to UniProt;
7              choice at Local{
8                  Rhea() from Local to Wikidata;
9                  /* Rhea */
10                 Select(reaction) from Local to Rhea;
11                 Select(protein,proteinFullName) from UniProt to Rhea;
12                 Select(chemical,chemicalLabel,treatedLabel) from Wikidata to Rhea;
13             }
14         }
15     }
16     or {
17         /* Local tells everyone that it chooses to act as UniProt */
18         UniProt() from Local to Rhea;
19         choice at Local{
20             UniProt() from Local to UniProt;
21             choice at Local{
22                 UniProt() from Local to Wikidata;
23                 /* UniProt */
24                 Select(reaction) from Rhea to UniProt;
25                 Select(protein,proteinFullName) from Local to UniProt;
26                 Select(chemical,chemicalLabel,treatedLabel) from Wikidata to UniProt;
27             }
28         }
29     }
30     or {
31         /* Local tells everyone that it chooses to act as Wikidata */
32         Wikidata() from Local to Rhea;
33         choice at Local{
34             Wikidata() from Local to UniProt;
35             choice at Local{
36                 Wikidata() from Local to Wikidata;
37                 /* Wikidata */
38                 Select(reaction) from Rhea to Wikidata;
39                 Select(protein,proteinFullName) from UniProt to Wikidata;
40                 Select(chemical,chemicalLabel,treatedLabel) from Local to Wikidata;
41             }
42         }
43     }
44 }

```

Fig. 6. Scribble script of global type for UniProt query number 45

with G_i shows the role of **local** which can either become **rhea** following with G_{rhea} , **uniprot** with G_{uniprot} , or **wikidata** with G_{wikidata} .

The Algorithm 1 at line number 2 evaluates $G_i \upharpoonright p$ whether it is an empty (**end**) or continues to be translated. In the case of $p = q$, a *message type* in $G_i \upharpoonright p = r?[U]$ and **r** is a **local** then translates a message type to a query Q . This query will be executed locally at **p**. On the other hand, the message type is translated to service federation $S_p.Q$, when **p** is neither **local** or current target participant **q**. In summary, $G_{\text{uniprot}} \upharpoonright \text{uniprot}$ translates to a local query in **uniprot** and $G_{\text{uniprot}} \upharpoonright \text{rhea}$ explains that **uniprot** sends and evaluates a query to **rhea**.

Table 1 only considers a *message type* to be translated directly into a query. The **end** type is translated to nothing and represents the end of a query. However,

Table 1. Translation of local type and SPARQL relational algebra

Local type ($T_q = G \upharpoonright q$)		SPARQL algebra (Q_q)
$\llbracket \text{end} \rrbracket$	=	$\emptyset$
$\llbracket p?[U]; G \upharpoonright q \rrbracket \quad p \neq \text{Local}$	=	$\mathcal{S}_p. \llbracket G \upharpoonright p \rrbracket \wedge \llbracket G \upharpoonright q \rrbracket$
$\llbracket p?[U]; G \upharpoonright q \rrbracket \quad p = \text{Local}$	=	$Q_q \wedge \llbracket G \upharpoonright q \rrbracket$
$\llbracket p![U]; G \upharpoonright q \rrbracket$	=	$\llbracket G \upharpoonright q \rrbracket$
$\llbracket p \oplus \{l_i : G_i \upharpoonright q\}_{i \in I} \rrbracket$	=	$\llbracket G_i \upharpoonright q \rrbracket_{i \in I}$
$\llbracket p \& \{l_i : G_i \upharpoonright q\}_{i \in I} \rrbracket$	=	$\llbracket G_i \upharpoonright q \rrbracket_{i \in I}$

internal choice and *external choice*, defined in Table 1, do not have a query translation. Instead, it continues to translate the following local types. Although it does not have query translation, it plays an essential part in the reduction process by deciding what is **local** participant role will be.

In a simple federation query where **uniprot** calls federation service to both **rhea** and **wikidata**. Here, **rhea** and **wikidata** send their result directly to **uniprot** and do not expect to receive a result from calling another federation service. Both local types T_{rhea} and T_{wikidata} define **rhea**![U] which does not translate to any queries. The algorithm continues to read recursively and reduces the local type from all participants until it is finished.

In a nested service federation, each local type for three participants, including **rhea**, **uniprot** and **wikidata** may have called a federation service to other participants. For example, **wikidata** expects a federation service result from **uniprot**, $T_{\text{wikidata}} = \text{uniprot}?[U]$. Meanwhile, **uniprot** is waiting federation service result from **rhea** before it can be sent to **wikidata**, $T_{\text{uniprot}} = \text{rhea}?[U]; \text{wikidata}![U]$. The SQT algorithm requires definitions from all participant's local types in order

Algorithm 1. Construct SPARQL Query Template

```

1: function SQT( $G \upharpoonright p, q$ ) ▷ begin with  $p = q$  where  $p, q \in pt(G) \setminus \{\text{local}\}$ 
2:   while ( $G \setminus \text{local} \xrightarrow{l} \bar{s}$ )  $\upharpoonright p \neq \text{end}; l = q$  do ▷  $\bar{s} = pt(G) \setminus \{\text{local}\}$ 
3:     if  $T_p = r?[U]$  then ▷  $p \neq r; r \in pt(G)$ 
4:       if  $r = \text{local}$  then
5:          $Q_q \leftarrow \llbracket r?[U] \rrbracket$  ▷ local query  $Q$ 
6:       else
7:          $\mathcal{S}_r.(Q_q) \leftarrow \text{SQT}(G \upharpoonright r, q)$  ▷ federated query  $\mathcal{S}.Q$ 
8:       end if
9:     end if
10:     $Q'_q = Q_q \wedge \mathcal{S}_r.Q_q \wedge \dots \wedge \mathcal{S}_{r'}.Q_q$  ▷ Conjunctive query
11:  end while
12:  return  $Q'_q$ 
13: end function

```

to generate a complete query template. The algorithm begins by generating a local query and service federation for target participant q then calling a related local type to generate a query in federated service.

5 Evaluation

We conduct experiments in the docker environment. We build a small-scale environment with three Apache Fuseki containers and a triple dataset to imitate Rhea, UniProt and Wikidata. For statistics **rhea** contains 2051 triples, **uniprot** 129902 triples and **wikidata** 6592 triples. We record the elapsed projection time, constructing SQT and performance differences between variants of federated query communication that follow the current use case.

Example 5. We formalise the general structure of a global protocol (type) for a current evaluation as follows:

$$G \setminus \text{local} \xrightarrow{l} \bar{q} = G' \quad \text{where} \quad l \in pt(G) \setminus \{\text{local}\}$$

We prepare global type G following Example 5. We assign varying numbers of participants to global type G : 10, 30, and 60. Table 2 shows the elapsed time for constructing SQT. The projection column shows the time needed to generate all local types. It shows that increasing the number of participants will result in a longer elapsed time.

We prepared three global types following Example 5 for service federation performance. We define each global type to have three specifications for **rhea**, **uniprot** and **wikidata**. We have nine different test cases in total. Global type G_1 calls two remote service federations from one participant, following Example 4. However, global types G_2 and G_3 call nested service federations. Global type G_2 is given in Example 6.

Example 6. Based on the global type in Example 5, we define G_{uniprot} such that **uniprot** can only call **rhea** service federation through **wikidata**. The complete

Table 2. Construction time for single query

Number of participants in G	Time (second)	
	Projection (total)	SQT (average)
10 participants	0.445	0.321
30 participants	2.720	0.724
60 participants	30.624	1.429

continuation of global type G_i for nested service federation defined as follows:

```

 $I$            = {rhea, uniprot, wikidata}
 $G_{\text{rhea}}$     = local  $\rightarrow$  rhea : [U]; uniprot  $\rightarrow$  rhea : [U];
               wikidata  $\rightarrow$  uniprot : [U]; end
 $G_{\text{uniprot}}$   = rhea  $\rightarrow$  wikidata : [U]; wikidata  $\rightarrow$  uniprot : [U];
               local  $\rightarrow$  uniprot : [U]; wikidata  $\rightarrow$  uniprot : [U]; end
 $G_{\text{wikidata}}$  = rhea  $\rightarrow$  wikidata : [U]; uniprot  $\rightarrow$  rhea : [U];
               local  $\rightarrow$  wikidata : [U]; end

```

Table 3. Query elapsed time from different global protocol

Query Template	Elapsed time (second)
1. $\llbracket (G_1 \setminus 1 \xrightarrow{r} r, u, w) \upharpoonright r \rrbracket = Q_r \wedge S_u.Q_r \wedge S_w.Q_r$	23.160
2. $\llbracket (G_1 \setminus 1 \xrightarrow{u} r, u, w) \upharpoonright u \rrbracket = S_r.Q_u \wedge Q_u \wedge S_w.Q_u$	12.255
3. $\llbracket (G_1 \setminus 1 \xrightarrow{w} r, u, w) \upharpoonright w \rrbracket = S_r.Q_w \wedge S_u.Q_w \wedge Q_w$	11.334
4. $\llbracket (G_2 \setminus 1 \xrightarrow{r} r, u, w) \upharpoonright r \rrbracket = Q_r \wedge S_u(Q_r \wedge S_w.Q_r)$	24.500
5. $\llbracket (G_2 \setminus 1 \xrightarrow{u} r, u, w) \upharpoonright u \rrbracket = S_w(S_r.Q_u) \wedge Q_u \wedge S_w.Q_u$	14.824
6. $\llbracket (G_2 \setminus 1 \xrightarrow{w} r, u, w) \upharpoonright w \rrbracket = S_r(Q_w \wedge S_u.Q_w) \wedge Q_w$	13.076
7. $\llbracket (G_3 \setminus 1 \xrightarrow{r} r, u, w) \upharpoonright r \rrbracket = Q_r \wedge S_w(S_u.Q_r \wedge Q_r)$	23.084
8. $\llbracket (G_3 \setminus 1 \xrightarrow{u} r, u, w) \upharpoonright u \rrbracket = S_r.Q_u \wedge Q_u \wedge S_r(S_w.Q_u)$	19.786
9. $\llbracket (G_3 \setminus 1 \xrightarrow{w} r, u, w) \upharpoonright w \rrbracket = S_u(S_r.Q_w \wedge Q_w) \wedge Q_w$	4.241

We define a simplified notation in Table 3, l for **local**, r for **rhea**, u for **uniprot** and w for **wikidata**. We show federated query elapsed time with different communication sequences. Suppose we take the first case in Table 3 which is generating a query template for **rhea**, $Q'_r = Q_r \wedge S_u.Q_r \wedge S_w.Q_r$. Based on Algorithm 1 and an input data $\llbracket (G_1 \setminus 1 \xrightarrow{r} r, u, w) \upharpoonright r \rrbracket$, global type G_1 is reduced by telling every participant that **local** will acts as **rhea** under reduction of $G_1 \setminus \text{local} \xrightarrow{\text{rhea}} \{\text{rhea}, \text{uniprot}, \text{wikidata}\}$. This syntactically will produce G_{rhea} . Global type is projected to **rhea**, $G_{\text{rhea}} \upharpoonright \text{rhea} = T_{\text{rhea}}$, before being translated into a query with Table 1.

The first query template in Table 3 has an overhead communication exchange between participants, as shown in Fig. 7a and local type definitions. Based on the query execution sequence in Fig. 4, a result from **rhea** must join with **uniprot** to retrieve reactions and related human enzymes. Next, a result must join with **wikidata** to find interactions between those previous enzymes and drug descriptions. Ideally, **uniprot** sends the join result directly to **wikidata** but local type

$T_u = \mathbf{r}![U];\mathbf{end}$ defines a communication that **uniprot** can send only to **rhea**. The communication **uniprot** to **wikidata** passes through **rhea** first and continues to **wikidata** together with federated query $S_w.Q_r$. Finally, **wikidata** sends result back to **rhea**. This causes a longer execution time because **uniprot** cannot send directly to **wikidata** and must pass **rhea** first.




Fig. 7. Federated service implementation between two global protocols

In contrast, case number 9 makes the most efficient communication as shown in Fig. 7b, as the join result is sent directly to the next target participant without an intermediary. According to local type $T_r = \mathbf{u}![U];\mathbf{end}$, **rhea** sends directly to **uniprot**. Next, **uniprot** sends a query result to **wikidata** as a final destination, $T_u = \mathbf{r}?[U];\mathbf{w}![U];\mathbf{end}$.

6 Conclusion

The current SQT (SPARQL Query Template) algorithm is a convenient and faster way of generating queries than writing them manually. The use of global types helps identify RDF sources by checking the interaction between participants. Global types ensure that executing queries on all projections will provide a complete result set. Moreover, the communication sequence described in global types can also be used to check for any overhead communication during a federated query. These benefits can significantly reduce network costs and ultimately make federated queries faster. In addition, we introduce a broadcast reduction for global types and prove its soundness against the original semantics of global types.

Acknowledgements. The first author is funded by the LPDP Indonesian Endowment Fund for Education from the Ministry of Finance of the Republic of Indonesia. This research is funded in part by EPSRC EP/T006544/2, EP/K011715/1, EP/K034413/1, EP/L00058X/1, EP/N027833/2, EP/N028201/1, EP/T014709/2, EP/V000462/1, EP/X015955/1, NCSS/EPSRC VeTSS and Horizon EU TaRDIS 101093006.

References

1. Mendes, P.N., Jakob, M., García-Silva, A., Bizer, C.: DBpedia spotlight: shedding light on the web of documents. In: Proceedings of the 7th International Conference on Semantic Systems, I-Semantics 2011, pp. 1–8. Association for Computing Machinery, New York (2011). <https://doi.org/10.1145/2063518.2063519>
2. Bateman, A., et al.: UniProt: the universal protein knowledgebase in 2023. *Nucleic Acids Res.* **51**(D1), D523–D531 (2023)
3. Bansal, P., et al.: Rhea, the reaction knowledgebase in 2022. *Nucleic Acids Res.* **50**(D1), D693–D700 (2022). <https://academic.oup.com/nar/article/50/D1/D693/6424769>
4. Aleksander, S.A., et al.: The Gene Ontology knowledgebase in 2023. *GENETICS* **224**(1) (2023). <https://academic.oup.com/genetics/article/doi/10.1093/genetics/iyad031/7068118>
5. Bauer-Mehren, A., Bundschuh, M., Rautschka, M., Mayer, M.A., Sanz, F., Furlong, L.I.: Gene-disease network analysis reveals functional modules in mendelian, complex and environmental diseases. *PLoS ONE* **6**(6), e20284 (2011). <https://dx.plos.org/10.1371/journal.pone.0020284>
6. Mendez, D., et al.: ChEMBL: towards direct deposition of bioassay data. *Nucleic Acids Res.* **47**(D1), D930–D940 (2019). <https://academic.oup.com/nar/article/47/D1/D930/5162468>
7. Potter, A., Motik, B., Nenov, Y., Horrocks, I.: Distributed RDF query answering with dynamic data exchange. In: Groth, P., et al. (eds.) ISWC 2016. LNCS, vol. 9981, pp. 480–497. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46523-4_29
8. Sidirourgos, L., Goncalves, R., Kersten, M., Nes, N., Manegold, S.: Column-store support for RDF data management. *Proc. VLDB Endow.* **1**(2), 1553–1563 (2008). <https://dl.acm.org/doi/10.14778/1454159.1454227>
9. Abadi, D.J., Marcus, A., Madden, S.R., Hollenbach, K.: SW-Store: a vertically partitioned DBMS for Semantic Web data management. *VLDB J.* **18**(2), 385–406 (2009). <https://doi.org/10.1007/s00778-008-0125-y>
10. Atre, M., Chaoji, V., Zaki, M.J., Hendler, J.A.: Matrix “Bit” loaded. In: Proceedings of the 19th International Conference on World Wide Web, pp. 41–50. ACM, New York (2010). <https://dl.acm.org/doi/10.1145/1772690.1772696>
11. Neumann, T., Weikum, G.: The RDF-3X engine for scalable management of RDF data. *VLDB J.* **19**(1), 91–113 (2010). <https://doi.org/10.1007/s00778-009-0165-y>
12. Rohloff, K., Schantz, R.E.: Clause-iteration with MapReduce to scalably query data graphs in the SHARD graph-store. In: Proceedings of the 4th International Workshop on Data-Intensive Distributed Computing, DIDC 2011, pp. 35–44. ACM (2011)
13. Zhang, X., Chen, L., Tong, Y., Wang, M.: EAGRE: towards scalable I/O efficient SPARQL query evaluation on the cloud. In: Proceedings of the International Conference on Data Engineering, pp. 565–576. IEEE (2013)

14. Hassan, M., Bansal, S.: S3QLRDF: Distributed SPARQL Query Processing Using Apache Spark—A Comparative Performance Study, vol. 41. Springer, Cham (2023). <https://doi.org/10.1007/s10619-023-07422-4>
15. Zeng, K., Yang, J., Wang, H., Shao, B., Wang, Z.: A distributed graph engine for web scale RDF data. *Proc. VLDB Endow.* **6**(4), 265–276 (2013)
16. Gurajada, S., Seufert, S., Miliaraki, I., Theobald, M.: TriAD: a distributed shared-nothing RDF engine based on asynchronous message passing. In: *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pp. 289–300 (2014)
17. Daga, E., Asprino, L., Mulholland, P., Gangemi, A.: Facade-X: an opinionated approach to SPARQL anything. In: Alam, M., Groth, P., de Boer, V., Pellegrini, T., Pandit, H.J. (eds.) *Further with Knowledge Graphs*, vol. 53, pp. 58–73. IOS Press (2021). <http://oro.open.ac.uk/78973/>
18. Hernawan, A., Sunarwidhi, A., Prasedya, E., Widyastuti, S.: A generalization SPARQL federated query: an initial step towards machine-readable web of data for halal food products. In: *IOP Conference Series: Earth and Environmental Science*, vol. 913, no. 1, p. 012040 (2021). <https://iopscience.iop.org/article/10.1088/1755-1315/913/1/012040>
19. Quilitz, B., Leser, U.: Querying distributed RDF data sources with SPARQL. In: Bechhofer, S., Hauswirth, M., Hoffmann, J., Koubarakis, M. (eds.) *ESWC 2008*. LNCS, vol. 5021, pp. 524–538. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-68234-9_39
20. Schwarte, A., Haase, P., Hose, K., Schenkel, R., Schmidt, M.: FedX: a federation layer for distributed query processing on linked open data. In: Antoniou, G., et al. (eds.) *ESWC 2011*. LNCS, vol. 6644, pp. 481–486. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-21064-8_39
21. Khan, Y., et al.: SAFE: SPARQL federation over RDF data cubes with access control. *J. Biomed. Semant.* **8**(1), 5 (2017). <http://jbiomedsem.biomedcentral.com/articles/10.1186/s13326-017-0112-6>
22. Görlitz, O., Staab, S.: SPLENDID: SPARQL endpoint federation exploiting VOID descriptions. In: *Proceedings of the Second International Conference on Consuming Linked Data, COLD 2011*, vol. 782, pp. 13–24. CEUR-WS.org, Aachen (2011)
23. Troumpoukis, A., et al.: Developing a benchmark suite for semantic web data from existing workflows. In: Fundulaki, I., Krithara, A., Ngomo, A.N., Rentoumi, V. (eds.) *Proceedings of the Workshop on Benchmarking Linked Data (BLINK 2016) co-located with the 15th International Semantic Web Conference (ISWC)*, Kobe, Japan, 18 October 2016. *CEUR Workshop Proceedings*, vol. 1700. CEUR-WS.org (2016). <https://ceur-ws.org/Vol-1700/paper-04.pdf>
24. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. *J. ACM* **63**(1), 1–67 (2016)
25. Yoshida, N.: Programming language implementations with multiparty session types. In: de Boer, F.S., Damiani, F., Hähnle, R., Johnsen, E.B., Kamburjan, E. (eds.) *Active Object Languages: Current Research Trends. Lecture Notes in Computer Science*, vol. 14360, pp. 147–165. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-51060-1_6
26. Coppo, M., Dezani-Ciancaglini, M., Padovani, L., Yoshida, N.: A gentle introduction to multiparty asynchronous session types. In: Bernardo, M., Johnsen, E. (eds.) *SFM 2015*. LNCS, vol. 9104, pp. 146–178. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-18941-3_4

27. Yoshida, N., Gheri, L.: A very gentle introduction to multiparty session types. In: Hung, D.V., D'Souza, M. (eds.) ICDCIT 2020. LNCS, vol. 11969, pp. 73–93. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-36987-3_5
28. Cyganiak, R.: A relational algebra for SPARQL. HP (2005). <http://www.hpl.hp.com/techreports/2005/HPL-2005-170.pdf>
29. Salas, J., Hogan, A.: Semantics and canonicalisation of SPARQL 1.1. *Semant. Web* **13**(5), 829–893 (2022)

Author Queries

Chapter 1

Query Refs.	Details Required	Author's response
AQ1	This is to inform you that corresponding author has been identified as per the information available in the Copyright form.	