

Migrating to GraphQL: A Practical Assessment

Gleison Brito, Thais Mombach, Marco Tulio Valente

ASERG Group, Department of Computer Science, Federal University of Minas Gerais, Brazil

{gleison.brito, thaismombach, mtov}@dcc.ufmg.br

Abstract—GraphQL is a novel query language proposed by Facebook to implement Web-based APIs. In this paper, we present a practical study on migrating API clients to this new technology. First, we conduct a grey literature review to gain an in-depth understanding on the benefits and key characteristics normally associated to GraphQL by practitioners. After that, we assess such benefits in practice, by migrating seven systems to use GraphQL, instead of standard REST-based APIs. As our key result, we show that GraphQL can reduce the size of the JSON documents returned by REST APIs in 94% (in number of fields) and in 99% (in number of bytes), both median results.

Index Terms—GraphQL, REST, APIs, Migration Study

I. INTRODUCTION

GraphQL is a novel query language for implementing Web-based APIs [1]. Proposed by Facebook in 2016, the language represents an alternative to popular REST-based APIs [2]–[4], shifting from servers to clients the decision on the precise data returned by API calls. To illustrate the usage of the language, suppose the REST API currently implemented by arXiv, the popular preprint service maintained by Cornell University. This API includes a *search* endpoint that allows clients to retrieve metadata about preprints with a given title. The result is a complex and large JSON document, with at least 33 fields. However, clients might need only a few ones (e.g., only the paper’s URL). Despite that, the mentioned endpoint returns all fields in a JSON document, which should be parsed by the clients. After that, the unneeded fields are discarded, although they have consumed server, client, and network resources. By contrast, suppose arXiv decides to support GraphQL. Using the language, clients formulate a simple call like this one:

```
1 search(title: "A Solution of the P versus NP Problem"){
2   pdfUrl
3 }
```

By means of this query, the client request a single field (*pdfUrl*) of a preprint entitled “*A Solution of the P versus NP Problem*”. The result is a JSON file with just this URL. Therefore, instead of receiving a full document, with 33 fields, the client receives exactly the single field it needs to process.

GraphQL is gaining momentum and it is now supported by important web services, as the ones provided by GitHub and Pinterest [5]. Despite that, we have few studies investigating the real benefits of using GraphQL for implementing Web-based APIs. Therefore, in this paper we ask the following research questions: (*RQ1*) What are the key characteristics and benefits of GraphQL? (*RQ2*) What are the main disadvantages of GraphQL? (*RQ3*) When using GraphQL, what is the reduction in the number of API calls performed by clients? (*RQ4*) When using GraphQL, what is the reduction

in the number of fields of the documents returned by servers? (*RQ5*) When using GraphQL, what is the reduction in the size of the documents returned by servers? To answer *RQ1* and *RQ2*, we conduct a grey literature review, covering 28 popular Web articles (mostly blog posts) about GraphQL. Since the query language has just two years, we focus on grey literature, instead of analysing scientific papers, as usually recommended for emerging technologies [6]–[8]. As a result, we confirmed two key characteristics of GraphQL: (a) support to an *hierarchical data model*, which can contribute to reduce the number of endpoints accessed by clients; (b) support to *client-specific queries*, i.e., queries where clients only ask for the precise data they need to perform a given task. Motivated by these findings, we also assess the benefits achieved by GraphQL in terms of a reduction in the number of API calls (*RQ3*) and in the number of fields returned by service providers (*RQ4*). To answer these questions, we manually migrated five clients of the GitHub REST API to use the new GraphQL API provided by GitHub. We also implemented a GraphQL wrapper for two endpoints of arXiv’s REST API and migrated two open source clients to use this wrapper. Finally, to answer *RQ5*, we reimplemented in GraphQL 14 queries used in seven recent empirical software engineering papers, published at two major software engineering conferences (ICSE and MSR).

Our contributions are twofold: (1) we reveal that GraphQL does not lead to a reduction in the number of queries performed by API clients in order to perform a given task, when compared to the number of required REST endpoints. For example, in our migration study, we migrated 29 API calls that access REST endpoints (distributed over seven clients) to 24 GraphQL queries, which therefore does not represent a major reduction; (2) we reveal that client-specific queries can lead to a drastic reduction in the size of JSON responses returned by API providers. On the median, in our study, JSON responses have 93.5 fields, against only 5.5 fields after migration to GraphQL, which represents a reduction of 94%. In terms of bytes, we also measure an impressive reduction: from 9.8 MB (REST) to 86 KB (GraphQL). Altogether, our findings suggest that API providers should seriously consider the adoption of GraphQL. We also see space for tool builders and researchers, with interest on providing support and improving the state-of-the-practice on GraphQL-based API development.

The rest of this paper has seven sections. Section II provides a brief introduction to GraphQL. Section III describes a grey literature review, covering popular Web articles about GraphQL. Section IV presents the migration study conducted to answer *RQ3* and *RQ4*. Section V describes a study to eval-

uate the runtime gains achieved by GraphQL (and therefore answer *RQ5*). Threats to validity are discussed in Section VI; and related work is discussed in Section VII. Section VIII concludes the paper and summarizes lessons learned.

II. GRAPHQL IN A NUTSHELL

This section introduces the key concepts of GraphQL. The goal is to make this paper self-contained; for a detailed and throughout presentation of the language we refer the reader to its documentation [1]. Essentially, GraphQL allows clients to query a database represented by a *schema*. Therefore, this design represents a major shift from REST APIs: in REST, server applications implement a list of endpoints (operations) that can be called by clients; by contrast, in GraphQL servers export a database, which can be queried by clients. A GraphQL database is defined by a *schema*, which is a multi-graph [9]. In this multi-graph, nodes are *objects*, which define a type and include a list of *fields*; each field has a name and also a type. Edges appear when objects define fields whose types are other object types.¹ GraphQL proposes a simple DSL for defining schemas. To illustrate the usage of this language, we use a simple blogging system, with two object types: *Post* and *Author*. As presented in Listing 1, object types are defined using the keyword `type`. In this example, *Post* has four fields: *id*, *author*, *title*, and *body*. The first field is a non-null *String* (the `!` after the type discards null values). The second field (*author*) references another object type in the schema, called *Author* (lines 8-12). The remaining two fields in *Post* have a *String* type. Finally, schemas usually include a predefined type, called *Query*, which represents the entry point of GraphQL APIs (lines 14-16). A *Query* object exposes the object types that can be queried by clients and the arguments that must be provided by them. For example, `post` (line 15) is a query that accepts a non-null string as argument and returns the *Post* object having this string as id.

```

1 type Post {
2   id: String!
3   author: Author
4   title: String
5   body: String
6 }
7
8 type Author {
9   id: String!
10  name: String
11  email: String
12 }
13
14 type Query {
15   post(id: String!): Post
16 }

```

Listing 1. Schema with two types (*Post* and *Author*) and a *Query* end-point

GraphQL also defines a query language, used by clients. Listing 2 shows three examples of queries in this language. In the first query (`PostByTitle`), the client asks for the *Post* object with *id* equals to 1000; specifically, the client only requests the *title* field of this object. The second query (`PostByTitleAndBody`) is similar, but in this case the client

¹Since an object T_1 can have multiple fields of type T_2 , multiple edges can connect T_1 to T_2 , leading to a multi-graph.

asks for two fields, *title* and *body*. Finally, the last query (`PostByTitleAndAuthor`) asks for the *title* and *author* of the same *Post*. Since *author* is another object, we have to specify its queried fields, in this case only *name*. The result of this third query is presented in Listing 3. As we can see, the result is a JSON object, which should be parsed and possibly deserialized by clients.

```

1 query PostByTitle{
2   post(id:"1000"){
3     title
4   }
5 }
6 query PostByTitleAndBody{
7   post(id:"1000"){
8     title
9     body
10  }
11 }
12 query PostByTitleAndAuthor{
13   post(id:"1000"){
14     title
15     author {
16       name
17     }
18   }
19 }

```

Listing 2. Querying distinct data from a given *Post* object

```

1 { "data": {
2   "post": {
3     "title": "GraphQL: A data query language"
4     "author": {
5       "name": "Lee Byron"
6     }
7   }
8 } }
9

```

Listing 3. JSON object returned by `PostByTitleAndAuthor` query

To respond to queries, the developer of a GraphQL server must implement a *resolver* function for each query declared in the *Query* type. These functions are called each time the GraphQL server engine needs to retrieve an object type specified in a query. Typically, these functions retrieve these objects from an underlying data structure, which can be any kind of database (relational, non-relational, in-memory, etc). Finally, it is also possible to define another predefined type in schemas, called *Mutation*, which is used to insert new objects on the server's database or modify existing ones. Listing 4 shows an example that defines an `addPost` mutation, which receives a *Post* object (and returns the object, just to confirm the operation has been successfully executed). Each endpoint (operation) in a *Mutation* type must have a resolver function, which implements the operation. In our running example, this function must insert the *Post* object received as argument in the underlying database.

```

1 type Mutation {
2   addPost(post: Post): Post
3 }

```

Listing 4. Mutation operation for *Post* objects

III. GREY LITERATURE REVIEW

This section reports the results of a systematic analysis of the grey literature about GraphQL, covering documents and discussions on blogs, tutorials, and similar Web articles.

Our goal is to better understand the key characteristics, benefits, and shortcomings appointed by practitioners who had a real experience with the language. Since it is a new technology, papers about GraphQL are not common in the scientific literature. Therefore, for such emerging technologies, a grey literature tends to provide a better coverage of relevant documents than a traditional literature review [7], [10].²

A. Study Design

To retrieve an initial list of Web articles considered in this review, we used Hacker News,³ which is a news aggregator site widely used by practitioners [11]. Recently, other similar reviews have used Hacker News as data source, e.g., a grey literature review on clouding computing services [12]. To retrieve Hacker News documents, we used the Algolia search engine,⁴ querying for posts containing *graphql* in their titles, as in September, 2018. We found 1,242 articles. We then sequentially removed articles that do not include a valid URL (286 articles), that do not contain comments (760 articles), or that are just promoting a tool or project (168 articles). After this filtering step, we selected 28 articles for analysis (1,242 - 286 - 760 - 168), which we refer as *A1* to *A28*.⁵ Figure 1 shows the year when these articles appeared on Hacker News. We can see an increase in the interest on GraphQL in the last four years. Interestingly, we found four articles in 2015, i.e; before GraphQL official release. Figure 2 shows violion plots with the distribution of the number of comments and upvotes of the 28 articles. The median number of comments is 9.5; and the median number of upvotes is 53 (which is usually enough to put the article in the front page of Hacker News).




Fig. 1. Grey literature articles by year of appearance on Hacker News




Fig. 2. Number of comments and upvotes on Hacker News (for the articles included in the grey literature review)

²Nevertheless, peer-reviewed articles are also discussed in this paper, but in Section VII (Related Work).

³<https://news.ycombinator.com>

⁴<https://hn.algolia.com>

⁵Detailed information at <https://github.com/gleisonbt/migrating-to-graphql>

After collecting the articles, the first author of this paper carefully read them and followed an open coding protocol to provide answers to the first two research questions:

RQ1: What are the *characteristics* and *benefits* of GraphQL?

RQ2: What are the main *disadvantages* of GraphQL?

B. Results

RQ1: Key Characteristics and Benefits

GraphQL is strongly typed, since all objects and fields have types (as mentioned in A1, A2, A5, A6, A10, and A28). This contributes to better tooling support, as reported in this article:

GraphQL is strongly-typed. Given a query, tooling can ensure that the query is syntactically correct and valid within the GraphQL type system before execution. (A5)

A related benefit is the possibility of having better error messages, e.g., *[types] allow GraphQL to provide descriptive error messages before executing a query (A28)*.

GraphQL enables client-specified queries, as mentioned by almost half of the articles (A1, A2, A5, A6, A10, A14, A17, A19, A21, A25, A26, and A27). The following article nicely describes this characteristic:

In GraphQL, the specification for queries are encoded in the client rather than the server. These queries are specified at field-level granularity. In the vast majority of applications written without GraphQL, the server determines the data returned in its various scripted endpoints. A GraphQL query, on the other hand, returns exactly what a client asks for and no more. (A5)

This characteristic makes GraphQL particularly interesting for mobile applications, which often face limited bandwidth and speed (A4, A5, A14, and A26). It also *moves the focus of development to client apps, where designers and developers spend their time and attention (A1)*. Finally, client-specific queries allow servers to better understand the needs of clients (A9, A12) and therefore improve the quality of their service:

It's great for service operators too, because its explicitness allows them to get a better understanding of exactly what their users are trying to do. (A9)

GraphQL data model is hierarchical, as mentioned in five articles (A1, A2, A3, A5, and A8) and defined as:

As its name would suggest, GraphQL models objects as a graph. Technically, the graph starts with a root node that branches into query and mutation nodes, which then descend into API-specific resources. (A2)

This characteristic allows clients to retrieve data from multiple sources (or endpoints) in a single request, therefore acting as gateways for different APIs (A3, A4, A14, and A20):

GraphQL makes it easy to combine multiple APIs into one, so you can implement different parts of your schema as independent services. (A20)

Introspection, which allows clients to inspect the types and fields defined in a schema, at runtime (A1, A3, A9, A16, and A28). Combined with a static type system, introspection allows clients to *learn and explore an API quickly without grepping the codebase or wrangling with cURL*. (A28). It also frees servers to support an interface description language, which are usually not available for REST; and when available they are *often not completely accurate because the description is not tied directly to the implementation* (A28).

Deprecation: As common in mainstream programming languages, it is possible to deprecate fields, using a `@deprecated` annotation (A1, A2 and A19). However, in GraphQL, new fields added to a type do not lead to breaking changes (as in standard APIs [13], [14]); and deprecations can be restricted to deleted fields. As a result, the pressure for versioning is less frequent, as mentioned in this article:

This process removes the need for incrementing version numbers. We still support three years of released Facebook applications on the same version of our GraphQL API. (A1)

RQ2: Disadvantages

GraphQL does not support information hiding. GraphQL does not support private fields, i.e., all fields are visible to client applications (A8, A11, A18, A20 and A24). Furthermore, according to A18, GraphQL queries tend to be more complex to implement, since they require a detailed understanding of the data schema, which can be a time-consuming task in large APIs:

By design, a developer who integrates against GraphQL needs to know the names of the fields to access, their relation to other objects and when to retrieve them. (A18)

Complex caching: In GraphQL, each query can be different, even though operating on the same type. This demands more sophisticated server-side caching, as mentioned in this article:

GraphQL does not follow the HTTP specification for caching and instead uses a single endpoint. Thus, it's up to the developer to ensure caching is implemented correctly ... (A20)

Performance: GraphQL servers can have to process complex queries (e.g., queries with deep nesting) that can consume server resources (A8, A11, A20, A23, and A25), as mentioned in the following article:

Great care has to be taken to ensure GraphQL queries don't result in expensive join queries that can bring down server performance or even DDoS the server. (A20)

Figure 3 summarizes the grey literature review results, by presenting the key characteristics, benefits, and disadvantages of GraphQL, and the number of articles mentioning them.




Fig. 3. Summary of grey literature findings

IV. MIGRATION STUDY

With this second study, we aim to quantitatively evaluate two key characteristics associated to GraphQL in the grey literature: (1) clients can precisely request the data they need from servers (due to the support to client-specific queries) (2) clients rely on a single endpoint to retrieve the data they need (due to a hierarchical data model). In the study, we migrate seven client applications based on REST APIs to GraphQL. Then, we assess the gains achieved by the GraphQL version. Specifically, we answer two research questions:

RQ3: When using GraphQL, what is the reduction in the number of API calls performed by clients? GraphQL clients normally implement a single query to retrieve all data they need to perform a given task; by contrast, when using REST, clients frequently have to access multiple endpoints. Therefore, in this RQ, we compare the number of endpoints accessed by REST clients with the number of endpoints accessed by the same clients after refactored to use GraphQL.

RQ4: When using GraphQL, what is the reduction in the number of fields of the JSON documents returned by servers? In GraphQL, client-specific queries allow developers to inform precisely the fields they need from servers. Therefore, in this RQ, we compare the number of fields in the following JSON documents: (a) returned by servers when responding to requests performed by REST clients; (b) returned by servers when responding to queries performed by the same clients but after being migrated to use GraphQL.

A. Study Design

Selected APIs: First, to answer the proposed research questions, we selected the APIs provided by two widely popular services: GitHub and arXiv.⁶ GitHub is an interesting case study because the system provides both REST and GraphQL APIs (the latter since 2016). Moreover, GitHub’s GraphQL API is quite complete and large, including 120 object types, 21 queries, and 62 mutations. By contrast, arXiv only provides a REST API. Therefore, we implemented ourselves a small GraphQL API for the system, in the form of a wrapper for the original API. This wrapper supports two queries (`getPreprint` and `search`), as presented in Listing 5. The first query (`getPreprint`) returns metadata about a preprint, given its ID. This metadata includes the paper’s title, authors, DOI, summary, URL, etc (see Listing 6). The second query searches for preprints whose title match a given string; it is also possible to define the maximal number of results the query should return, the first result that should be returned and the sort order. The implemented GraphQL wrapper was installed in a private server, in our research lab.

```

1 type Query {
2   getPreprint(id: ID!): Preprint
3   search(query: String!, maxResults: Int!,
4     start: Int!, sortBy: String,
5     sortOrder: String): [Preprint]
6 }

```

Listing 5. arXiv Query

```

1 type Preprint {
2   id: ID
3   pdfUrl: String
4   published: String
5   arxivComment: String
6   title: String
7   authors: [String]
8   arxivUrl: String
9   doi: String
10  tags: [Tag]
11  arxivPrimaryCategory: ArxivPrimaryCategory
12  updated: String
13  summary: String
14 }

```

Listing 6. Preprint type

Selected Clients: When searching for GitHub API clients, we first found that they usually have the tag (or topic) *GitHub*. Therefore, we selected five projects with this tag and that have at least 100 stars, as described in Table I. In the case of arXiv, we selected two clients mentioned in the project’s page⁷ and that have their source code publicly available on GitHub (see also their names in Table I). Table II shows information about the programming language, number of stars, size (in lines of code), and contributors of the selected systems. The smallest project is BIBCURE/ARXIVCHECK (131 LOC, one contributor, and five stars); the largest projects are VDAUBRY/GITHUB-AWARDS (35,153 LOC, 15 contributors, and 1,296 stars) and DONNEMARTIN/GITSOME (17,273 LOC, 24 contributors and 5,913 stars).

⁶<https://arxiv.org>

⁷<https://arxiv.org/help/api/index>

TABLE I
SELECTED PROJECTS

Project	Description
DONNEMARTIN/VIZ	Visualization of GitHub repositories
DONNEMARTIN/GITSOME	Command line interface for GitHub
CSURFER/GITSUGGEST	A tool to suggest GitHub repositories
GUYZMO/GIT-REPO	Command line interface to manage Git services
VDAUBRY/GITHUB-AWARDS	Ranking of GitHub repositories
BIBCURE/ARXIVCHECK	A tool to generate B _I T _E X of arXiv preprints
KARPATY/ ARXIV-SANITY-PRESERVER	Web interface for searching arXiv submissions

TABLE II
STATS OF SELECTED PROJECTS

Project	Lang.	Stars	LOC	Contrib
DONNEMARTIN/VIZ	Python	627	9,556	1
DONNEMARTIN/GITSOME	Python	5,913	17,273	24
CSURFER/GITSUGGEST	Python	613	389	2
GUYZMO/GIT-REPO	Python	764	5,602	17
VDAUBRY/GITHUB-AWARDS	Ruby	1,296	35,153	15
BIBCURE/ARXIVCHECK	Python	5	131	1
KARPATY/ ARXIV-SANITY-PRESERVER	Python	2,322	2,431	19

Migration Step: After selecting the APIs and client projects, the paper’s first author exhaustively searched the code of each client looking for REST calls. He then migrated each one to use GraphQL. Just to show one example of migration, in CSURFER/GITSUGGEST the following REST endpoint is used to search GitHub for repositories matching a given string:

```
1 GET /search/repositories
```

This endpoint requires three parameters: `q` (a string with the search keywords), `sort` (the sort field, e.g., `stars`), and the `order` (`asc` or `desc`).⁸ The request returns a JSON document with 94 fields, containing data about a repository. However, only three fields are used by CSURFER/GITSUGGEST: `owner’s login`, `description`, and `stargazers_count`. Therefore, we changed the function that implements the search call to use the following GraphQL query, which retrieves exactly the three fields used by CSURFER/GITSUGGEST:

```

1 query searchRepos{
2   search(query:$query, type:REPOSITORY, first: 100){
3     nodes{
4       ... on Repository{
5         nameWithOwner
6         description
7         stargazers{
8           totalCount
9         }
10      }
11    }
12  }
13 }

```

Listing 7. Example of GraphQL query (CSURFER/GITSUGGEST)

In Listing 7, the `search` query returns an union type, which might be either a `Repository`, `User`, or an `Issue` type, depending on the `type` argument. We use a feature of GraphQL called *inline fragments* to access only the fields

⁸<https://developer.github.com/v3/search/>

TABLE III
REST CALLS

Project	Func	REST endpoints
CSURFER/GITSUGGEST	F1	GET /users/:user
		GET /users/:user/starred
		GET /users/:user/following
		GET /users/:user/starred
	F2	GET /search/repositories
DONNEMARTIN/GITSOME	F3	GET /users/:user/followers
	F4	GET /users/:user/following
	F5	GET /repos/:owner/:repo/issues
	F6	GET /users/:user/repos
		GET /repos/:owner/:repo/pulls
	F7	GET /users/:user/repos
	F8	GET /search/issues
	F9	GET /search/repositories
	F10	GET /users/:user/starred
	GUYZMO/GIT-REPO	F11
GET /users/:user/repos		
F12		GET /users/:user/repos
F13		GET /users/:user/gists
F14		GET /repos/:owner/:repo
		GET /repos/:owner/:repo/pulls
F15		GET /repos/:owner/:repo
F16	GET /repos/:owner/:repo	
DONNEMARTIN/VIZ	F17	GET /users/:user
	F18	GET /search/repositories
VDAUBRY/GITHUB-AWARDS	F19	GET /repos/:owner/:repo
	F20	GET /users/:user
		GET /users/:user/repos
BIBCURE/ARXIVCHECK	F21	GET /query/:search_query
KARPATHY/ ARXIV-SANITY-PRESERVER	F22	GET /query/:search_query

of the Repository variant type. This variant is labeled as `...onRepository` (line 4). Therefore, in this case one REST endpoint is replaced by one GraphQL query (RQ3’s answer) and 91 fields (= 94 – 3) are retrieved but not used by the REST code (RQ4’s answer).

In total, the first author migrated 29 REST endpoint calls—distributed over the seven projects (see Table III)—to use GraphQL queries. For the sake of legibility, we use labels F1 to F22 to refer to the functions including these REST calls (instead of the functions’ original names). This migration effort consumed around 60 working hours (of the paper’s first author), including the time to understand the clients code.

Number of JSON fields: To answer RQ4, we have to compute the number of fields returned by the original API calls (performed using REST) and by the migrated calls (using GraphQL). First, it is important to highlight that we only count root nodes, i.e., the ones that have a primitive value associated to them, instead of referring to another JSON entry. Second, when the returned fields are lists, we only consider a single list element. For example, Listing 8 shows a JSON object

that contains a list of users followed by a given GitHub user. The list contains three nodes elements, delimited by square brackets (lines 5-7). Each node contains only one root field called `name`. Therefore, we consider that the JSON document in Listing 8 has only one field (which appears three times). Essentially, we followed this strategy to allow computing the number of fields in each document without having to define a synthetic load for executing the systems, which is not a simple task. Instead, we executed the systems with a trivial load and input, which is sufficient for counting the number of unique root nodes, without considering their number. We leave a more detailed evaluation of the runtime gains achieved with GraphQL to Section V.

```

1 { "data": {
2   "user": {
3     "following": {
4       "nodes": [
5         { "name": "user_1", },
6         { "name": "user_2", },
7         { "name": "user_3", }
8       ]
9     }
10  }
11 }
12 }
```

Listing 8. JSON document with a single root field (`name`)

B. Results

RQ3: What is the reduction in the number of API calls?

The 29 REST calls migrated in the study are implemented in 22 functions (see Table III). For each function (identified by F1 to F22), Figure 4 shows the number of REST calls performed in the original code and the number of GraphQL queries implemented in the migrated code. As we can see, in 17 functions (77%), there is a single REST call, which was therefore migrated to a single GraphQL query. In another function (F20), the two existing REST calls were migrated to two GraphQL calls. In only four functions (F1, F6, F11 and F14), there is a reduction in the number of REST calls. The highest reduction was observed in F1, where four REST calls were replaced by a single GraphQL query. In this case, the REST calls retrieve the repositories starred by a user and by the users he/she follows; they were replaced by the following semantically equivalent GraphQL query:

```

1 query interestingRepos($username: String!){
2   user(login: $username){
3     starredRepositories{
4       nodes { ... }
5     }
6     following(first: 100){
7       nodes{
8         starredRepositories{
9           nodes{ ...}
10        }
11      }
12    }
13  }
14 }
```

Listing 9. Query that returns the repositories starred by an user (lines 3-5) and by the users he/she follows (lines 6-12)

Interestingly, in a single function (F17), there is an increase in the number of API calls after migrating to GraphQL. This

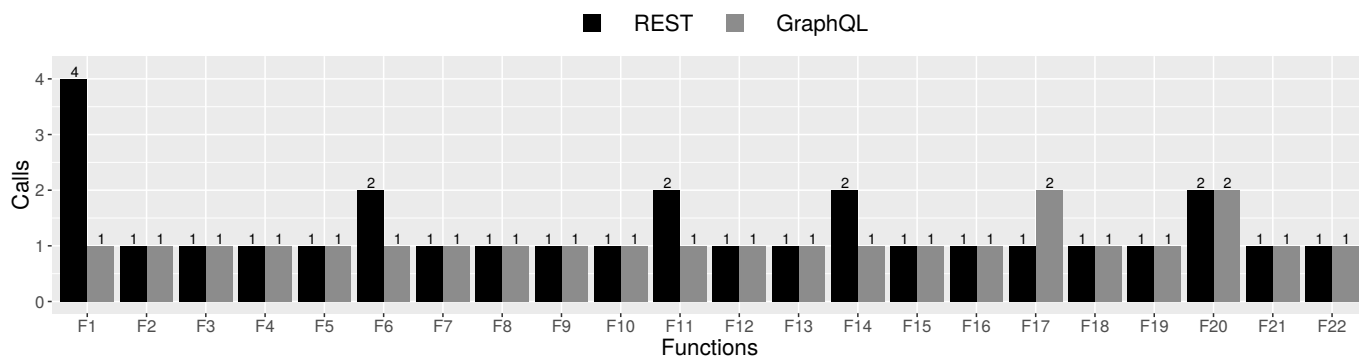


Fig. 4. RQ3 results: number of API calls (REST vs GraphQL) per function

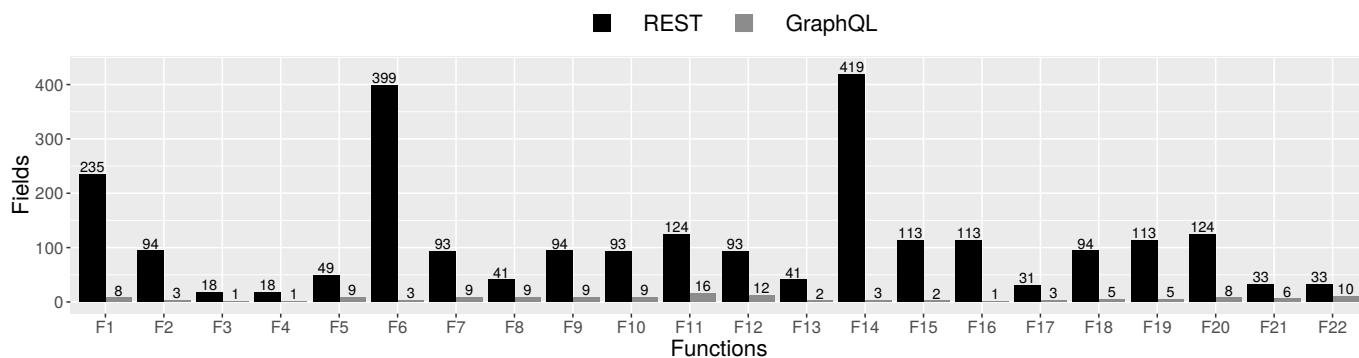


Fig. 5. RQ4: Number of fields returned by API calls (REST vs GraphQL) per function

function retrieves data about GitHub users; however, the required data depends on whether the user has an individual or an organizational account. In the REST API, there is a single endpoint that returns the whole set of fields about GitHub users, despite their account type. By contrast, in the GraphQL API, data about users is spread over three types: `User`, `Organization`, and `Actor`. The migrated code first queries `Actor` to retrieve the user’s category. Depending on the result, a second query targets `User` or `Organization`.

RQ3’s summary: The support to an hierarchical data model is a key characteristic of GraphQL, since it allows clients to retrieve data from multiple endpoints in a single request. However, in our migration study, we found very few opportunities to implement such queries. The reason is that most client functions access a single REST endpoint; the straightforward migration strategy is therefore to replace such calls by a single GraphQL query. Typically, client functions perform simple tasks, which reduces the demand for queries returning complex and nested data structures.

RQ4: What is the reduction in the number of JSON fields?

Figure 5 shows the number of unique root fields in the JSON documents returned by the original REST calls and by the same calls migrated to GraphQL. As we can see, in almost

all calls there is a major decrease in the number of returned fields when using GraphQL (and therefore client-specific queries). This reduction ranges from 17 fields (F3 and F4) to 416 fields (F14). Particularly, F14 is a function that returns data about the pull requests of a given repository. In the original code, the function relies on two REST endpoints to perform this task. The first endpoint returns all fields about the repository of interest. However, F14 consumes only the `pulls_url` field. Then, for each pull request returned by the second endpoint, F14 uses only three fields (`number`, `title`, and `html_url`); these are precisely the fields returned by the GraphQL query. Figure 6 shows violin plots with the distribution of the number JSON fields returned by REST and GraphQL. The REST calls return 93.5 fields (median values), against only 5.5 after migration to GraphQL. The 1st quartile measures are 41 (REST) and 3 (GraphQL); the 3rd quartiles are 113 (REST) and 9 (GraphQL).

RQ4’s summary: When using REST, clients need to process large JSON documents to consume just a few fields, which is often called *over-fetch* (A8 and A5). By contrast, when using GraphQL, clients specify exactly the fields they need from servers. In our study, there is a reduction from 93.5 to 5.5 in the number of JSON fields returned by REST endpoints when compared to equivalent GraphQL queries.

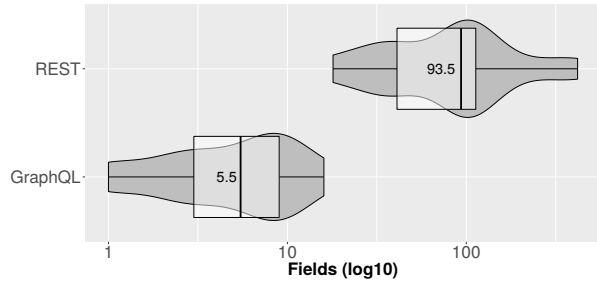


Fig. 6. Number of fields returned by REST and GraphQL calls

V. RUNTIME EVALUATION

In the previous section, we evaluated the number of unique fields returned by GraphQL-based APIs. However, the results were based on executing the systems with trivial input data. A more realistic execution requires the definition of a representative sample of users, repositories, preprints, etc; which is not trivial. For example, some systems retrieve the list of followers of a given user. The size of this list depends on the selected users (most GitHub users have few followers, but others have thousands of followers). Therefore, we postponed an evaluation of the runtime gains achieved by GraphQL to this section, where we ask this research question:

RQ5: When using GraphQL, what is the reduction in the size (in bytes) of the JSON documents returned by servers?

A. Study Design

To answer *RQ5*, we abandoned the idea of defining a realistic sample of input data for the migrated systems. Instead, we rely on real and precisely defined queries used on recent empirical software engineering papers. Typically, these papers depend on a dataset to evaluate their object of study. Therefore, we first retrieve a list of papers published in two recent and relevant software engineering conferences: International Conference on Software Engineering (ICSE, 2017 edition) and Mining Software Repositories Conference (MSR, 2017 edition). Then, we selected three papers from ICSE and four papers from MSR that rely on GitHub to create a dataset with data about open source projects. The advantage is that these papers carefully describe the criteria they use to select the projects and the data (fields) they collect for each one. For example, a paper by Floyd et. al [15] selects the top 100 projects whose *main language is C*. For each project, they collect the full name and the 1,000 most recent pull requests. For each pull request, they retrieve the number of edited files and the comments (*we considered the pull requests ... with at most two edited files ... [and] non-empty developer comments* [15]). In other words, the paper precisely specifies the amount of data retrieved from GitHub (100 C projects, 1,000 PRs per project, etc).

After selecting the papers, we carefully implemented queries to collect the datasets, first using GitHub’s REST API and then using the GraphQL API. The GraphQL queries retrieve

TABLE IV
PAPERS AND QUERIES

Paper	Query	Data
Floyd et al. [15]	Q1	Name of the top-100 C projects by stars
	Q2	For each project: <i>total number</i> and <i>body</i> of the 1K most recent PRs
	Q3	For each PR: <i>body</i> of comments
Xiong et al. [16]	Q4	Name and URL of the top-5 projects by stars (in any programming language)
Ma et al. [17]	Q5	For seven projects: <i>number of commits, branches, bugs, releases</i> and <i>contributors</i>
	Q6	For each project: <i>title, body</i> of closed bugs
	Q7	For each closed bug: <i>body</i> of comments
Osman et al. [18]	Q8	Name and URL of Java projects created before Jan, 2012, with 10+ stars, and 1+ commits
Zampetti et al. [19]	Q9	Number of stars of specific projects
Macho et al. [20]	Q10	Name of repositories with at least 1K stars
	Q11	Number of commits in a repository
Wan et al. [21]	Q12	For eight projects: number of <i>releases, stars, and language</i>
	Q13	Title, body, date and project name of open issues tagged with a bug tag
	Q14	For each issue: <i>body</i> of comments

only the data used in the papers. In total, we reimplemented 14 queries (denoted by Q1 to Q14, see Table IV), which are used by seven papers. Finally, we executed the queries and computed the size in bytes of the returned JSON documents.

B. Results

Figure 7 shows the size of the JSON documents returned by the selected queries, when implemented in REST and GraphQL. In almost all queries, there is a drastic difference after migrating to GraphQL. For example, when using REST, Q11 returns JSON documents that add up to almost 400 MB; when running the same query in GraphQL the size of the answer drops to 77 KB. This query counts the number of commits in a repository. In GraphQL, lists have a field called *totalCount* that returns their size (this field is similar to a COUNT function in SQL, for example). Therefore, it is straightforward to recover the total number of commits in the master branch of a given repository, using GraphQL, as presented in the following listing:

```

1 query totalCountCommits($owner:String!, $name:String!) {
2   repository(owner:$owner, name:$name) {
3     ref(qualifiedName: "master") {
4       target {
5         ... on Commit {
6           history {
7             totalCount
8           }
9         }
10      }
11    }
12  }
13 }

```

By contrast, using the REST API, the client needs to receive data about all commits and then locally compute the number of commits. The *totalCount* field also explains the




Fig. 7. RQ5: Size of JSON documents returned by API calls (REST vs GraphQL) per query




Fig. 8. Size of JSON documents returned by REST and GraphQL calls

reduction in the size of the JSON responses in queries Q5 (from 5.4 MB to 1.7 KB), Q12 (from 62.2 to 1.1 KB), and Q13 (from 14 MB to 205 KB). In the remaining queries, the papers only need a small subset of the fields in the returned documents. For example, in Q1 only the repositories’ names are needed; the remaining fields are discarded. Figure 8 shows violin plots with the distribution of the size of the JSON documents returned by REST and GraphQL. The REST responses have around 9.8 MB (median values), against only 86 KB after moving to GraphQL. The 1st quartile measures are 1.5 MB (REST) and 2.2 KB (GraphQL); the 3rd quartiles are 85 MB (REST) and 699 KB (GraphQL)

RQ5’s summary: When comparing the size of the JSON documents returned by REST and GraphQL calls—implemented to reproduce queries performed in recent empirical software engineering papers—we observed a major difference, from 9.8 MB (REST) to 86 KB (GraphQL), on the median; which represents a reduction of 99%. As in RQ4, this difference happens due to the over-fetching problem typical of REST clients, which receive several fields they do not need at all. This problem is amplified in queries that only need to compute the number of elements in lists of commits, releases, and branches.

VI. THREATS TO VALIDITY

In this section, we present threats to validity, separated by the three studies conducted in this paper.

Grey Literature Review: In this first study, we only review articles discussed on Hacker News. Although it is a popular technology news aggregator, we might have missed interesting articles that did not appear on this site. Further, the open coding protocol to identify key characteristics, advantages, and disadvantages of GraphQL was performed by a single paper’s author. Therefore, he might have missed important codes or incorrectly classified some of the articles discussions. However, this threat is minimized by two facts. First, because the number of reviewed articles is not high (28 articles). Second, because the classification was partially reviewed and checked by the paper’s third author.

Migration Study: First, the study is based on seven clients, of two APIs, which therefore should be considered before generalizing the presented results. Second, the GraphQL wrapper for arXiv’s API cover only two endpoints. Finally, the migration from REST to GraphQL was manually performed by one of the paper’s author and it is therefore error-prone. To minimize this threat, we performed functional tests in all systems, after migration to guarantee their behavior. We are also making the source code publicly available, to allow inspection, replication, and testing by other researchers and by practitioners.

Runtime Evaluation: In this study, we consider two software engineering conferences: a general conference (ICSE) and a topic-specific conference (MSR), whose papers normally depend on large datasets. Furthermore, the queries documented in these papers might not be representative of real data retrieved by software applications. In fact, since the studied papers depend on large datasets, the amount of data consumed by them would probably compare with the data retrieved by an application over days or weeks. Finally, we reimplemented (and not reused the code) of the queries, which is an error-prone task. Particularly, in the case of the GraphQL queries, we had to define exactly the data (fields) used in the papers, which is also subjected to errors and (in some cases) subjective interpretation. To reduce this threat, this task was performed by two authors, who read the papers independently and them discussed together the data effectively consumed by them.

VII. RELATED WORK

Research on simple and easy-to-use programmatic interfaces to computer systems dates from the 70s. For example, Query by Example (QBE) [22] was proposed in mid-1970s to facilitate writing queries to database systems. QBE allows users to specify the fields they want to recover from a relational database, by filling a template form and therefore without having to write SQL code. To some extent, GraphQL shares the same goals of QBE, but putting less emphasis on the presence of a graphical interface to formulate the queries. Tuple spaces—as proposed by Linda [23], [24], in the 80s—is another data structure to facilitate the access to a computer system by distributed and parallel clients. When using Linda, clients communicate by inserting (*out*), reading (*rd*), or removing (*in*) ordered sequences of data, called tuples, from a centralized data structured (the tuple space). Clients perform queries (*in* or *rd*) by means of a template, where wild cards designate any value. However, unlike supported by GraphQL, all fields are returned when a matching tuple is found in the server. In the early 2000s, REST (REpresentational State Transfer) [2]–[4] was proposed as a set of principles and architectural styles for implementing APIs based on Web standards and protocols, such as HTTP and URIs. For example, in REST-based architectures, all resources have URIs and communication is fully stateless. Due to its flexibility, robustness, and scalability, REST is largely used by major Internet companies to implement Web-based APIs. However, REST interfaces—in order to reduce the need of frequent access by clients—tend to rapidly become coarse-grained services. As a result, clients tend to receive superfluous data as a result of REST calls. This problem—called *over-fetching*—was the main motivation for GraphQL design.

Despite its recent popularity, GraphQL is an understudied topic in the scientific literature. In a workshop paper, Hartig and Perez were one of the first to study and provide a formal definition for GraphQL [9]. Later, they complemented and finished this formalization in a conference paper [5]. In this second paper, the authors also prove that evaluating the complexity of GraphQL queries is a NL-problem (i.e., a decision problem that can be solved by a nondeterministic Turing machine under a logarithmic amount of memory). In practical terms, this result shows that it is possible to implement efficient algorithms to estimate the complexity of GraphQL queries before their execution; which is important for example to handle the performance problems normally associated to GraphQL (as reported in our grey literature review) and particularly to avoid denial-of-service attacks. Vogel et al. [25] present a case study on migrating to GraphQL part of the API provided by a smart home management system. They report the runtime performance of two endpoints after migration to GraphQL. For the first endpoint, the gain was not relevant; but in the second one GraphQL required 46% of the time required by the original REST API. The authors also comment that parallel operation of REST and GraphQL services is possible without restrictions. Romano et al. [26]

propose a genetic algorithm for refactoring “fat interfaces”, i.e., coarse-grained interfaces whose clients rely on different subsets of their methods. The authors argue that such interfaces should be refactored into fine-grained interfaces, containing only methods effectively called by groups of clients. Therefore, they focus on superfluous methods, while GraphQL focus on superfluous data returned by REST-based APIs. Wittern et al. [27] propose a tool to generate GraphQL wrappers from REST-like APIs with OpenAPI Specification (OAS). Their tool takes as input a specification that describes a REST API and generates a GraphQL wrapper. They evaluate the proposed tool with 959 publicly available REST APIs; and it was able to generate GraphQL wrappers for 89.5% of these APIs, but with limitation in some cases.

VIII. CONCLUSION

As our key finding, we show that there is a drastic reduction in the number of fields and size of the returned JSON documents when using GraphQL, instead of REST. Probably to avoid frequent client/server interactions [28] (or to avoid the implementation of slightly different endpoints), REST-based interfaces are usually coarse-grained components, designed to provide at once all possible data needed by clients. However, specific clients require only a small subset of the data provided by such interfaces; and therefore simply discard the unneeded information. Our results show that the proportion of data received but discarded by clients is outstanding: GraphQL can reduce the size of the JSON documents returned by REST-based APIs in 94% (measured in number of fields) and in 99% (measured in bytes); both measures are median values. To our knowledge, we are the first to reveal such numbers, by means of a study involving 24 queries performed by seven open source clients of two popular REST APIs (GitHub and arXiv) and 14 queries performed by seven recent empirical papers published in two software engineering conferences.

As our secondary finding, we show that it is not straightforward to refactor API clients to use complex GraphQL queries. The reason is that developers tend to organize their code around small functions that consume small amounts of data. Refactoring these programs to request at once large graph structures is probably a complex reengineering task.

Our work paper can be extended as follows: (a) by evaluating the runtime performance of GraphQL queries, particularly the ones used in Section V; (b) by interviewing developers to reveal their views and experience with GraphQL; (c) by migrating more systems to GraphQL and studying the logs they produce during normal operation; (d) by investigating the benefits of GraphQL in specific domains, such as mobile applications and microservices orchestration [29].

The dataset used in this paper—including the articles of the grey literature, the source code of the migrated systems, and the queries used in the runtime evaluation—is publicly available at <https://github.com/gleisonbt/migrating-to-graphql>.

ACKNOWLEDGMENTS

Our research is supported by CNPq, CAPES, and FAPEMIG.

REFERENCES

- [1] Facebook Inc., “GraphQL specification (draft),” <https://facebook.github.io/graphql/draft/>, 2015, [accessed 15-October-2018].
- [2] R. T. Fielding and R. N. Taylor, “Principled design of the modern Web architecture,” *ACM Transactions on Internet Technology (TOIT)*, vol. 2, no. 2, pp. 115–150, 2002.
- [3] —, “Principled design of the modern web architecture,” in *22nd International Conference on Software Engineering (ICSE)*, 2000, pp. 407–416.
- [4] R. T. Fielding, “Architectural styles and the design of network-based software architectures,” Ph.D. dissertation, University of California, 2000.
- [5] O. Hartig and J. Pérez, “Semantics and complexity of GraphQL,” in *27th World Wide Web Conference on World Wide Web (WWW)*, 2018, pp. 1155–1164.
- [6] R. T. Ogawa and B. Malen, “Towards rigor in reviews of multivocal literatures: Applying the exploratory case study method,” *Review of Educational Research*, vol. 61, no. 3, pp. 265–286, 1991.
- [7] V. Garousi, M. Felderer, and M. V. Mäntylä, “The need for multivocal literature reviews in software engineering: complementing systematic literature reviews with grey literature,” in *20th International Conference on Evaluation and Assessment in Software Engineering (EASE)*, 2016, p. 26.
- [8] —, “Guidelines for including the grey literature and conducting multivocal literature reviews in software engineering,” *arXiv preprint arXiv:1707.02553*, 2017.
- [9] O. Hartig and J. Pérez, “An initial analysis of Facebook’s GraphQL language,” in *11th Alberto Mendelzon International Workshop on Foundations of Data Management and the Web (AMW)*, 2017, pp. 1–10.
- [10] T. Barik, B. Johnson, and E. Murphy-Hill, “I heart Hacker News: expanding qualitative research findings by analyzing social news websites,” in *10th Foundations of Software Engineering Conference (FSE)*, 2015, pp. 882–885.
- [11] M. Aniche, C. Treude, I. Steinmacher, I. Wiese, G. Pinto, M.-A. Storey, and M. A. Gerosa, “How modern news aggregators help development communities shape and share knowledge,” in *40th International Conference on Software Engineering (ICSE)*, 2018, pp. 499–510.
- [12] P. Leitner, E. Wittern, J. Spillner, and W. Hummer, “A mixed-method empirical study of function-as-a-service software development in industrial practice,” *PeerJ PrePrints*, vol. 6, pp. 1–24, 2018.
- [13] A. Brito, L. Xavier, A. Hora, and M. T. Valente, “Why and how Java developers break APIs,” in *25th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2018, pp. 255–265.
- [14] L. Xavier, A. Brito, A. Hora, and M. T. Valente, “Historical and impact analysis of API breaking changes: A large scale study,” in *24th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2017, pp. 138–147.
- [15] B. Floyd, T. Santander, and W. Weimer, “Decoding the representation of code in the brain: An fMRI study of code review and expertise,” in *39th International Conference on Software Engineering (ICSE)*, 2017, pp. 175–186.
- [16] Y. Xiong, J. Wang, R. Yan, J. Zhang, S. Han, G. Huang, and L. Zhang, “Precise condition synthesis for program repair,” in *39th International Conference on Software Engineering (ICSE)*, 2017, pp. 416–426.
- [17] W. Ma, L. Chen, X. Zhang, Y. Zhou, and B. Xu, “How do developers fix cross-project correlated bugs? a case study on the GitHub scientific Python ecosystem,” in *39th International Conference on Software Engineering (ICSE)*, 2017, pp. 381–392.
- [18] H. Osman, A. Chiş, C. Corrodi, M. Ghafari, and O. Nierstrasz, “Exception evolution in long-lived Java systems,” in *14th International Conference on Mining Software Repositories (MSR)*, 2017, pp. 302–311.
- [19] F. Zampetti, S. Scalabrino, R. Oliveto, G. Canfora, and M. D. Penta, “How open source projects use static code analysis tools in continuous integration pipelines,” in *14th International Conference on Mining Software Repositories (MSR)*, 2017, pp. 334–344.
- [20] C. Macho, S. McIntosh, and M. Pinzger, “Extracting build changes with builddiff,” in *14th International Conference on Mining Software Repositories (MSR)*, 2017, pp. 368–378.
- [21] Z. Wan, D. Lo, X. Xia, and L. Cai, “Bug characteristics in blockchain systems: a large-scale empirical study,” in *14th International Conference on Mining Software Repositories (MSR)*, 2017, pp. 413–424.
- [22] M. M. Zloof, “Query-by-example: A data base language,” *IBM Systems Journal*, vol. 16, no. 4, pp. 324–343, 1977.
- [23] D. Gelernter, “Generative communication in Linda,” *ACM Transactions on Programming Languages and Systems (TOPLAS)*, vol. 7, no. 1, pp. 80–112, 1985.
- [24] N. Carriero and D. Gelernter, “Linda in context,” *Communications of the ACM*, vol. 32, no. 4, pp. 444–458, 1989.
- [25] M. Vogel, S. Weber, and C. Zircpins, “Experiences on migrating RESTful Web Services to GraphQL,” in *15th International Conference on Service-Oriented Computing (ICSOC)*, 2017, pp. 283–295.
- [26] D. Romano, S. Raemaekers, and M. Pinzger, “Refactoring fat interfaces using a genetic algorithm,” in *30th IEEE International Conference on Software Maintenance and Evolution (ICSME)*, 2014, pp. 351–360.
- [27] E. Wittern, A. Cha, and J. A. Laredo, “Generating GraphQL-Wrappers for REST (-like) APIs,” in *International Conference on Web Engineering*, 2018, pp. 65–83.
- [28] S. Baker and S. Dobson, “Comparing service-oriented and distributed object architectures,” in *7th International Symposium on Distributed Objects and Applications (DOA)*, 2005, pp. 631–645.
- [29] P. Jamshidi, C. Pahl, N. C. Mendonça, J. Lewis, and S. Tilkov, “Microservices: The journey so far and challenges ahead,” *IEEE Software*, vol. 35, no. 3, pp. 24–35, 2018.