

# Planning Ahead: Stream-Driven Linked-Data Access under Update-Budget Constraints

Shen Gao<sup>1</sup>, Daniele Dell’Aglio<sup>2</sup>, Soheila Dehghanzadeh<sup>3</sup>,  
Abraham Bernstein<sup>1</sup>, Emanuele Della Valle<sup>2</sup>, Alessandra Mileo<sup>3</sup>

<sup>1</sup>Department of Informatics, University of Zurich, Switzerland

<sup>2</sup>DEIB, Politecnico di Milano, Italy

<sup>3</sup>INSIGHT Research Center, NUI Galway, Ireland

**Abstract.** Data stream applications are becoming increasingly popular on the web. In these applications, one query pattern is especially prominent: a join between a continuous data stream and some background data (BGD). Oftentimes, the target BGD is large, maintained externally, changing slowly, and costly to query (both in terms of time and money). Hence, practical applications usually maintain a local (cached) view of the relevant BGD. Given that these caches are not updated as the original BGD, they should be refreshed under realistic budget constraints (in terms of latency, computation time, and possibly financial cost) to avoid stale data leading to wrong answers. This paper proposes to model the join between streams and the BGD as a bipartite graph. By exploiting the graph structure, we keep the quality of results good enough without refreshing the entire cache for each evaluation. We also introduce two extensions to this method: first, we consider a continuous join between recent portions of a data stream and some BGD to focus on updates that have the longest effect. Second, we consider the future impact of a query to the BGD by proposing to delay some updates to provide fresher answers in future. By extending an existing stream processor with the proposed policies, we empirically show that we can improve result freshness by 93% over baseline algorithms such as Random Selection or Least Recently Updated.

## 1 Introduction

Real-time processing of massive, dynamically generated stream-data has become increasingly popular on the Web [18]. In stream processing, one common task is to enrich the streams with external background data (BGD). This kind of tasks has to deal with two V’s of “Big Data” at the same time: *Velocity*, the rapidly changing nature of the stream data; *Variety*, integrating data from different sources<sup>1</sup>. RDF Stream Processing (RSP) has provided necessary languages to declare this task. Current RSP languages, such as C-SPARQL [3], SPARQL<sub>stream</sub> [4], and CQELS-QL [16], support complex queries that involve both streams and remote BGD. However, these RSP engines are not optimized for remote BGD access. Usually, they continuously fetch BGD to match newly arrived stream data ignoring the communication and potential financial cost of such operations. To improve BGD access, RSP engines may adopt local views (or caches), as done in database systems [9]. However, the remote BGD is not always static. Indeed, even in the mostly static linked-data realm, information changes [13]. Hence, the *freshness* of local

<sup>1</sup> <http://www.ibmbigdatahub.com/infographic/four-vs-big-data>

views in the RSP engine degrades over time as updates in BGD do not propagate to the local view. To address this problem, RSP engines have to *maintain* the local view, by identifying the out-of-date (or *stale*) data items and replacing them with the up-to-date (or *fresh*) values retrieved from the remote. Examples of such updating behavior include the identification of opinion makers in social media based on a stream of posts and (slowly-changing) contact-networks as BGD or traffic prediction based on position data fetched from mobile phones.

Maintaining a local view can take time. Given that a federated query evaluation can spend up to 95% of its time on accessing remote data [19], query evaluation under response time constraints becomes a major challenge. To ensure a certain response time, only a limited number of remote accesses can be allowed. Additionally, BGD providers may impose constraints such as API rate limits, e.g., Twitter<sup>2</sup>. Lastly, other communication and financial constraints may have to be considered, since accessing BGD can cost money, computation power or energy (in both the RSP engine and the remote service). Returning to the above examples, computing updated network metrics for opinion makers is computationally expensive, and fetching location updates from cell phones burdens scarce battery power. In this paper, we consider these constraints as a limited *budget* that restricts the number of BGD accesses. We study the problem of how to utilize the limited budget so that it can provide fresher response to the query.

To optimally manage BGD accesses under realistic budget constraints, this paper proposes to allocate budget only to carefully selected “important” data that could lead to more *fresh* join results. Our algorithms exploit characteristics of the join between the stream and the BGD to improve the response freshness. Specifically, our contribution is threefold. First, we propose an algorithm that employs a bipartite graph to model the join selectivity between stream and BGD. It favors the update of data items with a higher selectivity within a budget constraint. This problem decomposes to two scenarios: one can be tackled with a local optimal approach; a second is NP-hard requiring a greedy heuristic approach. This encodes Hypothesis **H1: A maintenance processes exploiting join selectivity improves response freshness**. Second, we extend the above model to favor data items that have a longer impact on the response freshness, which leads to hypothesis **H2: Leveraging the definition of the sliding window and BGD change frequencies can improve response freshness**. Third, we explore the trade-off between the current and future importance of data elements. We present an algorithm that exploits the change frequencies, join selectivity, and the sliding window all together to delay some current refreshes in favor of future, more important ones. It encodes hypothesis **H3: Considering both current and future evaluations for budget allocation can further improve response freshness**.

*Outline:* Section 2 introduces some background of RSP and BGD access. Section 3 reviews related work. Section 4 formalizes the problem. Our solutions and their optimization are in Section 5. Section 6 provides evaluation results of our hypotheses on both real and synthetic data sets.

---

<sup>2</sup> <https://dev.twitter.com/rest/public/rate-limiting>

## 2 Background

An **RDF stream**  $S$  is a potentially unbounded sequence of timestamped informative units  $(d_i, t_i)$  ordered by the temporal dimension, where  $t_i$  is the timestamp (as in [3,4,16]), we consider the time as discrete) and  $d_i$  is a set of RDF statements. An RDF statement is a triple  $(s, p, o) \in (I \cup B) \times I \times (I \cup B \cup L)$ , where  $I$ ,  $B$ , and  $L$  identify the sets of IRIs, blank nodes and literals, respectively. An **RDF term** is an element of the set  $T = I \cup B \cup L$ .

**RSP Query Languages** [3,4,7,16] extend SPARQL<sup>3</sup> with operators to cope with streams. They enable the registration of queries over RDF streams. RSP queries are evaluated in a continuous fashion, i.e., results are computed at different time instances as the data flows in the streams. Given a query  $q$ , the answer  $Ans(q)$  is a stream, to which the results of the evaluations are appended. This work focuses on the RSP query languages that support the **time-based sliding window** operator  $\mathbb{W}$ , which is defined through the parameter  $\omega$ , the width, and  $\beta$ , the slide, and generates a sequences of fixed windows, i.e., portions of  $S$  in a time interval  $(o, c]$  [3,4,16]. Given a time-based sliding window and two generated consecutive windows  $W_i$  and  $W_{i+1}$ , defined in  $(o_i, c_i]$  and  $(o_{i+1}, c_{i+1}]$ , two constraints hold:  $c_i - o_i = c_{i+1} - o_{i+1} = \omega$  and  $o_{i+1} - o_i = \beta$ .

Let  $V$  be a set of variables (disjoint with  $I$ ,  $B$  and  $L$ ), graph patterns are expressions defined recursively as: 1) a basic graph pattern, i.e., a set of triple patterns  $(ts, tp, to) \in (I \cup B \cup V) \times (I \cup V) \times (I \cup B \cup L \cup V)$ , is a graph pattern; 2) let  $P_1$  and  $P_2$  be graph patterns,  $P_1$  *JOIN*  $P_2$  or  $P_1$  *UNION*  $P_2$  is a graph pattern; 3) let  $P$  be a graph patterns and  $u \in I \cup V$ , *SERVICE*  $u$   $P$  or *WINDOW*  $u$   $P$  is also a graph pattern. Other graph pattern expressions are possible (e.g. *OPTIONAL*, *FILTER*) but are not presented for the sake of space<sup>3</sup>.

Like SPARQL, the evaluation semantics of RSP Query Languages rely on the notion of **solution mapping**, i.e., a partial function that maps variables to RDF terms, i.e.,  $\mu : V \rightarrow T$ . A full formalization of RSP Query Languages is in [7]. We briefly describe the semantics of *WINDOW*, *SERVICE*, and *JOIN* in RSP Query Languages. Evaluating a **WINDOW** clause results the content of a sliding window, similarly to what *GRAPH* does in SPARQL, which refers to the content of a named graph in the data set. The **SERVICE** retrieves mappings from SPARQL endpoints by submitting a graph pattern [2]. **JOIN** can be formally defined as: let  $dom(\mu) \subset V$  be the set of variables mapped by  $\mu$ , two mappings  $\mu_1$  and  $\mu_2$  are **compatible** (denoted with  $\mu_1 \sim \mu_2$ ) if they assign the same values to the common variables, i.e.,  $\forall v \in dom(\mu_1) \cap dom(\mu_2), \mu_1(v) = \mu_2(v)$ . We name **joining variables** the elements in  $dom(\mu_1) \cap dom(\mu_2)$ .

As explained, this paper focuses on queries containing the graph pattern:

$$(WINDOW\ u^W\ P^W)\ JOIN\ (SERVICE\ u^S\ P^S),$$

where  $P^W$  and  $P^S$  are two graph patterns that share one or more variables,  $u^S$  is the address of a service BGD in remote and  $u^W$  is an IRI denoting a sliding window operator  $\mathbb{W}$  defined through  $\omega, \beta$  and applied to a stream  $S$ .

**Local view.** Existing RDF stream engines leverage a nested loop join strategy to fetch data from BGD. It follows that evaluating the above graph pattern can be expensive: each request to BGD has a latency, computational and, possibly, financial cost. In the SPARQL

<sup>3</sup> Cf. <https://www.w3.org/TR/sparql11-query/> for additional reference.

endpoint of our experiments (see Section 6), each invocation takes  $4.6ms$ . Hence, during one second, it can only accommodate up to 200 requests. In real scenarios, SPARQL endpoints are exposed over Internet, and each quest can take more than  $500ms$  [19].

For this reason, we previously proposed to use a local view  $\mathcal{R}$  to store the result of  $P^S$  in the RSP engine [5].  $\mathcal{R}$  stores the results of the SERVICE clause so that the engine computes the results of the query without invoking the SPARQL endpoint of BGD at each evaluation. However, given that the content of BGD changes over time, the mappings in  $\mathcal{R}$  become outdated, and the evaluation of the SERVICE clause produces different solution mappings can leading to wrong results. We consider these outdated results invalid. Therefore, each mapping  $\mu^{\mathcal{R}} \in \mathcal{R}$  can be classified as *fresh* or *stale*:  $\mu^{\mathcal{R}}$  is *fresh* at time  $t$ , if it is contained in the result set by evaluating the SERVICE clause over BGD at  $t$ ; it is *stale* otherwise (i.e., if BGD changes, it produces different results when evaluating of the SERVICE clause over  $\mu^{\mathcal{R}}$  and the remote BGD). In the following, we assume that mappings in BGD change with fixed intervals. This happens, e.g., in data warehouses, where updates are scheduled, or in data generated by sensors or automatic processes, where data is updated with fixed interval. As in [6], we define the freshness of an answer  $Ans(q)$  as  $\frac{|fresh(Ans(q))|}{|Ans(q)|}$ .

**Maintenance process.** To ensure the freshness of the local view over time, we introduce a maintenance process  $MP$  that refreshes a portion of  $\mathcal{R}$ .  $MP$  selects a set of mappings  $\mathcal{E} \subseteq \mathcal{R}$  to refresh within each evaluation of the queries over BGD. The design of  $MP$  is the key to the freshness of  $Ans(q)$ : if the process correctly identifies the stale mappings and puts them in  $\mathcal{E}$ , then both the freshness of  $\mathcal{R}$  and  $Ans(q)$  increase. Note, however, that if the number of refresh queries sent to BGD is too high, the presence of  $\mathcal{R}$  does not bring any advantage. In practice,  $MP$  has to consider (i) Quality of Service requirements associated to the query, e.g., responsiveness; (ii) system reactivity, e.g., each evaluation should terminate before the next one starts; (iii) constraints imposed by the BGD providers on the number of requests during a time interval. We capture these aspects by introducing a notion of **refresh budget** value  $\Gamma$ , defined as the number of refresh queries that can be sent to BGD in a given time period under the above constraints. In our Hypotheses 1 and 2, we assume that  $\Gamma$  is evenly distributed over  $n$  evaluations, when the stream rate is stable. In Hypothesis 3, in order to deal with unstable stream rate, we relax such assumption by allowing to move budget between evaluations. We use  $\gamma = \lfloor \Gamma/n \rfloor$  to denote the maximum refresh budget available in one evaluation.

### 3 Related Work

Traditional databases usually materialize remote BGD locally. Sophisticated optimizations of retrieving remote data on-demand have been introduced to improve availability, scalability and query processing performance [8,9,14]. The drawback of materialization is that local data becomes stale when the remote data changes. Those works are neither in stream processing context, nor considering budget constraints on remote access.

In Complex Event Processing (CEP), the incoming events not only need to be matched with specified event patterns, but also need to be enriched [10,22]. During enrichment, it usually needs to access remote BGD through APIs defined by service providers [11]. These API providers usually apply constraints on the number of accesses to restrict the massive loads of requests, as the computation and communication costs

involved are shown to be intensive. Given the repetitive nature of the access to BGD [17], caching techniques can improve on response latency. However, when a cache becomes outdated, refreshing it raises the trade-off between latency and freshness [1]. More remote accesses could provide fresher response, but take longer time. Authors in [14] addresses this trade-off in a web setting, where updates of the remote BGD are pushed into the system [9]. However, this work does not consider the constraints of service providers or the view maintenance without updates being pushed into the system.

In RDF processing, SPARQL 1.1 standardizes the access to remote BGD by introducing the federated extension [2] and the `SERVICE` clause. Broadly, there are two ways of accessing BGD: either one pulls the whole data into the query processor [15] or one ‘federates’ query-execution and transfers the data for individual operations over the network [12], defining new join strategies that can efficiently process both local and remote data [15]. Extending static RDF processing, RSP technologies deal with data of different velocity and variety. C-SPARQL [3] performs query matching on subsets of the information flow defined by windows. CQELS [16] implements its native query operators, which can be adaptively optimized to improve performance. MorphStream [4] allows querying relational data streams over a set of stream-to-ontology mappings. INSTANS [20] is a semantic event processing platform, which compiles a query into a Rete-like structure. All those systems are optimized for processing streams. They support the `SERVICE` clause as described above but do not consider budget-constrained updates in the local view. Hence, our solution is orthogonal to these and other RSP engines.

Our previous work [5] studied the maintenance process of local view for queries where each mapping in the `WINDOW` clause joins with exactly one mapping in the `SERVICE` one. In this paper, we tackle a more general join relationship between `WINDOW` and `SERVICE` clauses, i.e., we extend the 1:1 join relationship to M:N and propose a flexible budget allocation method that further improves the maintenance process.

## 4 Problem Definition

Given the graph pattern expression  $P^S$  in the `SERVICE` clause, we define two sets of variables: first,  $V^{SR} \subset \text{var}(P^S)$  contains the variables in  $\text{var}(P^S)$  that are related to the changing part in BGD. In other words,  $V^{SR}$  captures the dynamicity of BGD and contains the information needed to construct the refresh queries that are sent to remote BGD. Second,  $V^{SN}$  are the common variables that join the  $P^S$  and  $P^W$  clauses, i.e.,  $V^{SN} = \text{var}(P^S) \setminus V^{SR}$ . We model the relationship between  $V^{SR}$  and  $V^{SN}$  as a bipartite graph. The maintenance process  $MP$  exploits the graph to identify the candidate set  $\mathcal{E}$  for refreshing. The  $MP$  builds a bipartite graph (*maintenance graph*, Figure 1) out of  $\mathcal{C}$ , which is a subset of  $\mathcal{R}$ . Mappings in  $\mathcal{C}$  are (1) stale and (2) belong to the candidate set of the current window (i.e., they have compatible mappings in the result set  $\Omega^W$  of the `WINDOW` clause). The maintenance graph has signature  $G^{\mathcal{C}} = (\Omega^{SN}, \Omega^{SR}, E)$ , where  $\Omega^{SN}$  ( $\Omega^{SR}$ ) is the set of mappings with domain  $V^{SN}$  ( $V^{SR}$ ), and  $E$  are the mappings  $\mu^{\mathcal{R}}$  in  $\mathcal{C}$ , modeled as edges connecting elements of  $\Omega^{SN}$  and  $\Omega^{SR}$ .

Different subqueries in  $P^S$  have different optimization goals. In this work, we consider: 1)  $P^S$  is a Basic Graph Pattern (BGP) query; 2)  $P^S$  is an aggregate query<sup>4</sup>.

<sup>4</sup> We assume that the aggregation is performed locally in the query processor and not in the remote BGD. It happens, e.g., when BGD is not SPARQL 1.1 compliant.

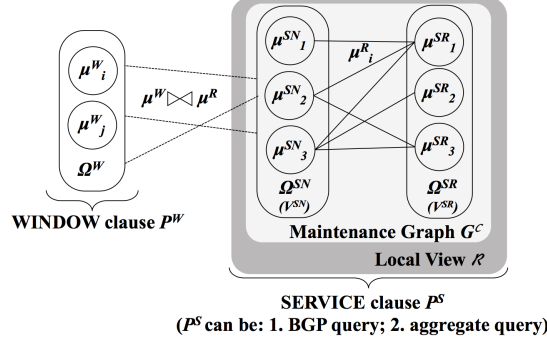


Fig. 1: WINDOW/SERVICE clauses and the Maintenance graph

**Case 1:  $P^S$  is a BGP query.** By differentiating  $V^{SR}$  and  $V^{SN}$ , we split  $\mu^R$  into two mappings  $\mu = \mu^{SR} \cup \mu^{SN}$  such that  $\text{dom}(\mu^{SR}) \subseteq V^{SR}$  and  $\text{dom}(\mu^{SN}) \subseteq V^{SN}$ . As  $P^S$  is a BGP query, each mapping  $\mu_k^R$  consists a  $\mu_i^{SN}$  and a  $\mu_j^{SR}$ . Updating one  $\mu_j^{SR}$  can ensure all its corresponding  $\mu_k^R$  are fresh. As an example, consider the graph in Figure 1, where  $\mathcal{C} = \{\mu_1^R, \dots, \mu_6^R\}$ .  $\Omega^{SR}$  contains the mappings with the variables in  $V^{SR}$ , i.e.,  $\{\mu_1^{SR}, \mu_2^{SR}, \mu_3^{SR}\}$  (on the right);  $\Omega^{SN}$  contains the other mappings, i.e.,  $\{\mu_1^{SN}, \mu_2^{SN}, \mu_3^{SN}\}$  (in the middle). The mappings in  $\mathcal{R}$  are encoded as the edges in  $E$  (e.g.,  $(\mu_1^{SN}, \mu_1^{SR})$  represents  $\mu_1^R$ ). Updating  $\mu_1^{SR}$  will make all its three corresponding mappings to be fresh:  $(\mu_1^{SN}, \mu_1^{SR})$ ,  $(\mu_2^{SN}, \mu_1^{SR})$ , and  $(\mu_3^{SN}, \mu_1^{SR})$ . Given  $\Omega^W$  (on the left) as the solution of the WINDOW clause and  $\gamma$  as the refresh budget at the current iteration, the maintenance process can be summarized as: what is the subset of  $\Omega^{SR}$  to refresh can maximize the number of fresh join results between  $\mu^W$  and  $\mu^R$ ? Formally, it can be modeled as the following optimization problem:

$$\text{Sub. } u_j^{SR} = 0 \text{ or } 1 \quad \forall j = [1, |\Omega^{SR}|] \quad (1)$$

$$\sum_{j=1}^{|\Omega^{SR}|} u_j^{SR} \leq \gamma \quad (2)$$

$$f_i^{SN} = \sum \mu_j^{SR} \quad \forall \mu_j^{SR} : (\mu_i^{SN}, \mu_j^{SR}) \in E \quad \forall i = [1, |\Omega^{SN}|] \quad (3)$$

$$c_i^{SN} = |\{\mu^W : \mu^W \in \Omega^W \wedge \mu^W \text{ comp. with } (\mu_i^{SN}, \mu_j^{SR})\}| \quad \forall i = [1, |\Omega^{SN}|] \quad (4)$$

$$\text{Max. } \sum_{i=1}^{|\Omega^{SN}|} f_i^{SN} * c_i^{SN} \quad (5)$$

The optimization is subject to: in Formula (1), the value of  $u_j^{SR}$  shows whether the  $j$ -th stale mapping is updated ( $u_j^{SR} = 1$ ) or not ( $u_j^{SR} = 0$ ). The total number of updates is limited by  $\gamma$ , as in Formula (2). Formula (3) defines  $f_i^{SN}$  as the number of fresh mappings  $\mu_i^{SN}$  will have. Each  $\mu_i^{SN}$  may have several related  $\mu_j^{SR}$ . By summing all its refreshed  $\mu_j^{SR}$ , we have the total number of fresh mappings for  $\mu_i^{SN}$ . As discussed above, this is because each updated  $\mu_j^{SR}$  produces one fresh  $\mu_k^R$  ( $\mu_i^{SN}, \mu_j^{SR}$ ). Overall, Formula (1) to (3) give the total number of fresh  $\mu^R$  in the SERVICE clause. Since each  $\mu^R$  may have several compatible mappings in the WINDOW clause, Formula (4) introduce  $c_i^{SN}$  to represent the number of compatible mappings of  $\mu_k^R$  in the window. Finally, our optimization goal is to maximize the total number of join results between WINDOW and SERVICE clauses, which could be defined as the product of  $c_i^{SN}$  and  $f_i^{SN}$ , as shown in Formula (5).

**Case 2.  $P^S$  is an aggregate query.** In this case, the maintenance graph  $G^C$  is constructed as the previous case:  $\Omega^{SN}$  contains mappings with variables used for join, and  $\Omega^{SR}$  contains mappings with dynamic values. However,  $\Omega^{SR}$  in this case does not directly participate in the join, but are needed for aggregation.

Consider the example in Figure 1:  $\mathcal{C} = \{\mu_1^R, \mu_2^R, \mu_3^R\}$ :  $\mu_1^R$  contains the value of the aggregate variables by using the data stored in  $\mu_1^{SR}$ ;  $\mu_2^R$  has an aggregate computed from  $\mu_1^{SR}$  and  $\mu_3^{SR}$ ;  $\mu_3^R$  is computed from  $\mu_1^{SR}$ ,  $\mu_2^{SR}$  and  $\mu_3^{SR}$ . The edges in this case represent the mappings required to compute the aggregates, e.g.,  $(\mu_2^{SN}, \mu_1^{SR})$  and  $(\mu_3^{SN}, \mu_3^{SR})$  indicate that the mapping  $\mu_2^R$  should be computed by using both the fresh values of  $\mu_1^{SR}$  and  $\mu_3^{SR}$ . The maintenance problem is still to choose a subset of  $\mu^{SR}$  to maximize the fresh join results. However, in this case, updating one  $\mu_j^{SR}$  cannot ensure its corresponding  $\mu_k^R$  is fresh. To have a fresh  $\mu_k^R$ , we need all its related  $\mu^{SR}$  to be fresh. Therefore, the problem can be modeled as:

$$\text{Sub. } u_j^{SR} \leq 1 \quad \forall j = [1, |\Omega^{SR}|] \quad (6)$$

$$\sum_{j=1}^{|\Omega^{SR}|} u_j^{SR} \leq \gamma \quad (7)$$

$$f_i^{SN} = \prod \mu_j^{SR} \quad \forall \mu_j^{SR} : (\mu_i^{SN}, \mu_j^{SR}) \in E \quad \forall i = [1, |\Omega^{SN}|] \quad (8)$$

$$c_i^{SN} = |\{\mu^W : \mu^W \in \Omega_W \wedge \mu^W \text{ comp. with } (\mu_i^{SN}, \mu_j^{SR})\}| \quad \forall i = [1, |\Omega^{SN}|] \quad (9)$$

$$\text{Max. } \sum_{i=1}^{|\Omega^{SN}|} f_i^{SN} * c_i^{SN} \quad (10)$$

The constraints in Formula (6) and (7) are same with Case 1. Formula (8) uses  $f_i^{SN}$  to model the fact that the  $i$ -th mapping  $\mu_i^{SN}$  is fresh ( $f_i^{SN} = 1$ ) *iff* all its related  $\mu^{SR}$  are refreshed. For example, to have a fresh result of  $\mu_2^{SN}$ , both  $\mu_1^{SR}$  and  $\mu_3^{SR}$  have to be 1; otherwise,  $f_i^{SN} = 0$ . Formula (9) is same with Case 1. Finally, the objective function in Formula (10) maximizes the number of fresh mappings produced by the join.

Overall, both Case 1 and 2 can be treated as binary integer programming problems. However, Case 2 can be seen as an extension of the knapsack problem, which is NP-hard, e.g., packing a  $\mu^{SN}$  has a cost (the number of its  $\mu^{SR}$ ). We can only afford a certain number of  $\mu^{SR}$ , but need to maximize the number of  $\mu^{SN}$ . Furthermore, after choosing a  $\mu^{SN}$  and its related  $\mu^{SR}$  to pack, those  $\mu^{SR}$  might contribute to other  $\mu^{SN}$ . Therefore, choosing different  $\mu^{SR}$  will have different influence on the following decisions. Currently, there is no optimal way to find the best subset of  $\mu^{SR}$ .

## 5 Maintenance Algorithms

In this section, we propose a set of budget allocation algorithms. Section 5.1 proposes two greedy algorithms,  $\text{SBM}_{BGP}$  and  $\text{SBM}_{Agg}$ , for the problems in Case 1 and 2, respectively. They aim at maximizing the freshness of the current slide evaluation. Because the sliding window operator supplies information about future evaluations (i.e., elements stay in the window for different periods), Section 5.2 shows how to exploit this information to improve the maintenance process. Section 5.3 discusses how to flexibly manage the budget to optimize the overall response freshness. The basic idea is to uniformly allocate  $\Gamma$  to  $n$  evaluations (i.e.,  $\gamma = \lfloor \Gamma/n \rfloor$ ). When it is worthwhile, the solution trades the current remote accesses for the future ones.

### 5.1 Selectivity-Based Maintenance (SBM)

To maximize the number of fresh join results, we propose the  $\text{SBM}_{BGP}$  algorithm for Case 1, where  $P^S$  is a BGP query; and the  $\text{SBM}_{Agg}$  for Case 2, where  $P^S$  is an aggregate query. In both cases, we start from the maintenance graph  $G^C$  defined above. **SBM<sub>BGP</sub>**. The objective function of Case 1 (Formula (5)) aims at maximizing the number of fresh mappings produced by the join. Based on  $G^C$ ,  $\text{SBM}_{BGP}$  first computes a score for each mapping  $\mu^{SR} \in \Omega^{SR}$ , which represent the total number of the fresh join mappings that would be generated if  $\mu^{SR}$  is updated:

$$\text{score}_{SBM}(\mu^{SR}) = \sum_{\mu_i^{SN}: (\mu_i^{SN}, \mu^{SR}) \in E} c_i \quad (11)$$

Based on the selectivity of  $\mu^{SR}$ , the number of results it will have equals to the sum of each its connected  $\mu^{SN}$  times  $\mu^{SN}$ 's compatible mappings in the window. Then,  $\text{SBM}_{BGP}$  picks  $\mu^{SR}$  with the highest scores under the budget  $\gamma$  to refresh. If there are more than  $\gamma$  data with the same highest score, our algorithm chooses among them randomly.

**SBM<sub>Agg</sub>**. This case aims to maximize the number of fresh aggregate results. A mapping  $\mu^{SN}$  produces a fresh aggregate result only if all its connected  $\mu^{SR}$  are fresh. As discussed, finding the optimal set of  $\mu^{SN}$  is a NP-hard problem. We propose a heuristic algorithm:  $\text{SBM}_{Agg}$ . It tries to utilize the budget on those ‘‘cheap’’  $\mu^{SN}$ , which connects to less stale  $\mu^{SR}$ . Specifically,  $\text{SBM}_{Agg}$  picks the mapping  $\mu^{SN}$  with the smallest amount of connected  $\mu^{SR}$  and puts those  $\mu^{SR}$  in  $\mathcal{E}$ . Then,  $\mu^{SN}$  and the mappings in  $\mathcal{E}$  are removed from the maintenance graph  $G^C$ , and a new iteration starts again. It ends when  $\gamma$  elements have been moved into  $\mathcal{E}$ . If the budget left  $\gamma'$  is less than  $\mu^{SN}$ , we will randomly choose  $\gamma'$  amount of stale  $\mu^{SR}$ .

### 5.2 The Impact-Based Maintenance (IBM)

The two SBM algorithms maximize the freshness of the current evaluation, but do not consider future evaluations. As shown in [5], a maintenance process  $MP$  can take into account the sliding window and the changing frequency of the background data to have a prediction on what will be stale in future. We combine this idea with SBM to improve the performance of  $MP$ .

Before presenting the solution, we first introduce the concept of ranking data by a score based on two properties from [5], which quantify the impact of a mapping in future window evaluations. Consider a set of solution mappings  $\Omega^W$  resulted from the evaluation of a WINDOW clause and a local view  $\mathcal{R}$ , where each mapping in  $\Omega^W$  can have *only one* compatible mapping in  $\mathcal{R}$ .

The first property is the *remaining lifetime*, denoted with  $L$ . Let  $\mu^{\mathcal{R}}$  be a mapping in  $\mathcal{R}$ , and let  $\mu^W$  be its only compatible mapping in  $\Omega^W$  computed at time  $t$  in a sliding window  $\mathbb{W} = (S, \omega, \beta)$ . The  $L$  value of  $\mu^S$  at time  $t^{now}$  is computed as  $\lceil (t + \omega - t^{now}) / \beta \rceil$ . It represents the number of evaluations, in which  $\mu^S$  will be involved. For example, given a sliding window  $\mathbb{W} = (S, \omega = 150, \beta = 30)$  and a mapping  $\mu^W$  with timestamp  $t = 100$ , the  $L$  value of the compatible mapping  $\mu^{\mathcal{R}}$  at time 100 is  $L(\mu^{\mathcal{R}}, 100) = \lceil (100 + 150 - 100) / 30 \rceil = 5$ ; at time 160, it is  $L(\mu^{\mathcal{R}}, 100) = \lceil (100 + 150 - 160) / 30 \rceil = 3$ . The second property is the *number of evaluations before the next expiration*, denoted with  $B$ . Given a stale mapping  $\mu^{\mathcal{R}}$ ,  $B$  represents



the number of evaluations that  $\mu^{\mathcal{R}}$  would be fresh, if refreshed now.  $B$  is computed as  $B(\mu^{\mathcal{R}}, t^{now}) = \lceil (t^{exp} - t^{now}) / \beta \rceil$ , where  $t^{exp}$  is the next time on which  $\mu^{\mathcal{R}}$  would become stale.  $t^{exp}$  is processed by exploiting the change rate interval information of  $\mu^{\mathcal{R}}$ . At time  $t^{now} = 100$ , the value of  $B$  is  $B(\mu^{\mathcal{R}}, 100) = 3$ , i.e., if  $\mu^{\mathcal{R}}$  is refreshed now, it would remain fresh for the next three evaluations (evaluations at 100, 130, and 160; at 190,  $\mu^{\mathcal{R}}$  will be stale).

Now,  $L$  and  $B$  can be combined to assign a *score* to the elements in  $\mathcal{C}$  (i.e., the stale mappings in the local view currently involved). Intuitively, the *score* of the mapping  $\mu^{\mathcal{R}}$  represents *how many future correct results are attainable if  $\mu^{\mathcal{R}}$  is refreshed now*. The *score* of  $\mu^{\mathcal{R}}$  at time  $t^{now}$  is computed as  $score(\mu^{\mathcal{R}}, t^{now}) = \min\{L(\mu^{\mathcal{R}}, t^{now}), B(\mu^{\mathcal{R}}, t^{now})\}$ . If  $B(\mu^{\mathcal{R}}, t^{now}) < L(\mu^{\mathcal{R}}, t^{now})$ ,  $\mu^{\mathcal{R}}$  can generate at most  $B(\mu^{\mathcal{R}}, t^{now})$  fresh join mappings, before it becomes stale while remaining in the window; otherwise, it generates  $L(\mu^{\mathcal{R}}, t)$  fresh join results and will leave the window before it becomes stale. Based on this *score*, we extend the two SBM algorithms so that they also consider the future impact of a refresh. Given the maintenance graph  $G^{\mathcal{C}} = (\Omega^{SN}, \Omega^{SR}, E)$  as defined in Section 4 (M:N bipartite graph), the extensions, namely  $IBM_{BGP}$  and  $IBM_{Agg}$ , can cope with the stale mappings  $\mu^{SR}$  appearing in different mappings  $\mu^{\mathcal{R}}$  of the local view.

**$IBM_{BGP}$ .** We assign a score for the stale mappings in  $\Omega^{SR}$ , as with  $SBM_{BGP}$ . The formula proposed above for  $B$  is still valid for the elements in  $\Omega^{SR}$ . However,  $L$  cannot be directly associated with mappings in  $\Omega^{SR}$  because they are related to the mappings computed by the WINDOW clause  $\Omega^W$  through  $\Omega^{SN}$ .

$$L(\mu^{\mathcal{R}}, \mu^W, t^{now}) = \lceil (t_{\mu^W} + \omega - t^{now}) / \beta \rceil \quad (12)$$

$$score(\mu^{\mathcal{R}}, \mu^W, t^{now}) = \min\{L(\mu^{\mathcal{R}}, \mu^W, t^{now}), B(\mu^{SR}, t^{now})\} \quad (13)$$

$$score(\mu^{\mathcal{R}}, t^{now}) = \sum_{\mu^W \in c.w.\mu^{\mathcal{R}}} score(\mu^{\mathcal{R}}, \mu^W, t^{now}) \quad (14)$$

$$score_{IBM_{BGP}}(\mu^{SR}, t^{now}) = \sum_{\mu^{\mathcal{R}} = (\mu^{SN}, \mu^{SR}) \in E} score(\mu^{\mathcal{R}}, t^{now}) \quad (15)$$

$IBM_{BGP}$  associates the remaining lifetime  $L$  to the pair of compatible mappings  $(\mu^{\mathcal{R}}, \mu^W)$  as defined in Formula (12): the function considers the arriving time  $t_{\mu^W}$  of  $\mu^W$  as well, in order to cope with the fact that there are multiple compatible mappings for a  $\mu^{SR}$ . This extension allows defining a score for each pair  $(\mu^{\mathcal{R}}, \mu^W)$ , as in Formula (13). It represents the number of fresh mappings that are potentially generated by joining  $\mu^{\mathcal{R}}$  and  $\mu^W$  in the current and the following evaluations, if a  $\mu^{SR}$  is refreshed. Formula (14) computes the score of a mapping  $\mu^{\mathcal{R}}$  in the local view, which sums the scores of  $\mu^{\mathcal{R}}$  with compatible mappings in  $\Omega^W$ . Finally,  $IBM_{BGP}$  assigns the score to the mappings in  $\Omega^{SR}$  by Formula (15): it represents the total number of fresh join mappings that will be generated, if  $\mu^{SR}$  is refreshed. We select  $\gamma$  mappings of  $\mu^{SR}$  with the highest scores to refresh. Section 5.4 discusses why  $IBM_{BGP}$  is a local optimal solution.

**$IBM_{Agg}$ .** As discuss in Section 4, budget allocation in this case is a NP-hard problem. When future evaluations of the current data are considered, the complexity increases further, due to the additional level of combinatorial optimization. Therefore,  $IBM_{Agg}$  exploits a *score* function to improve the basic  $SBM_{Agg}$  algorithm. An aggregate value for a  $\mu^{SN} \in \Omega^{SN}$  is fresh only when all the required mappings  $\mu^{SR} \in \Omega^{SR}$  are fresh.

$$L(\mu^{SN}, t^{now}) = \lceil (\max_{t: \mu^W \in \Omega^W \wedge \mu^W \text{ c.w. } \mu^{SN}} \{t\} + \omega - t^{now}) / \beta \rceil \quad (16)$$

$$score_{IBM_{agg}}(\mu^{SN}, t^{now}) = \min \{L(\mu^{SN}, t^{now})\}_{\mu^{SR}: (\mu^{SN}, \mu^{SR}) \in E} \min \{B(\mu^{SR}, t^{now})\} \quad (17)$$

$IBM_{agg}$  computes the score of the mappings in  $\Omega^{SN}$  when two or more of them have the same lowest amount of connected  $\mu^{SR}$ . Specifically, Formula (16) computes the remaining lifetime of  $\mu^{SN}$ , which takes the most recent timestamp of the compatible mappings of a  $\mu^{SN}$  in  $\Omega^W$ . Formula (17) reports the function to compute the score, which considers two factors: (1)  $\mu^{SN}$  will continue to generate fresh mappings as long as all its related mappings  $\mu^{SR}$  are fresh; (2) their compatible mappings of  $\mu^{SN}$  still remain in the window.

### 5.3 Flexible Budget Allocation (FBA)

Above solutions only consider the fixed amount of refresh budget  $\gamma$  assigned in the current evaluation. However, fixing  $\gamma$  may be inefficient as the number of refresh requests changes over time. Saving current budget for future updates may improve result freshness, if the future ones can generate more results.<sup>5</sup> The semantics of the sliding window allow inferring how long each element in the current window will be involved in future joins. We propose FBA to allocate the refresh budget by considering both current and future evaluations. Specifically, FBA iterates from the current to the future  $\omega/\beta$  slides (window length/slide length). At each iteration, it identifies the maintenance graph  $G_i^C$  and the stale data  $\Omega_i^{SR}$ . It calculates the number of future fresh results for each  $\mu^{SR}$  in every  $\Omega_i^{SR}$  at their corresponding evaluation time and orders  $\mu^{SR}$  by their scores. FBA allocates total  $n \times \gamma$  budgets to the Top- $(n \times \gamma)$   $\mu^{SR}$  with the largest scores. Note that this set contains both current and future stale  $\mu^{SR}$ . If the number of  $\mu^{SR}$  in the current evaluation is less than  $\gamma$ , it means FBA delays the budgets of current  $\mu^{SR}$  to some future ones.

### 5.4 Discussion

**SBM<sub>BGP</sub> and IBM<sub>BGP</sub> are optimal for Case 1.** For a BGP query, choosing the top- $\gamma$  data in  $\Omega^{SR}$  based on  $deg^{SR}$ , which is the number of  $\mu^{SR}$ 's associated elements in  $\Omega^{SN}$ . It gives the local optimal solution at the current time without considering the future impact of  $\Omega^{SR}$ . This is because the top- $\gamma$  of  $\Omega^{SR}$  is the set with the largest sum of  $deg^{SR}$ , since the sum of  $deg^{SR}$  exactly equals the number of fresh results. Therefore, SBM<sub>BGP</sub> gives the local optimal solution. The same reason applies to IBM<sub>BGP</sub>, where  $\gamma$  mappings with the largest *score* also gives the most results, as *score* accurately reflects the number of future results. Note that since the future elements in stream are not predictable (with certainty), there is no global optimal solution for BGP query.

**Complexity.** Both the SBM and IBM only consider data in the current evaluation. SBM<sub>BGP</sub>/IBM<sub>BGP</sub> visits each  $\mu^W$  and  $\mu^R$  to count the number of mappings and calculate scores for  $\mu^{SR}$ , which both take linear time of  $\mathcal{O}(|\Omega^W| + |\Omega^R| + |\Omega^{SR}|)$ . Then, choosing the Top- $\gamma$  mapping take  $\mathcal{O}(|\Omega^{SR}| \log |\Omega^{SR}|)$  time. SBM<sub>Agg</sub> takes  $\mathcal{O}(|\Omega^{SN}|^2 \log |\Omega^{SN}|)$  time, as whenever updating a  $\mu^{SN}$ , we have to update all its

<sup>5</sup> We acknowledge that not all types of budget can be saved for future (e.g., a fixed amount of bandwidth cannot be saved). Other types of budgets, such as a supplier charges per request, a limited data plan, or limited power can be saved.

related  $\mu^{SR}$ .  $IBM_{Agg}$ , as an extension of  $SBM_{Agg}$ , has the same complexity. FBA has the same time complexity as IBM, since they have the same way of ranking and choosing data to refresh, except that IBM chooses data only in the current slide; FBA does this for a fixed number of future slides.

## 6 Experiments

**Experiment environment.** We implemented the maintenance process in a real RSP system: C-SPARQL [3]. The system registers continuous federated queries with `WINDOW` and `SERVICE` clauses (as in Section 2) and continuously evaluates the query per window on the incoming stream. Each evaluation joins the content of the current `WINDOW` with the results of the `SERVICE` clause. For evaluating the `SERVICE` clause, we have implemented a local view in C-SPARQL to cache remote BGD data (as in Section 4). Before executing the `SERVICE` clause, different maintenance algorithms will select a candidate set  $\mathcal{E}$  from  $\Omega^{SR}$  to refresh. For each data in  $\mathcal{E}$ , the `SERVICE` clause will request its fresh value from the remote server. We used Fuseki 2.0.0 as the remote BGD server and ran it with the C-SRAPQL engine on the same machine. The delay of each remote access under this setting is much smaller than querying an actual remote server.

**Experiment data sets.** We employ a *real* data set and several *synthetic* data sets to investigate the performance of our solutions. The real data set was recorded from Twitter. The synthetic ones were constructed by resembling the real one, but using a generator that can alter its characteristics. Each data set broadly contains three kinds of data: the remote BGD, the local view  $\mathcal{R}$ , and the input stream. We discuss different parameters of our data sets below and report their values in each experiment.

**The remote BGD.** In BGD, data change according to each one’s own change interval  $ChR$ . In  $real_{twitter}$  data, we use the number of a user’s followers as the BGD. When a tweet mentions several users, the `SERVICE` clause will provide the number of followers for each mentioned user in Case 1 (BGP query); it will provide the sum of the followers for all mentioned users in Case 2 (aggregate query). We collected the follower number of 100 selected users every minute for four hours by using the Twitter search API [5]. We noticed that the distribution of  $ChR$  is highly skewed: only few users have a very dynamic changing number of followers, while others are stable. Roughly, it resembles a Beta distribution with  $\alpha=50$  and  $\beta=1$ . In *synthetic* data, our data generator outputs data with different  $ChR$ -distributions. The generator has two parameters: the skewness of the distributions and the correlation with the selectivity of data elements in the local view. The latter controls whether data that changes more frequently can have either a higher or a lower selectivity.

**The local view  $\mathcal{R}$ .** In  $\mathcal{R}$ , we model the relationship between  $\Omega^{SN}$  and  $\Omega^{SR}$  as a bipartite graph. Therefore, we are mainly interested in the  $deg^{SR}$  of  $\Omega^{SR}$ . After analyzing the  $real_{twitter}$  data set, we observed a skewed distribution of the  $deg^{SR}$  that can be modeled as a Zipf distribution with a skewness parameter of 0.2. In the *synthetic* data, we can tune two aspects of  $deg^{SR}$ : the skewness and the correlation with the stream/remote data.

**The stream and the sliding window.** By using the Twitter stream API, the *real* data set collects a stream of tweets that contains the mentions of the monitored Twitter users described above. The *synthetic* data set generates the streaming data through

a Poisson process [21]. To verify Hypothesis 3, the stream is generated with a non-homogeneous Poisson process, where the data arrival rate changes over time, e.g.,  $\lambda_i = 0.95\lambda_0 \cdot (i \bmod 2) + \lambda_0 \cdot ((i + 1) \bmod 2)$ , where  $\lambda_0$  is the initial expected arrival interval and  $i$  is incremented along the time. The input query has a sliding window length of  $\omega = 4$  seconds and slides every  $\beta = 1$  second. Each experiment has 50 evaluations, and the first 10% is used as a warm-up period.

We first use synthetic data sets to verify our hypotheses and study the performance of our algorithms. The performance and the computational overhead on the real data set are reported as well. The average response freshness is used as the Key Performance Indicator (KPI). As discussed in Section 2, it is the ratio of fresh results to total number of results, within each evaluation. The number of fresh results is acquired by comparing the current result set to the corresponding set acquired by the original C-SPARQL engine<sup>6</sup>, where all results are fresh, since it queries BGD without budget constraints.

**The baseline algorithms.** We choose two baseline algorithms: 1) Least Recently Update (LRU), which selects the least recently updated *stale* data from  $\mathcal{R}$ ; 2) Random (RAND), which randomly chooses *stale* data from  $\mathcal{R}$ . All algorithms pick at most  $\gamma$  (the refresh budget) candidates to refresh.

**Resulting synthetic data sets.** The default settings of the synthetic data set are:  $\Omega^{SN}$  and  $\Omega^{SR}$  in  $\mathcal{R}$  contains 50 data elements each. There are 1000 edges  $\mu^R$  between  $\Omega^{SN}$  and  $\Omega^{SR}$ . Each  $\mu^R$  randomly connects a pair of  $\mu^{SN}$  and  $\mu^{SR}$ . Every  $\mu^{SR}$  has a change interval *ChR* randomly chosen from [100, 3000] *ms*. A stream trace generated from a Poisson distribution decides the arrival time of each  $\mu^{SN}$ . For the Poisson distribution, each  $\mu^{SN}$  chooses its  $\lambda$  (the expected arrival interval) randomly from [1000, 2000] *ms*. The default budget  $\gamma$  is 10.

## 6.1 Verifying hypotheses H1 and H2.

H1 and H2 are tested together by comparing the response freshness among RAND, LRU, SBM and IBM in both subquery cases:

**Case 1.** In the four settings of Figure 2, both  $SBM_{BGP}$  and  $IBM_{BGP}$  greatly improve the response freshness of the baselines by 22%/43%/93% (Min/Average/Max). These different settings show how the performance improvement generalizes.

In Figure 2(a), we show the performance of using  $\mathcal{R}$  with different densities, i.e., the number of edges  $|\mu^R|$  in  $\mathcal{R}$  is set to be 500, 1000, 1500, and 2000. Note that 2500 ( $|\Omega^{SN}| \times |\Omega^{SR}|$ ) edges will form a fully connected  $\mathcal{R}$ . First, we observe that the performance of RAND and LRU remain roughly stable over different densities. The reason is that they select the refresh candidates  $\mathcal{E}$  “blindly” without considering  $deg^{SR}$ —the selectivity of  $\mu^{SR}$ . Therefore, the percentage of edges being updated remains the same for different densities. On the other hand, in higher densities, the performance improvement of  $SBM_{BGP}$  and  $IBM_{BGP}$  decreases a bit. The reason is that in a denser graph the difference of  $deg^{SR}$  among  $\mu^{SR}$  becomes less significant.  $SBM_{BGP}$  and  $IBM_{BGP}$  always choose the  $\mu^{SR}$  with the highest  $deg^{SR}$ , however, the percentage of the chosen  $\mu^{SR}$  to the total number of  $\mu^{SR}$  in  $G^C$  becomes smaller. Hence,  $SBM_{BGP}$  and  $IBM_{BGP}$  favor sparse graphs.

<sup>6</sup> <http://streamreasoning.org/larkc/csparql/CSPARQL-ReadyToGoPack-0.9.zip>

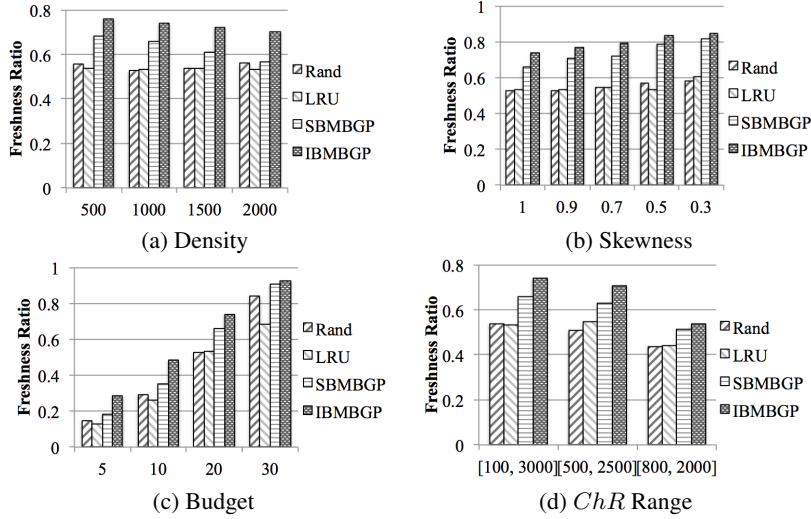


Fig. 2:  $SBM_{BGP}$  and  $IBM_{BGP}$  outperform baselines under different settings

Figure 2(b) plots the performance on graphs with different distributions of  $\mu^{SR}$ 's selectivity. We set the selectivity to follow different Zipf's distributions, with skewness parameter  $s$  to be 1 (uniform), 0.8 (slightly skewed), 0.5 (skewed), and 0.3 (highly skewed). Figure 2(b) shows that the performance improvement of  $SBM_{BGP}$  and  $IBM_{BGP}$  is more significant in skewed graphs. The reason is same with the second observation above:  $SBM_{BGP}$  and  $IBM_{BGP}$  refresh the  $\mu^{SR}$  with the highest  $deg^{SR}$ . In a skewed graph, the percentage of the selected  $\mu^{SR}$  increases, which leads to more fresh results. Therefore,  $SBM_{BGP}$  and  $IBM_{BGP}$  favor skewed  $\mu^{SR}$  selectivity distribution.

Figure 2(c) shows the performance with different budgets i.e.,  $\gamma = 5, 10, 20,$  and  $30$ . With a larger budget, the performance improvement of  $SBM_{BGP}$  and  $IBM_{BGP}$  becomes less. In an extreme case of having a large enough budget to cover most of the *stale*  $\mu^{SR}$ , different subsets of  $\mathcal{E}$  do not affect the freshness anymore. Therefore,  $SBM_{BGP}$  and  $IBM_{BGP}$  can achieve significant improvement with less budget. The above three experiments verify **H1**: considering the selectivity  $deg^{SR}$  of  $\mu^{SR}$  enables choosing better candidates for refreshing and improves response freshness.

Regarding H2, Figure 2(d) shows the performance results of BGD change intervals that are randomly chosen from different ranges:  $[100, 3000]$ ,  $[500, 2000]$ , and  $[800, 1200]$  *ms*. We can make these comparisons: first,  $IBM_{BGP}$  always has a higher freshness than  $SBM_{BGP}$ . Second, having a wider range for  $ChR$  leads to better improvement in  $IBM_{BGP}$ . The reason is that  $IBM_{BGP}$  chooses  $\mu^{SR}$  with larger ‘‘impact’’, i.e., larger


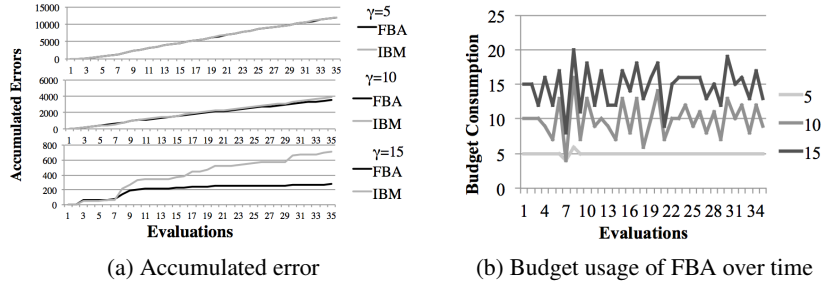


Fig. 3:  $SBM_{Agg}$  and  $IBM_{Agg}$  outperform baselines in subquery Case 2



(a) Accumulated error (b) Budget usage of FBA over time  
 Fig. 4: The performance of FBA under different budgets

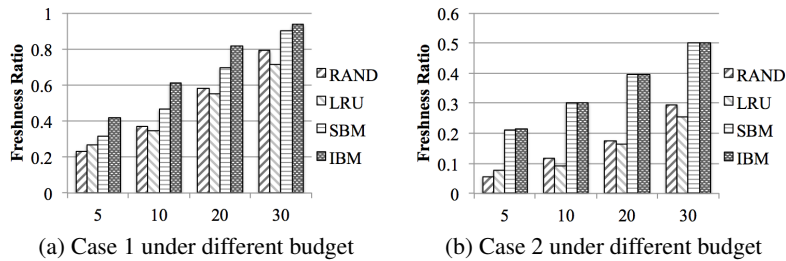
score, since the score indicates that  $\mu^{SR}$  makes more results in the current and future slides. Therefore, this experiment verifies **H2** and shows that  $IBM_{BGP}$  favors larger ranges of change intervals.

**Case 2.** When  $P^S$  is an aggregate query, in all of the above cases, we observed similar performance improvements of  $SBM_{Agg}$  over RAND and LRU. To save space, we just show the results with different skewnesses to demonstrate the performance in Figure 3(a). Besides the freshness improvement, we notice that in most cases  $IBM_{Agg}$  performs similarly with  $SBM_{Agg}$ . This is because  $IBM_{Agg}$  is designed to be at least as good as  $SBM_{Agg}$ . Only when several  $\mu^{SN}$  have the same amount of connected  $\mu^{SR}$ ,  $SBM_{Agg}$  will choose the one with the lowest score. Furthermore, for different  $\mu^{SN}$ , when the overlapping between their associated  $\mu^{SR}$  is small, the chance of  $\mu^{SN}$  s have different scores is larger and the effect of  $SBM_{Agg}$  is, therefore, more significant. Figure 3(b) investigates this by plotting the results of a special case: a very sparse graph (100 edges) and a tiny budget, e.g., 3 to 5. In these cases,  $IBM_{Agg}$  outperforms  $SBM_{Agg}$  by up to 12.5%.

## 6.2 Verifying hypothesis H3

We compare the performance of FBA with IBM in Case 1 with three refresh budgets,  $\gamma = 5, 10, 15$  in Figure 4(a). They track the accumulated number of stale results over time. By increasing the budget, the gap between FBA and IBM becomes more significant. Furthermore, when  $\gamma = 15$ , after the first 15 iterations, FBA makes the accumulated stale result increases very slowly, i.e., the freshness ratio of the answer is almost 100%, while IBM still keep producing stale results. To explain the improvement, Figure 4(b) plots the actual amounts of budget that are consumed over time. For IBM, the consumption of budget will always be a vertical line for different budgets. For FBA, when  $\gamma = 5$ , the line fluctuates a bit. With larger budgets, the lines fluctuate more. It shows that FBA moves budgets between different slides to improve freshness.

**Results on a real data set.** Figure 5 plots the results on a real data set with different budgets for both cases. We can observe that IBM always achieves the best freshness



(a) Case 1 under different budget (b) Case 2 under different budget  
 Fig. 5: SBM and IBM outperform baselines in a real data set

and SBM also outperforms the two baseline algorithms. When we decrease the budget, the performance improvement of IBM and SBM increases. These results confirm our findings in Figure 2(c).

**Computational overhead.** We finally report the computational overhead and the average remote access delay. Under the default setting  $\gamma = 20$ , the total latency of a slide evaluation is about 94.4 *ms*. The delay of querying the BGD server accounts for 92 *ms* on average (4.6 *ms* per request); the computational overhead is only about 2.3 *ms* (2.5% of the overall latency). Note that, the current setting has the remote BGD server running locally. When requests are sent over Internet, the computational overhead will become even more negligible while the performance gain will become more substantial.

## 7 Conclusions and Future Work

In this paper, we studied the problem of accessing remote background data (BGD) from an RDF Stream Processing (RSP) context. When BGD is large, stored remotely, and/or changing over time, accessing it can be expensive, waste resources, and deteriorate the response time. Hence, a local view is often used to speed up the BGD accesses, but maintaining it is often subject to refresh budget constraints. This paper proposes to efficiently allocate the budget for refreshing the local view. Specifically, our solution relies on a bipartite graph to model the join between stream data and BGD. It exploits the graph structure to improve response freshness for two kinds of `SERVICE` subqueries: a BGP query (Case 1) and an aggregate query (Case 2). Our solution, SBM, exploits a set of basic algorithms that leverage the selectivity of the join between the stream and the background data. Experiments show that it can significantly improve the response freshness up to 25% compared to baseline algorithms (i.e., RAND and LRU). An also introduce an improved approach, IBM, that takes the future impact of refreshes into account and improves the performance up to 55.6% over the SBM. Finally, we propose the FBA optimization that flexibly allocates budget considering not only the current but also future data. As a result FBA significantly improves over all other solutions and maintains a freshness of close to 100% even in the light of limited update budget.

Our findings have the following limitations: first, we propose a greedy heuristic algorithm for Case 2. We hope to investigate a more advanced approximate approach in the future. Second, the current approach focuses on BGP. Some SPARQL operators (e.g., `OPTIONAL`) can introduce new challenges and require non-trivial extensions of our model. Third, we currently focus on stream querying. In future, we plan to extend our current optimization problem to reasoning over both stream and BGD. For example, we plan to investigate how to ensure stream consistency over a background knowledge base under a given budget.

Even in the light of these limitations we believe that this paper highlights an important problem in RSP—the joint evaluation of stream and BGD under budget constraints—and provides solutions for different subqueries. As such it paves the way for truly scalable RSP systems in real-world environments, where the integration of stream and BGD is ubiquitous.

**Acknowledgments.** This research has been partially funded by Science Foundation Ireland (SFI) grant No. SFI/12/RC/2289, EU FP7 CityPulse Project grant No.603095 and the IBM Ph.D. Fellowship Award 2014 granted to Dell’Aglío.

## References

1. D. J. Abadi. Consistency tradeoffs in modern distributed database system design: Cap is only part of the story. *Computer*, (2):37–42, 2012.
2. C. B. Aranda, M. Arenas, Ó. Corcho, and A. Polleres. Federating queries in SPARQL 1.1: Syntax, semantics and evaluation. *J. Web Sem.*, 18(1):1–17, 2013.
3. D. F. Barbieri, D. Braga, S. Ceri, E. Della Valle, and M. Grossniklaus. Querying RDF streams with C-SPARQL. *SIGMOD Record*, 39(1):20–26, 2010.
4. J. Calbimonte, H. Jeung, Ó. Corcho, and K. Aberer. Enabling query technologies for the semantic sensor web. *Int. J. Semantic Web Inf. Syst.*, 8(1):43–63, 2012.
5. S. Dehghanzadeh, D. Dell’Aglío, S. Gao, E. Della Valle, A. Mileo, and A. Bernstein. Approximate Continuous Query Answering Over Streams and Dynamic Linked Data Sets. In *ICWE 2015*, pages 307–325, 2015.
6. S. Dehghanzadeh, J. Parreira, M. Karnstedt, J. Umbrich, M. Hauswirth, and S. Decker. Optimizing SPARQL query processing on dynamic and static data based on query time/freshness requirements using materialization. In *JIST*, pages 257–270, 2014.
7. D. Dell’Aglío, E. Della Valle, J. Calbimonte, and Ó. Corcho. RSP-QL semantics: A unifying query model to explain heterogeneity of RDF stream processing systems. *Int. J. Semantic Web Inf. Syst.*, 10(4):17–44, 2014.
8. S. Gançarski, H. Naacke, E. Pacitti, and P. Valduriez. The leganet system: Freshness-aware transaction routing in a database cluster. *Info. Systems*, pages 320–343, 2007.
9. H. Guo, P.-Å. Larson, and R. Ramakrishnan. Caching with good enough currency, consistency, and completeness. In *VLDB*, pages 457–468. VLDB Endowment, 2005.
10. S. Hasan, S. O’Riain, and E. Curry. Towards unified and native enrichment in event processing systems. In *DEBS*, pages 171–182. ACM, 2013.
11. A. Hinze, K. Sachs, and A. Buchmann. Event-based applications and enabling technologies. In *DEBS*, page 1. ACM, 2009.
12. Y. Ji, Z. Jerzak, A. Nica, G. Hackenbroich, and C. Fetzer. Optimization of continuous queries in federated database and stream processing systems. pages 403–422, 2015.
13. T. Käfer, J. Umbrich, A. Hogan, and A. Polleres. Towards a dynamic linked data observatory. *LDOW at WWW*, 2012.
14. A. Labrinidis and N. Roussopoulos. Exploring the tradeoff between performance and data freshness in database-driven web servers. *VLDB J.*, 13(3):240–255, 2004.
15. G. Ladwig and T. Tran. Sihjoin: querying remote and local linked data. In *The Semantic Web: Research and Applications*, pages 139–153. Springer, 2011.
16. D. Le Phuoc, M. Dao-Tran, J. X. Parreira, and M. Hauswirth. A native and adaptive approach for unified processing of linked streams and linked data. In *ISWC’11*, pages 370–388, 2011.
17. R. Lee and Z. Xu. Exploiting stream request locality to improve query throughput of a data integration system. *IEEE Trans. on Computers*, 58(10):1356–1368, 2009.
18. A. Margara, J. Urbani, F. van Harmelen, and H. Bal. Streaming the web: Reasoning over dynamic data. *J. Web Sem.*, 25:24–44.
19. G. Montoya, M.-E. Vidal, O. Corcho, E. Ruckhaus, and C. Buil-Aranda. Benchmarking federated sparql query engines: Are existing testbeds enough? In *ISWC*, pages 313–324, 2012.
20. M. Rinne, M. Solanki, and E. Nuutila. Rfid-based logistics monitoring with semantics-driven event processing. In *DEBS*, pages 238–245, 2016.
21. M. Sharaf, P. Chrysanthi, and A. Labrinidis. Preemptive rate-based operator scheduling in a data stream management system. In *AICCSA*, pages 46–59, 2005.
22. K. Teymourian and A. Paschke. Plan-based semantic enrichment of event streams. In *ESWC*, volume 8465, pages 21–35, 2014.