

Bindings-restricted triple pattern fragments

Olaf Hartig and Carlos Bull-Aranda

Conference Publication



N.B.: When citing this work, cite the original article.

Original Publication:

Olaf Hartig and C Bull-Aranda, Bindings-restricted triple pattern fragments, On the Move to Meaningful Internet Systems, 2016. (), pp.762-769.

http://dx.doi.org/10.1007/978-3-319-48472-3_48

Copyright: www.springer.com

Postprint available at: Linköping University Electronic Press

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-134276>



brTPF: Bindings-Restricted Triple Pattern Fragments (Extended Version)*

Olaf Hartig^{1,2} and Carlos Buil-Aranda³

¹ Hasso Plattner Institute, University of Potsdam, Germany

² Dept. of Computer and Information Science (IDA), Linköping University, Sweden
olaf.hartig@liu.se

³ Informatics Department, Universidad Técnica Federico Santa María, Chile
cbuil@inf.utfsm.cl

Abstract The Triple Pattern Fragment (TPF) interface is a recent proposal for reducing server load in Web-based approaches to execute SPARQL queries over public RDF datasets. The price for less overloaded servers is a higher client-side load and a substantial increase in network load (in terms of both the number of HTTP requests and data transfer). In this paper, we propose a slightly extended interface that allows clients to attach intermediate results to triple pattern requests. The response to such a request is expected to contain triples from the underlying dataset that do not only match the given triple pattern (as in the case of TPF), but that are guaranteed to contribute in a join with the given intermediate result. Our hypothesis is that a distributed query execution using this extended interface can reduce the network load (in comparison to a pure TPF-based query execution) without reducing the overall throughput of the client-server system significantly. Our main contribution in this paper is twofold: we empirically verify the hypothesis and provide an extensive experimental comparison of our proposal and TPF.

1 Introduction

Recent years have witnessed a large and constant growth in the amount of data that is structured based on the data model of the Resource Description Framework (RDF) [5] and made available on the Web through HTTP interfaces [3,12,15]. A prevalent (and standardized) type of such interfaces that provides query-based access to RDF data are SPARQL endpoints [6]; that is, Web services that accept queries written in the SPARQL query language [9]. While a SPARQL endpoint enables users to query its RDF dataset by using the full potential of SPARQL, providing such a comparably complex functionality presents a serious challenge (the evaluation problem of a core fragment of SPARQL has been shown to be PSPACE complete [14]). As a consequence, many public endpoints suffer from frequent downtime; for instance, by monitoring over 400 such endpoint for 27 months, Buil-Aranda et al. show that only 32.3% of the endpoints offer an availability of more than 99%, and 50.4% have an availability of less than 95% [4]. Furthermore, if many client applications start to access such an endpoint concurrently, then the performance of the endpoint (in terms of average query execution times and per-client query throughput) drops significantly [19].

* This document is an extended preprint of a paper published in the proceedings of the ODBASE 2016 conference [11]. In contrast to the proceedings version, this document contains Appendixes A and B which present additional experimental results.

To address these problems, Verborgh et al. recently proposed the Triple Pattern Fragment (TPF) interface [19,20]. This proposal restricts the type of queries supported by the server to single triple patterns. Then, to support arbitrary SPARQL queries over a dataset exposed by such a TPF server, a major part of the query processing effort has to be shifted to the clients. As a result, the server load is reduced and query execution times are more stable. However, everything comes at a cost and, thus, the price for the aforementioned benefits of the TPF approach is not only a higher client-side load but also a significant increase in network load. More precisely, to execute a given SPARQL query, a TPF-based client usually sends many more HTTP requests than a client that sends the whole query with a single request to a SPARQL endpoint. Additionally, the overall amount of data returned in response to these TPF requests is much greater than what a SPARQL endpoint returns (namely, just the query result).

To mitigate this drawback of the TPF approach (without losing the benefits) we propose a slightly extended interface that supports so-called *Bindings-Restricted Triple Pattern Fragments (brTPF)*. That is, in addition to pure TPF requests, the brTPF interface allows clients to attach intermediate results to TPF requests. The response to such a brTPF request is expected to contain RDF triples from the underlying dataset that do not only match the given triple pattern (as in the case of TPF), but that are guaranteed to contribute in a join with the given intermediate result. Hence, given the brTPF interface, it becomes possible to distribute the execution of joins between client and server by using the well-known bind join strategy [8]. Our hypothesis is given as follows:

Hypothesis. *In comparison to pure TPF-based query executions, query executions that are based on the brTPF interface can reduce the network load without reducing the overall throughput of the client-server system significantly.*

Our main contribution in this paper is twofold: First, we empirically verify the aforementioned hypothesis. Second, we also provide an extensive experimental comparison of the brTPF approach and the TPF approach. The queries that we focus on in this study are expressed using SPARQL basic graph patterns (BGPs). We chose this focus because we believe that achieving a comprehensive understanding of how TPF and brTPF behave for this fundamental fragment of SPARQL is essential before attempting to focus on more expressive fragments. All digital artifacts related to our study (e.g., software, test data, etc.) are available on-line at: <http://olafhartig.de/brTPF-ODBASE2016>.

The remainder of the paper is organized as follows: Section 2 provides an overview of the related work on existing HTTP interfaces to access RDF data on the Web. Section 3 presents a formalization of the concepts used in this paper, Section 4 describes the implementation of the prototypes for the brTPF server and client, as well as a brief description of the TPF server and client implementations. Sections 5, 6 and 7 present the experiments that we have done to evaluate TPF and brTPF in terms of the network load, throughput, and cache use, respectively. Finally, we conclude in Section 8.

2 Related Work

One of the most common ways for accessing RDF data on the Web is through HTTP interfaces. Verborgh et al. classified such interfaces along an axis as illustrated in Figure 1 [19]. In the following we describe the currently most common of these interfaces.

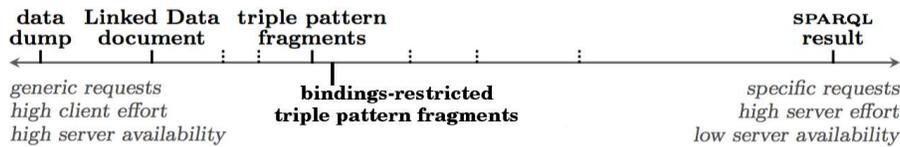


Figure 1. HTTP interfaces to RDF data (adapted from [19]).

2.1 SPARQL Endpoints

SPARQL endpoints are Web services that implement the SPARQL Protocol [6]. A SPARQL endpoint is usually an HTTP interface that accepts SPARQL queries (like Jena Fuseki⁴) and the query processing is done by the endpoint’s triple store (e.g., Virtuoso⁵, Jena TDB⁶). As it was shown in [4], public SPARQL endpoints may suffer from low availability and poor performance.

2.2 Triple Pattern Fragments

Verborgh et al. [19] proposed the *Triple Pattern Fragments* (TPF) interface to access RDF data which relies on client processing power to execute SPARQL queries while the server only provides the data for the triple patterns in the queries. By relying on the clients to execute the actual operations on the data, it is possible to provide higher availability with the same server infrastructure, since servers only perform operations that require minimal effort. The originally proposed client-side algorithm for the TPF interface is based on a decomposition of a given query into triple patterns; the resulting subqueries are then executed recursively in dynamically generated pipelines that order the subqueries based on result size metadata as provided by the TPF server [19,20]. In comparison to SPARQL endpoints, the downside of this approach is not only a higher client-side load but also an increase in network load (in terms of both the number of HTTP requests and data transfer).

To reduce the impact of this downside some new approaches to execute SPARQL queries over the TPF interface appeared. For instance, Van Herwegen et al. extend the original TPF client algorithm by injecting every partial solution into the next subquery with which it shares variables [17]. Acosta et al. [1] introduce a new TPF client algorithm using a query optimizer that focuses on reducing both the amount of intermediate results and the requests posed to the TPF server, and it also implements an adaptive routing query engine that is able to dynamically adapt the query plan according to execution conditions. These works focus only on optimizing the client algorithm while the server remains unchanged, and even though the approaches improve TPF-based query execution considerably, we argue that there is room for improvement at the server side.

There have been proposals to extend the TPF server instead of focusing on the clients. Van Herwegen et al. extended the original TPF controller to deal with substring

⁴ https://jena.apache.org/documentation/serving_data/

⁵ <http://virtuoso.openlinksw.com/>

⁶ <https://jena.apache.org/documentation/tdb/>

filtering [16]. The authors extended the server capabilities using an Elastic Search service, allowing the server to execute FILTER queries besides the usual triple patterns. Another approach, by Vander Sande et al., is to extend the TPF metadata by using approximate membership functions (Bloom filters and Golomb-coded sets) [18]. As a result of using these statistical techniques, the amount of HTTP requests can be reduced by 25%, barely increasing the server load. On the other hand, due to an increased response size, query execution times could not be improved by these techniques.

In addition to a single-server scenario, it has also been shown that distributed query processing is possible over a federation of multiple TPF servers [20]. In this context, Montoya et al. introduce a TPF-based federation system that is aware of replicated data [13]. Given a set of TPF servers with data replicated across these servers, the system creates an efficient execution plan for a given query by identifying which fragments will contribute to the final query result. Again, this work focuses mainly on the client, and tries to minimize the amount of data transferred from the servers to the client.

2.3 Linked Data Querying

Hartig et al. [10] proposed a Linked Data traversal approach, in which the follow your nose concept is applied for following links between linked RDF datasets on the Web. The implementation consists first in a source selection process for next traversing the links obtained from the initial query to the selected data sources. The link traversal operator implements an asynchronous pipeline of iterators executing first the most selective iterator. The query engine is also able to adapt the execution to source availability by detecting whenever an HTTP server stops responding. This is one of RDF data access approaches that imposes less load on the server, however it is among the slowest.

3 Formal Definition of brTPF

In this section, we provide a formal definition of brTPF. We assume that the reader is familiar with the fundamental concepts of RDF [5] and SPARQL [9,14].

Then, for our formalization we adopt the general formal framework for defining interfaces to RDF datasets as provided by Verborgh et al. [20]. As the basis for formalizing any type of such an interface, the authors introduce the notion of a *selector function* that captures conditions for selecting sets of triples from datasets (cf. [20, Definition 1]). Formally, any selector function is a mapping between RDF graphs.

We define such a selector function for brTPFs as follows.

Definition 1. *Given a triple pattern tp and a finite sequence of solution mappings Ω , the **bindings-restricted triple pattern selector** for tp and Ω , denoted by $s_{(tp,\Omega)}$, is the selector function that, for every RDF graph G , is defined by*

$$s_{(tp,\Omega)}(G) = \begin{cases} \{t \in G \mid t \text{ is a matching triple for } tp\} & \text{if } \Omega \text{ is empty,} \\ \{t \in G \mid t \text{ is a matching triple for } tp, \text{ and there} & \text{else.} \\ \quad \text{exists a solution mapping } \mu \text{ such that} & \\ \quad \text{the application of } \mu \text{ to } tp \text{ results in } t & \\ \quad \text{and } \mu \text{ is compatible with some } \mu' \text{ in } \Omega\} & \end{cases}$$

Given the general notion of a selector function, in [20, Definition 3], Verborgh et al. define the abstract concept of a *Linked Data Fragment* (LDF) of an RDF graph G as a tuple $f = \langle u, s, \Gamma, M, C \rangle$ where u is a URI (from which f can be retrieved); s is a selector function; Γ is a set of (blank-node-free) RDF triples that is the result of applying the selector function s to G , i.e., $\Gamma = s(G)$; M is a finite set of (additional) RDF triples, including triples that represent metadata for f ; C is a finite set of hypermedia controls (cf. [20, Definition 2]). Additionally, by [20, Definition 5], for a hypermedia control c (cf. [20, Definition 2]), a *c-specific LDF collection* over an RDF graph G is a set F of LDFs such that, for each LDF $f \in F$ with $f = \langle u, s, \Gamma, M, C \rangle$, the following three properties hold: i) f is an LDF of G , ii) $s \in \text{dom}(c)$, and iii) $c(s) = u$.

By using these preliminaries, we can now define brTPFs as follows.

Definition 2. *Given a positive integer maxMpR and a control c (as per [20, Definition 2]), a c -specific LDF collection F (cf. [20, Definition 5]) is called a **bindings restricted triple pattern fragment collection** for maxMpR if, for any possible triple pattern tp and any finite sequence Ω of at most maxMpR distinct solution mappings, there exists an LDF $\langle u, s, \Gamma, M, C \rangle \in F$ (as per [20, Definition 3]), referred to as a **bindings restricted triple pattern fragment (brTPF)**, that has the following three properties:*

1. *Selector s is the bindings-restricted triple pattern selector for tp and Ω ;*
2. *There exists a (metadata) RDF triple $\langle u, \text{void:triples}, cnt \rangle \in M$ with cnt representing an estimate of the cardinality of Γ ; that is, cnt is an integer that has the following two properties:*
 - (a) *If $\Gamma = \emptyset$, then $cnt = 0$.*
 - (b) *If $\Gamma \neq \emptyset$, then $cnt > 0$ and $\text{abs}(|\Gamma| - cnt) \leq \epsilon$ for some F -specific threshold ϵ .*
3. *$c \in C$.*

Observe that our definition assumes a positive integer maxMpR that presents a well-defined restriction on the number of distinct solution mappings that can be attached to any brTPF request supported by any specific brTPF interface. Furthermore, note that we require Ω to be a *sequence* of solution mappings (rather than a set). This requirement reduces the complexity of a correct (i.e., deterministic) implementation of paging for the brTPF interface.

4 Evaluation Prototypes

This section describes the implementations that we used for our experiments.

4.1 Server Implementation

As a basis for the server, we used an established Java servlet implementation of the TPF interface⁷ and extended it with the functionality to also support the brTPF interface. As a result, both interface implementations coexist within the Java servlet, which chooses which of them it invokes depending on the HTTP GET request it receives: If the request

⁷ <https://github.com/LinkedDataFragments/Server.Java>

contains a bindings-restricted triple pattern selector, then the brTPF implementation is used to generate the response; if the HTTP request just contains a TPF selector, the TPF implementation is used. Having both implementations in a single software component has the advantage that commonly used basic functionality (such as serializing RDF triples for a response) is shared and, thus, experimental results are not affected by potential differences in how efficient the implementation of such basic functionality is.

TPF Server Implementation: The TPF server implementation is a Java servlet that accepts HTTP GET and POST requests. The servlet identifies the SPARQL triple pattern enclosed within such a request and evaluates the triple pattern using an internal storage component that contains the dataset exposed via the TPF interface. The particular storage component used by the servlet in our experiments is based on an RDF-HDT back-end [7], which manages a highly compressed, in-memory representation of RDF data. The TPF server returns to the client an HTTP response containing the RDF triples from the evaluation of the triple pattern divided in pages (whose size is configurable) and an estimation of the entire result set size. These result size estimations are obtained by just asking the RDF-HDT back-end how many matching triples there are for the given triple pattern. The TPF client uses such result set size estimation in its query execution algorithm.

brTPF Server Implementation: The brTPF server implementation extends the TPF implementation servlet by enabling it to process brTPF requests as follows. Given such a request, the servlet internally generates a stream of data triples for the requested (brTPF) fragment and processes this stream in the same way as the TPF implementation processes the stream of matching triples returned by the RDF-HDT backend. To generate the data triples for a brTPF request, the servlet iterates over the sequence of solution mappings that comes with the request. For each such mapping, the servlet applies this mapping to the triple pattern of the request by replacing variables in the triple pattern according to the mapping. From the resulting sequence of potentially more concrete triple patterns, the servlet removes all duplicates. Next, the servlet uses each remaining triple pattern, one after the other, to query the RDF-HDT backend. The resulting streams of matching triples are then concatenated into the desired stream of all data triples for the brTPF request.

4.2 TPF Client Implementation

For our experiments we use a TPF client that is implemented using Node.js and that uses the TPF-based query execution algorithm for SPARQL basic graph patterns (BGP) as proposed by Verborgh et al. [19]. This algorithm is based on iterators that are arranged in pipelines. Query results are computed recursively by executing the pipelines. Each of these pipelines is generated for a subquery obtained from a decomposition of the initial BGP. Each iterator executes one of these subqueries returning as well an estimation of the size of its response. The algorithm uses this estimation to adapt its execution dynamically so the subqueries with a smaller result are executed first. In this way it is possible to rapidly return a first subset of the query result.

4.3 brTPF Client Implementation

To develop a brTPF client we used the aforementioned Node.js-based TPF client and added a brTPF-based query execution algorithm to it. Hence, as for the servers, all basic functionality required for both client implementations is based on the same source code. This way, we ensure the comparability of our experimental results and allow for a fair comparison of both approaches. For the same purpose, the brTPF-based query execution algorithm that we developed is kept deliberately simple. In fact, we did not attempt to integrate any sophisticated query optimization technique such as the adaptive, intermediate result based approach to generate subplans at query execution runtime as used by the TPF client. Instead, our brTPF algorithm simply chooses a fixed query execution plan upfront. This plan represents a left-deep join tree that is implemented using a fixed pipeline of iterators such that each of these iterators is responsible for a different triple pattern of the query. The join order is decided based on intermediate result cardinality estimates for every triple pattern of the query. These estimates can be obtained from the server by requesting the first TPF page for each of the triple patterns.

During query execution, every iterator receives chunks of solution mappings from its predecessor. The size of these chunks corresponds to the value of `maxMpR` as specified by the brTPF interface. Given such a chunk, the iterator issues a brTPF request consisting of the triple pattern that the iterator is responsible for and the solution mappings from the chunk. Upon arrival of the data for the requested brTPF, the iterator uses this data to generate chunks of solution mappings for the next iterator in the pipeline.

5 Experimental Comparison of Network Load

Our first group of experiments focuses on comparing TPF and brTPF in terms of the network load that the interfaces may cause when accessed by clients that execute SPARQL queries. In this section, we first introduce the metrics and the experimental setup that we use for these experiments, and, thereafter, we present the results of the experiments.

5.1 Metrics

For this comparison we focus on two metrics: First, the *number of requests* (`#req`) that a client sends to the server during the execution of a query. Recall that both the TPF interface and the brTPF interface split fragments into pages. Therefore, the measurements for `#req` do not correspond to the number of fragments (TPF or brTPF) requested during query executions but to the number of pages requested for the fragments that the client chooses to access. Notice furthermore that such requests do not necessarily have to reach the server if there exists an HTTP cache in the network between the client and the server. Nonetheless, unless the cache is located directly in front of the client (in which case it may not be very effective), the requests are sent into the network.

The second network-related metric that we focus on is the *amount of data received* (`dataRecv`) by the client during query executions. We measure this metric in terms of the number of RDF triples contained in all fragment pages that the client receives during a query execution. Observe that this metric is independent of whether the data comes directly from the server or from an HTTP cache that acts as a proxy server.

5.2 Setup

For this group of experiments we use a single-machine setup with a single client. That is, the combined TPF/brTPF Java servlet is deployed (using Jetty 9.2.5) on the same machine on which the client implementation performs the query executions (using either the TPF algorithm or the brTPF algorithm). This machine runs the Ubuntu 12.04.5 LTS operating system with Oracle Java 1.8.0_92 and Node.js 0.10.37, and it is equipped with an Intel Core i7-2620M CPU (2.7GHz) and 8 GB of main memory. In all of our experiments, every query execution is performed by using a separate operating system process (that is, we stop and restart the client software between any two executions).

As a basis for the experiments we used the 10M triples dataset provided on the project page⁸ of the Waterloo SPARQL Diversity Test Suite (WatDiv) [2], and we used a sequence of 145 BGP queries that we selected uniformly at random from the WatDiv stress test query workload⁹. This workload has been shown to be a challenging benchmark that is very diverse in terms of various structural and data-driven features [2].

5.3 Results

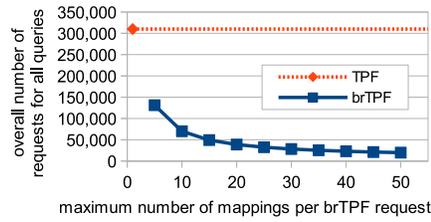
For our first experiment we use a page size of 100 data triples per fragment page (which is the default configuration of the TPF server implementation that we use as a basis for our evaluation) and execute the query sequence using the TPF client and the brTPF client, respectively. For the latter we repeat the execution of the query sequence using each of the following values for `maxMpR`: 5, 10, 15, . . . , 45, and 50.¹⁰ The charts in Figure 2(a) and 2(b) provide an aggregated view on the resulting measurements for WatDiv. In particular, Figure 2(a) illustrates the overall `#req` summed up for each client over the whole sequence of queries, respectively; similarly, Figure 2(b) illustrates the sums of the `dataRecv` measurements obtained for all queries, respectively.

Regarding these measurements, we make the following observations. By first focusing on brTPF only, we observe that the overall `#req` decreases with an increasing value for `maxMpR` (from about 131K for `maxMpR=5` to about 20K for `maxMpR=50`), and so does the overall `dataRecv` (from about 1,126K for `maxMpR=5` to about 756K for `maxMpR=50`). While for `#req` this observation is not surprising (if the fraction of any large intermediate result that can be sent with each request is smaller, the brTPF client has to send a greater number of such requests), for `dataRecv` we explain the observation as follows: Each fragment page contains not only data triples but also additional meta-data triples that refer to the next and the previous page, describe the controls of the LDF interface, etc. Therefore, when the number of fragment pages requested and received is greater (as is the case for smaller `maxMpR`; cf. Figure 2(a)), then so is the overall number of these additional triples that have to be received with each fragment page.

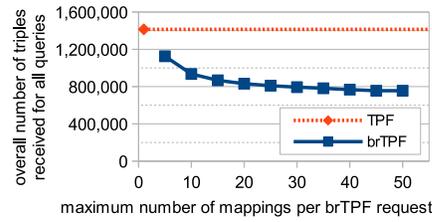
⁸ <http://dsg.uwaterloo.ca/watdiv/>

⁹ <https://cs.uwaterloo.ca/~galuc/watdiv/stress-workloads.tar.gz>

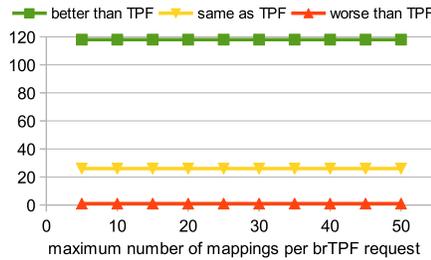
¹⁰ We do not increase `maxMpR` beyond a value of 50 because the client prototype uses the HTTP method GET, which causes problems for our server if `maxMpR>50`; that is, in preliminary tests with greater values we observed server responses with status code 414 (URI Too Long).



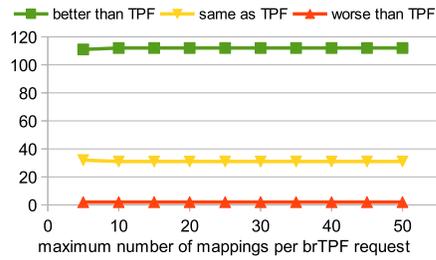
(a) sum of all #req for the WatDiv runs



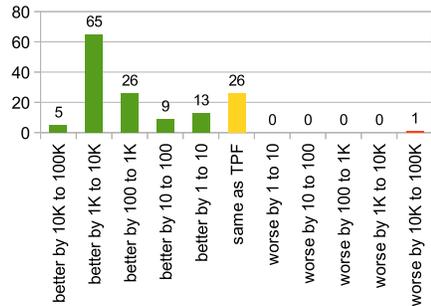
(b) sum of all dataRecv



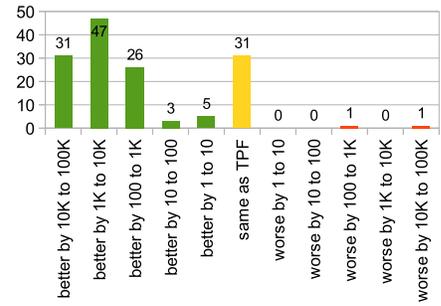
(c) number of queries for which brTPF has a better/same/worse #req than TPF



(d) number of queries where brTPF has a better/same/worse dataRecv than TPF



(e) breakdown of the number of queries in terms of the differences between #req of brTPF (maxMpR=30) and of TPF



(f) breakdown of the number of queries in terms of the differences between dataRecv of brTPF (maxM/R=30) and of TPF

Figure 2. Measurements of network-related metrics using WatDiv.

By now comparing the behavior of TPF vs. brTPF in the charts in Figures 2(a)–2(b), we notice that for both the overall #req and the overall dataRecv, brTPF achieves significantly smaller values. More specifically, regarding dataRecv, brTPF achieves between 53.5% (maxMpR=50) and 79.6% (maxMpR=5) of the overall dataRecv of TPF (which is about 1,414K), and for #req, it even goes down to 6.5% (maxMpR=50) of the overall #req of TPF (310K). At this point, we have to recall that these charts only show aggregated measurements. Hence, it might still be possible that the vastly superior behavior of brTPF as shown in these charts is actually only due to a small number of outliers. We can verify that this is not the case by drilling into the measurements: For the differ-

ent values of maxMpR , Figure 2(c) illustrates the number of queries for which brTPF has a smaller (i.e., better) or greater (i.e., worse) $\#\text{req}$ than TPF. Figure 2(d) presents a corresponding comparison for dataRecv . In Figures 2(e)–2(f), we drill in even deeper for $\text{maxMpR}=30$ (corresponding charts for the other values of maxMpR look very similar) and report the number of queries for which the difference between the $\#\text{req}$ (resp. dataRecv) of brTPF vs. TPF is between 100K to 10K, between 10K to 1K, etc. These charts show that, in terms of both $\#\text{req}$ and dataRecv , brTPF is not only better than TPF in an impressively high number of cases, but for a large majority of these cases in which brTPF is better, the differences are significant.

To investigate whether these results are different for a different page size we conducted another experiment in which we varied the page size (number of data triples per fragment page). That is, with both the TPF client and the brTPF client (using $\text{maxMpR}=15$ and $\text{maxMpR}=30$ as exemplars), we repeated the execution of the query sequences for each of the following page sizes: 100, 250, 500, 1000, and 2000. Due to space limitations, we do not include charts for this experiment in this paper. However, we highlight that the measurements obtained by this experiment show for both brTPF and TPF, the page size does not have any considerable impact on $\#\text{req}$ or on dataRecv . In other words, the relative differences between brTPF and TPF as identified by the first experiment are independent of the page size (and so are the relative differences between the different maxMpR configurations for brTPF). Hence, our main conclusion from these experiments is that, *independent of the page size (and the value of maxMpR), brTPF typically achieves a significantly smaller $\#\text{req}$ and dataRecv than TPF.*

6 Experimental Comparison of Performance under Load

With our second group of experiments we aim to compare the behavior of TPF and brTPF when multiple clients access a server concurrently,¹¹ and we aim to analyze how the two approaches scale to an increasing number of concurrent clients. To this end, we deliberately ignore the possibility of having an HTTP cache that may reduce the server load (we consider caching in our third group of experiments as presented in Section 7).

6.1 Setup

For our multi-client experiments we use a cluster of 17 identical machines that are connected via a 10 Gb Ethernet network (Dell PowerConnect 6248 switch). Each machine has an Intel Core i7 processor with 4 cores at 2.6Ghz and 8 GB of main memory. The operating system running on all machines is an Ubuntu 14.04 LTS.

One of these machines we used as the server and deployed the combined TPF/brTPF Java servlet on an Apache Tomcat 8 application server running with Java 1.8. The page size in this setup we fix to 100 data triples per fragment page.

The other 16 machines we used to simulate clients. For this purpose, we configured each of these machines to run on each of its 4 cores a thread with a single brTPF/TPF

¹¹ We also tried to conduct a comparison between TPF/brTPF and a Virtuoso-based SPARQL endpoint. Unfortunately, as detailed in Appendix A, the Virtuoso server failed to correctly serve multiple concurrent clients in our experimental setup.

client, which gives us a total of up to 64 brTPF/TPF clients that we used for executing different WatDiv query sequences in parallel. To this end, we used the aforementioned WatDiv stress test query workload which consists of a set of 12,400 different queries; we split this set into 64 distinct sets, and distributed these sets over our experiment cluster such that each of the 64 CPU cores had available a total of 193 queries that they always executed as a sequence in the same order. We also configured each client to terminate any query execution that did not complete within 5 minutes. These query executions were recorded in our experiments as 'timed out.' Once a client stopped a query execution after 5 minutes, the client starts executing the next query from its sequence.

6.2 Metrics

The metric that we consider to study the performance of the approaches under load is *query throughput* (throughput). We measure throughput in terms of the overall number of queries that all concurrent clients manage to execute within one hour.¹²

6.3 Results

The three charts on the left hand side of Figure 3 illustrate the measurements obtained by the experiment (ignore the right-hand-side figures for the moment).

Before we focus on the throughput values measured, we need to mention that we observed a varying, but often very high number of timed out queries across all runs. More specifically, as illustrated in Figure 3(c), for both TPF and brTPF, the number of timeouts increases with an increasing number of clients; and in all cases, the TPF clients run into a substantially higher number of timeouts than the brTPF clients. We identified the following cause of this behavior: Both client implementations asynchronously issue multiple HTTP requests in parallel. The TPF server has comparably less work to do to answer each of these multiple HTTP requests, returning the data as fast as possible. However both clients at this stage still have to do some data processing to join intermediate results. In the case of TPF, this client-side processing is significantly more work due to the larger amounts of data that the TPF client receives from the TPF server. As a result of this extra work that the TPF client has to do, there are more query executions that exceed the 5 mins timeout threshold before they complete.

Given the high amount of timeouts, it is important to report not only the throughput of all query execution attempts, including the ones the timed out, but also the throughput in terms of query executions that terminated normally after computing their complete query result. Therefore, Figure 3(a) illustrates the former and Figure 3(e) the latter. In these charts we make two main observations: First, the brTPF approach achieves a greater throughput than TPF, and, second, both approaches are able to achieve higher throughput values for an increasing number of clients. However, regarding the latter we also observe that the brTPF approach scales better if we look at the throughput of completed query executions, whereas, by comparing only the throughput in terms of all query execution attempts, both approaches seem to exhibit a similar scaling behavior.

¹² Appendix B complements our presentation of throughput-related results by also drilling in to the query execution times measured during our multi-client experiments.

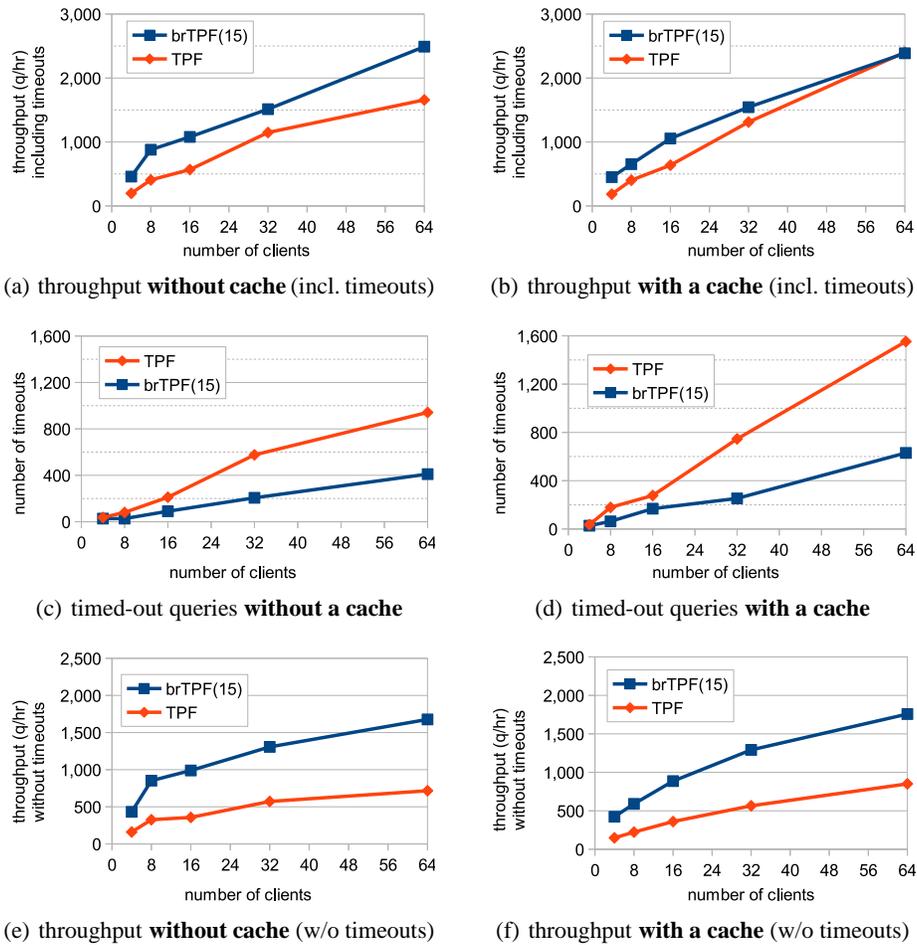


Figure 3. Measurements of overall system performance in the multi-client setup using WatDiv.

As a last observation regarding this group of experiments we note that the average query execution time (average QET) across all WatDiv queries that were executed completely (no timeouts) in the 4-clients setup was 17.9 secs for TPF (st.dev. 33.2) and 16.5 secs for brTPF (st.dev. 44.9). For the 64-clients setup this average QET increased to 45.3 secs for TPF (st.dev. 61.0) and to a much smaller 33.1 secs for brTPF (st.dev. 57.7). Appendix B drills into these numbers and discusses QETs in more detail.

7 Experimental Comparison of Performance with Cache

For our analysis in the previous section we ignore the possibility of HTTP caches. We recall that such caches are designed to reduce the load of Web servers by serving requests that are identical to earlier requests (instead of requiring the server to recompute

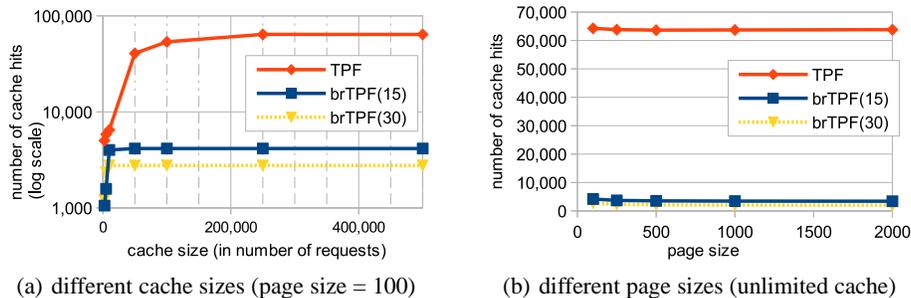


Figure 4. Comparison of the cacheability of TPF requests vs. brTPF requests.

the response for such identical requests over and over again). Therefore, our previous setup without such a cache can be conceived of as some kind of a worst-case scenario for both TPF and brTPF, and it might be worse for one than it is for the other depending on the respective likelihood for observing identical requests from different query executions. In fact, it seems reasonable to assume that this likelihood is higher for TPF than it is for brTPF. More precisely, it seems more likely that different TPF-based query executions request the same triple pattern than it is for different brTPF-based executions to request an identical pair of the same triple pattern and the same sequence of solution mappings. To verify this assumption and to identify how caching affects the performance of both approaches we conducted another set of experiments. In this section, we describe these experiments and present their results.

7.1 Cache Hit Potential

To systematically study and to compare the potential for TPF requests and for brTPF requests to be served from a cache we use the same single-machine setup as used for the network-related experiments (cf. Section 5.2). As a metric for this analysis, we use the *number of cache hits* (#hits) as could be achieved by a possible HTTP cache when executing the test sequence of WatDiv queries. We instrumented our combined TPF/brTPF server implementation to measure this number assuming either an unlimited cache or a cache whose size is limited to a given number of distinct requests (using LRU as replacement policy). For our analysis we first used the latter and varied the cache size from 2.5K, 5K, 10K, 50K, 100K, 250K, to 500K.

The chart in Figure 4(a) illustrates our measurements for these different cache sizes when executing the WatDiv query sequence with either the TPF client or the brTPF client (using $\text{maxMpR}=15$ and $\text{maxMpR}=30$ as exemplars, and a page size of 100). First and foremost, we observe that TPF always achieves a significantly higher #hits than brTPF (note that the y-axis is log scale). This observation verifies the aforementioned assumption that the likelihood for observing identical requests from different query executions is higher for TPF than it is for brTPF. By focusing on brTPF, we notice that a smaller value for maxMpR results in a higher #hits. More precisely, the #hits for $\text{maxMpR}=15$ in this experiment is always about 150% of the #hits for $\text{maxMpR}=30$,

which is not surprising given that brTPF requests with a greater number of solution mappings attached to them are more specialized than brTPF requests with a smaller number of solution mappings. As a final noteworthy observation regarding the measurements in Figure 4(a) we mention that the curves flatten out completely at some point and this point is different for each curve. The respective cache size at each of these points correlates with the overall number of requests issued during the respective WatDiv run (see the #req as reported for these runs in Figure 2(a)). Unsurprisingly, for cache sizes that are large enough to cover all distinct requests issued during such a run, the #hits is equivalent to the #req as achieved by an unlimited cache.

In a second experiment we assumed an unlimited cache and executed the query sequences with different page sizes to investigate whether the cache hit potential of both TPF and brTPF is affected by the page size. Our measurements, as reported in Figure 4(b), show that the page size has no impact in the #hits. Therefore, from these experiments we conclude that, *independent of the page size and the value of maxMpR, TPF has a significantly higher potential to benefit from HTTP caches than brTPF.*

7.2 Impact of Caching on Performance

After verifying that TPF is more likely to benefit from an HTTP cache than brTPF, the obvious question that arises is whether the use of a cache allows TPF to gain a measurable advantage in performance. To answer this question we extended our multi-machine setup (cf. Section 6.1) with an additional machine on which we run an Nginx proxy server (1.4.6) as an HTTP cache that is located between the client machines and the machine with the TPF/brTPF server. This additional machine has an Intel Core i7 CPU with 8 cores and 16 GB of main memory, and it also runs an Ubuntu 14.04 LTS operating system.

Given this extended setup, we repeated the same set of multi-client executions as in the corresponding experiment without the cache (cf. Section 6). The three charts on the right hand side of Figure 3 illustrate the resulting measurements for the number of timeouts and the throughput (with and without queries whose execution timed out).

As a first observation, by comparing Figures 3(c) and 3(d), we note that the number of timeouts for TPF has increased substantially in comparison to the experiment without the cache. Our explanation of such an increased number of timed out queries is similar to what we saw in Section 6.3, however more exacerbated. That is, for many of the additional query execution attempts, the amount of data that the TPF clients receives results in client-side processing work that exceeds the 5 minutes timeout threshold. Figures 3(a) and 3(b) show the query throughput by brTPF and TPF without a cache server and with a cache server, respectively. By looking at Figure 3(b) it is possible to see how the TPF query throughput is almost the same than brTPF's query throughput when 64 clients access the servers concurrently. This high TPF throughput is due to the large amount of timed out queries as Figure 3(d) shows and as Figures 3(e) and 3(f) confirm. By comparing Figures 3(e) and 3(f), we observe that, in the case of a high number of concurrent clients (i.e., the 64-clients runs in our experiments), the performance of *both approaches* benefits from the cache, whereas for smaller number of concurrent clients we do not observe any significant impact for either approach. To explain the latter we revisit the measurements of our single-machine experiments: Given the #hits achieved

by the WatDiv runs discussed in the previous section, cf. Figure 4(b), and the corresponding values for the overall number of requests during these runs, cf. Figure 2(a), we can compute hit rates. For TPF we obtain a hit rate of about 20.7% and for brTPF it is about 10.7% for $\text{maxMpR}=15$ and 11.8% $\text{maxMpR}=30$.¹³ It appears that these hit rates are too small to allow for the cache to take significant load from the server in the case of a smaller number of concurrent clients. Only when the overall system becomes more busy with a higher number of concurrent clients, the availability of the cache reduces the load on the server and the throughput increases. Then, due to the greater hit rates, this increase in throughput is greater for TPF than it is for brTPF.

However, the perhaps most surprising finding is that TPF cannot benefit enough from the cache to gain an advantage over brTPF. We explain this finding by the network load that is significantly higher for TPF (as shown in Section 5). In particular, the increased amounts of data that need to be transferred and processed prove to be the primary weakness of TPF, even if some of the data comes from the cache.

In conclusion, this experiment shows that *for both approaches, TPF and brTPF, caching helps to increase the overall performance in a multi-client setting (in particular for greater numbers of clients), but it does not help TPF to outperform brTPF.*

8 Conclusions

In this paper we present an interface to access RDF datasets that slightly extends the TPF interface. Our extended interface, which we call Bindings-Restricted Triple Pattern Fragments (brTPF), allows clients to attach intermediate results to triple pattern requests. By an extensive evaluation of brTPF and TPF we obtain the following results:

- Our main conclusion from the experimental comparison of network load is that, independent of the page size (and the value of maxMpR), brTPF typically achieves a significantly smaller number of requests than TPF, and less data is transferred.
- From the experimental comparison of server performance under load our conclusions are twofold: first, the brTPF approach achieves a greater throughput than TPF, and, second, both approaches are able to achieve higher throughput values for a greater number of clients. Regarding the latter we also observe that brTPF scales better if we look at the throughput in terms of completed query executions, whereas, by comparing the throughput in terms of all query execution attempts (which includes timed-out executions), both approaches exhibit a similar scaling behavior.
- From the experimental comparison of server performance with an HTTP cache we conclude that for both approaches, TPF and brTPF, caching helps to increase the overall performance in a multi-client setting (in particular for greater numbers of clients), but it does not help TPF to outperform brTPF.

In this paper we improved the idea of TPFs by significantly reducing the amount of HTTP requests and query execution times. Since we have shown these reductions by using a deliberately straightforward client implementation, we believe that there is even more room for improvement by using more sophisticated client-side query execution algorithms. Consequently, our future work includes investigating such algorithms.

¹³ Although the $\#hits$ is higher for $\text{maxMpR}=15$ than for $\text{maxMpR}=30$ (cf. Section 7.1), the former achieves a smaller hit rate because it has a higher $\#req$ (as we observe in Section 5.3).

Acknowledgements: Olaf Hartig’s work has been funded partially by the German Government, Federal Ministry of Education and Research (project no. 03WKJ4D). Carlos Buil-Aranda’s work has been supported by the Millennium Nucleus Center for Semantic Web Research (Grant NC120004) and by UTFSM DGIP Project no. 116.24.1

References

1. Acosta, M., Vidal, M.: Networks of Linked Data Eddies: An Adaptive Web Query Processing Engine for RDF Data. In: Proc. of the 14th Int. Semantic Web Conference (ISWC) (2015)
2. Aluç, G., Hartig, O., Özsu, M.T., Daudjee, K.: Diversified Stress Testing of RDF Data Management Systems. In: Proc. of the 13th Int. Semantic Web Conference (ISWC) (2014)
3. Bizer, C., Eckert, K., Meusel, R., Mühleisen, H., Schuhmacher, M., Völker, J.: Deployment of RDFa, Microdata, and Microformats on the Web – A Quantitative Analysis. In: Proc. of the 12th Int. Semantic Web Conference (ISWC) (2013)
4. Buil-Aranda, C., Hogan, A., Umbrich, J., Vandenbussche, P.Y.: SPARQL Web-Querying Infrastructure: Ready for Action? In: Proc. of the 12th Int. Semantic Web Conf. (ISWC) (2013)
5. Cyganiak, R., Wood, D., Lanthaler, M.: RDF 1.1 Concepts and Abstract Syntax. W3C Recommendation (Feb 2014)
6. Feigenbaum, L., Williams, G.T.: SPARQL protocol for RDF. W3C Recommendation (2013)
7. Fernández, J.D., Martínez-Prieto, M.A., Gutiérrez, C., Polleres, A., Arias, M.: Binary rdf representation for publication and exchange (hdt). Web Semantics: Science, Services and Agents on the World Wide Web 19, 22–41 (2013)
8. Haas, L.M., Kossmann, D., Wimmers, E.L., Yang, J.: Optimizing Queries Across Diverse Data Sources. In: Proc. of the 23rd Int. Conf. on Very Large Data Bases (VLDB) (1997)
9. Harris, S., Seaborne, A.: SPARQL 1.1 Query Language. W3C Recommendation (2013)
10. Hartig, O., Bizer, C., Freytag, J.C.: Executing SPARQL Queries over the Web of Linked Data. In: Proc. of the 8th Int. Semantic Web Conference (ISWC) (2009)
11. Hartig, O., Buil-Aranda, C.: Bindings-Restricted Triple Pattern Fragments. In: Proc. of the 15th Int. Conf. on Ontologies, Databases, and Applications of Semantics (ODBASE) (2016)
12. Mika, P., Potter, T.: Metadata Statistics for a Large Web Corpus. In: Proc. of the 5th Linked Data on the Web Workshop (LDOW) (2012)
13. Montoya, G., Skaf-Molli, H., Molli, P., Vidal, M.: Federated SPARQL Queries Processing with Replicated Fragments. In: Proc. of the 14th Int. Semantic Web Conf. (ISWC) (2015)
14. Pérez, J., Arenas, M., Gutierrez, C.: Semantics and Complexity of SPARQL. ACM Transactions on Database Systems 34(3) (2009)
15. Schmachtenberg, M., Bizer, C., Paulheim, H.: Adoption of the Linked Data Best Practices in Different Topical Domains. In: Proc. of the 13th Int. Semantic Web Conf. (ISWC) (2014)
16. Van Herwegen, J., De Vocht, L., Verborgh, R., Mannens, E., Van de Walle, R.: Substring Filtering for Low-Cost Linked Data Interfaces. In: Proc. of the 14th Int. Semantic Web Conference (ISWC) (2015)
17. Van Herwegen, J., Verborgh, R., Mannens, E., Van de Walle, R.: Query Execution Optimization for Clients of Triple Pattern Fragments. In: Proc. of the 12th European Semantic Web Conference (ESWC) (2015)
18. Vander Sande, M., Verborgh, R., Van Herwegen, J., Mannens, E., Van de Walle, R.: Opportunistic Linked Data Querying Through Approximate Membership Metadata. In: Proc. of the 14th Int. Semantic Web Conference (ISWC) (2015)
19. Verborgh, R., Hartig, O., De Meester, B., Haesendonck, G., De Vocht, L., Vander Sande, M., Cyganiak, R., Colpaert, P., Mannens, E., Van de Walle, R.: Querying Datasets on the Web with High Availability. In: Proc. of the 13th Int. Semantic Web Conference (ISWC) (2014)
20. Verborgh, R., Vander Sande, M., Hartig, O., Van Herwegen, J., De Vocht, L., De Meester, B., Haesendonck, G., Colpaert, P.: Triple Pattern Fragments: a Low-cost Knowledge Graph Interface for the Web. Journal of Web Semantics 37–38, 184–206 (2016)

Appendix

A Comparison to a Virtuoso-based SPARQL Endpoint

We also tried to conduct a comparison between brTPF and a Virtuoso-based SPARQL endpoint in our experimental setup. To this end, we installed the latest version of the Open Source edition of Virtuoso (v.7.2.4) on the machine that we used as the server in our multi-machine experiments (cf. Section 6.1), and we loaded the WatDiv dataset (cf. Section 5.2) into this Virtuoso endpoint. For the clients we developed a simple Python script that submits the sequence of WatDiv queries of the respective client using the SPARQL endpoint REST API, one query after another (recall from Section 6.1 that each of the up to 64 clients uses a distinct set of 193 queries that are always executed as a sequence in the same order). Since the result cardinality of some of these queries is between 100K and 1M, to obtain complete query results and, thus, to achieve a fair comparison, we configured the result size limit of the Virtuoso server to be 1M.

For the Virtuoso-based run with 4 clients, everything went smooth and Virtuoso achieved a throughput of 217 q/hr. With 8 clients, we saw 33 query executions hitting our 5 minutes timeout (over the course of one hour); the throughput without these timeouts was 1,420 q/hr, which is roughly twice the throughput achieved by brTPF and roughly 4x of what the TPF approach achieved in this experiment (cf. Figure 3(e)). However, when we increased the load to 16 concurrent clients, Virtuoso always managed to answer about 300 queries before it crashed, which always happened around 2 minutes after starting this run. Before the crash, Virtuoso wrote out log messages such as the following: "System is under high load. Adding cluster nodes or using more replicated copies may be needed." We provide the Virtuoso configuration file and the Virtuoso log files from this test on the Web page for this paper (as referred to in Section 1).

To investigate this issue we repeated the 16-clients run with a configuration in which we reduced the aforementioned result size limit to 100K. Hence, we allowed Virtuoso to return partial results. With this configuration, Virtuoso did not have any trouble serving the 16 concurrent clients, but we observed a high number of queries for which Virtuoso returned a result of the exact size of 100K. Hence, these all are incomplete query results! Therefore, from this test we conclude that the reason for why Virtuoso crashed when trying to correctly serve our 16 clients (with complete query results) is due to the difficulty of high-cardinality WatDiv queries.

Given these observations, we decided not to conduct any further comparison with the Virtuoso-based SPARQL endpoint in the context of this paper; we did not want to compare with a system that returns partial results, or crashes during the experiments.

B Query Execution Times in the Multi-Client Experiments

The average query execution time (average QET) across all WatDiv queries that were executed completely (no timeouts) in the 4-clients setup without cache was 17.9 s for TPF (st.dev. 33.2) and 16.5 s for brTPF (st.dev. 44.9). For the 64-clients setup this average QET increased to 45.3 s for TPF (st.dev. 61.0) and to a much smaller 33.1 s for brTPF (st.dev. 57.7). For the setup with a cache, the relative differences are similar.

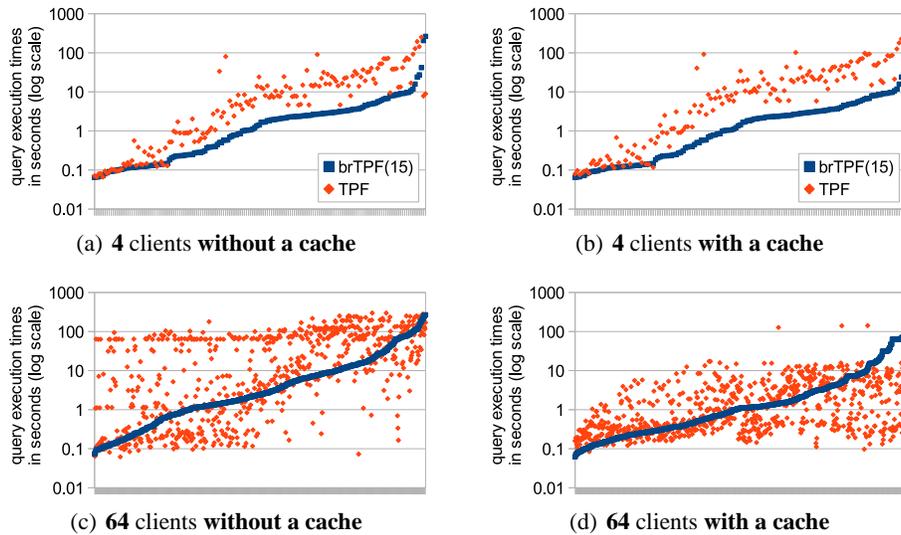


Figure 5. Execution times of all the WatDiv queries that were executed completely (without timeout) by *both* the respective TPF setup and the respective brTPF setup ($\text{maxMpr}=15$). Hence, these charts ignore all queries that were executed completely by only one of the approaches (due to a timeout or a smaller throughput of the other approach). In each chart the measurements are ordered from left to right by the time that brTPF required to execute the corresponding query.

To drill into these numbers, the charts in Figure 5 detail the individual execution times for all queries that were executed completely by both TPF and brTPF in the respective setup. More precisely, in each chart, the queries are organized along the x-axis by sorting them from left to right based on their respective QET in the brTPF case. Then, at the x-axis position of such a query, the chart contains two measurement points that indicate the brTPF-based QET and the TPF-based QET of the query, respectively.

For the 4-clients setups without cache (Figure 5(a)) and with cache (Figure 5(b)) we observe that for almost all queries, brTPF achieved a QET that is either similar or even smaller than the QET achieved by using TPF. For the 64-clients setups (Figures 5(c) and 5(d)), there is more variation, which we attribute to fact that, for both approaches, TPF and brTPF, individual QETs are more affected by the higher load of the whole system. By interpreting these charts, the reader should keep in mind that these charts only consider queries executed completely in both the corresponding TPF and brTPF runs. What these charts do not show is that the brTPF approach achieved more than twice as many complete query executions as TPF as we recall from Figures 3(e) and 3(f).