

# Restats: A Test Coverage Tool for RESTful APIs

Davide Corradini\*, Amedeo Zampieri†, Michele Pasqua‡ and Mariano Ceccato§

Department of Computer Science

University of Verona – Verona, Italy

Email: \*davide.corradini@univr.it, †amedeo.zampieri@studenti.univr.it, ‡michele.pasqua@univr.it, §mariano.ceccato@univr.it

**Abstract**—Test coverage is a standard measure used to evaluate the completeness of a test suite. Coverage is typically computed on source code, by assessing the extent of source code entities (e.g., statements, data dependencies, control dependencies) that are exercised when running test cases. When considering REST APIs, an alternative perspective to assess test suite completeness is with respect to the service definition.

This paper presents Restats, a test coverage tool for REST APIs that supports eight state-of-the-art test coverage metrics with a black-box perspective, i.e., only relying on the OpenAPI interface specification of the REST API under test. In fact, metrics are computed by only observing the HTTP requests and responses occurring at testing time, and no access to source/compiled code of the REST API is required.

These coverage metrics come in handy for: (i) developers and test engineers working at development and maintenance tasks; (ii) stakeholders and customers who want to evaluate the completeness of acceptance tests; (iii) researches interested in comparing different automated test case generation strategies.

Restats GitHub repository: <https://github.com/SeUnivr/restats>  
Restats demo video: <https://smarturl.it/restats-demo>

**Keywords**—REST API, Test coverage, Software testing

## I. INTRODUCTION

A RESTful API (or REST API for short) is an API that respects the REST (REpresentational State Transfer) architectural style [1]. REST APIs provide a uniform interface to create, read, update and delete (CRUD) a resource. A resource is generally identified by an HTTP URI, and CRUD operations are usually mapped to the HTTP methods POST, GET, PUT and DELETE to the resource URI.

REST APIs are becoming a de-facto industrial standard to interconnect different computer systems. They are used in a quite wide set of contexts, e.g., when exchanging data with the cloud [2], when connecting smartphone apps to their corresponding server [3], for identity provisioning [4] and when inter-operating different banks [5].

The correct integration among computer systems, potentially carried out by different parties, is a critical point. To monitor correct integrations and promptly reveal implementation defects, several tools are available for testing computer systems through their REST APIs. These tools include not just extensions to general purpose testing frameworks (e.g., REST-Assured [6], that extends the popular JUnit testing framework with REST API specific features), but also brand-new tools, explicitly developed to test REST APIs (e.g., Postman [7] and SoapUI [8]). However, they still require developers to spend valuable time in manually coding test cases. Additionally, to mitigate the cost of manually writing all the test cases, the

research community proposed several approaches to automatically generate test cases for REST APIs (e.g., QuickREST [9], RESTler [10] and RestTestGen [11]).

Despite several approaches being available, either to generate and run test cases on REST APIs, to the best of our knowledge no tool is available to *measure* test coverage, i.e., to quantify the extent of a REST API that has been subject to testing and what part of it still requires more tests. To fill this gap, we propose Restats, a tool meant to monitor the whole REST API testing process and compute the corresponding comprehensive coverage report.

Restats is meant to be helpful in many contexts.

- Supporting software developers when developing and/or evolving software systems that expose a REST API. In fact, before publishing an update, a REST API should be carefully assessed and its tests should meet an adequate level of coverage.
- Supporting stakeholders and customers in evaluating REST APIs delivered by contractors or by third parties, before fully committing and adopting them. In fact, stakeholders might be interested in assessing the coverage of their acceptance tests before accepting a product.
- Supporting researchers in conducting experimental surveys, when comparing different tools and alternative strategies to automatically generate test cases for REST APIs. In fact, an approach that delivers test cases with higher coverage might have potentially higher chances of exposing defects in the implementation of a REST API.

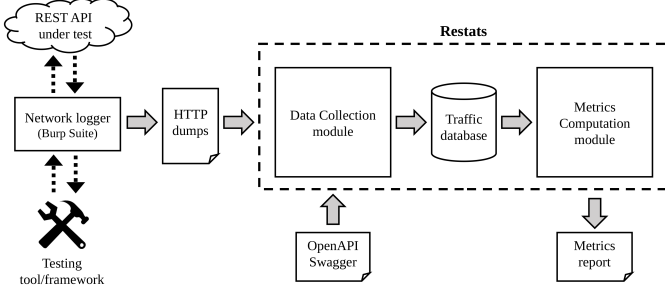
Despite theoretical metrics have been proposed to measure REST APIs test coverage [12], no implementation is available to actually compute them. So, although theoretically sound, in practice REST APIs test coverage metrics convey limited value to developers and practitioners. To the best of our knowledge, Restats is the first tool that computes coverage metrics for REST API test suites.

This paper is organized as follows. In Section II we present the architecture of Restats and the implemented coverage metrics. In Section III we show how to use Restats on a case study, reporting the output of the tool (i.e., the computed coverage metrics) for the REST API under test. Section IV presents related work and, finally, Section V closes the paper.

## II. ARCHITECTURE

Restats computes REST API test coverage metrics based on the theoretical measurement framework proposed by Martin-Lopez et al. [12]. An overview of the architecture of Restats

Fig. 1. Restats architecture overview.



is shown in Figure 1. The tool requires as input the REST API definition file (called OpenAPI specification or Swagger) and a network log containing the HTTP messages exchanged between the API and the testing tool. An external *Network logger* is in charge of collecting these messages. In particular, we used the HTTP proxy offered by the Burp Suite [13].

The network log is consumed by the *Data Collection module*, in which raw inputs are rearranged, processed and written to a database. Eventually, the *Metrics Computation module* computes the coverage metrics using data from the database and fills in a final coverage report.

#### A. Network Logger

To collect testing input and output, we use the *Burp Suite* [13]. It works as a proxy and collects all the HTTP messages exchanged between the testing tool and the REST API under test. Then, using the *burp-dump* [14] extension, we dump all HTTP requests and HTTP responses as plain text. Note that Restats is not dependent on the Burp Suite; any alternative tool that can emit a plain-text log of HTTP messages is fine. The user can choose an alternative component that best fits his/her needs to gather request/response samples, as long as they can be dumped as text file(s).

HTTP messages dumps are expected to follow a naming convention to pair requests and responses in the form (*n-request.txt*, *n-response.txt*), with a progressive identifier *n*.

#### B. Data Collection Module

This module is used to parse network dumps. Given the textual HTTP dumps as input, it extracts all the information needed for the coverage computation (e.g., request paths, HTTP methods, response status codes). Then, it populates a SQLite database.

To this aim, HTTP requests and responses are parsed with the following strategies.

*Message payload.* HTTP requests and responses might contain arbitrary content types in their body, either text (e.g., JSON, XML) or binary (e.g., JPEG, MP4), leading to possible message body parsing failures. Therefore, message body is safely parsed as a binary content. A second parser is called for the message body, depending on the content type declared in the message header. At the moment, only JSON payloads are supported. Hence,

if a different format is found in the header, the payload is ignored.

*Unpaired requests.* Not all HTTP requests match a corresponding HTTP response. In fact, some requests could have been ignored due to a too high server load or a server crash. Even if unpaired, these orphan requests are still recorded in the database. Indeed, they are used for the computation of input coverage metrics, which do not require a matching response.

*Path parameter matching.* The OpenAPI standard allows path templating, i.e., they allow to specify input data as part of the called URL. Therefore, Restats has to be able to group and recognize HTTP requests belonging to the same path template. This task is challenging to carry out in the case of OpenAPI specifications with colliding paths. For instance, */users/{id}* (where *{id}* is a placeholder for any string) and */users/auth* seems referring to the same path, and *auth* could be interpreted as both a specific path and a user id.

All the fields of HTTP requests and responses are then saved in the database to be quickly accessed by the subsequent module. We decided to use a SQLite database because of its performance, portability and configuration-less features.

#### C. Metrics Computation Module

This module takes the database populated by the Data Collection module and the OpenAPI specification as input. The database is used to recognize which elements of the REST API have been tested, while the OpenAPI is used to list all the potentially testable elements. The metrics are computed as the ratio of the number of tested elements to the total number of elements documented in the OpenAPI specification.

Following the approach of Martin-Lopez et al. [12], metrics are divided into input and output metrics. Input metrics are computed based on the observed HTTP requests sent to the REST API under test, and comprise the following.

- **Path coverage:** ratio of the number of tested paths to the total number of paths documented in the OpenAPI specification.
- **Operation coverage:** ratio of the number of tested operations to the total number of operations.
- **Parameter coverage:** ratio of the number of input parameters used by test cases to the total number of parameters.
- **Parameter value coverage:** ratio of the number of the exercised parameter values to the total number of possible values that parameters can assume. This metric applies, at the moment, only to domain-limited parameters such as `boolean` and `enum` types.
- **Request content-type coverage:** ratio of the number of tested content-types to the total number of accepted content-types. This metric is computed only when operations content-types have no wildcard (e.g., `application/*`), because otherwise the number of accepted content-types would be unbounded.

Similarly, the following output metrics are computed based

on the observed HTTP responses, emitted by the REST API under test.

- **Status code class coverage:** a test suite reaches 100% status code class coverage when it is able to trigger both correct status code (i.e., the 2XX class) and erroneous status codes (i.e., 4XX and 5XX classes).
- **Status code coverage:** the ratio of the number of obtained status codes to the total number of status codes documented in the OpenAPI specification.
- **Response content-type coverage:** the ratio of the number of obtained content-types to the total number of response content-types documented in the OpenAPI specification. As for the input case, this metric is computed only when specific content-types are defined with no wildcard.

Restats also automatically computes Test Coverage Levels (TCLs), as defined in Martin-Lopez et al. [12]. The TCL is a summary coverage indicator, ranking test suites based on the score they obtain for the defined metrics. TCLs are incremental: to achieve a TCL, the test suite also has to satisfy the requirements for the lower levels. In particular:

*TCL0* is the base level, with no coverage requirements;

*TCL1* requires the path coverage to be 100%;

*TCL2* requires the operation coverage to be 100%;

*TCL3* requires both request and response content-type coverages to be 100%;

*TCL4* requires parameters and status code classes coverages to be 100%;

*TCL5* requires the status code coverage to be 100%;

*TCL6* requires the body properties coverage to be 100%;

*TCL7* requires the operation flow coverage to be 100%.

### III. RUNNING EXAMPLE

To better understand the use of Restats, we now present a running example involving a REST API and its (manually written) test suite, for which Restats will compute coverage metrics.

#### A. The Pet Store and Its Test Suite

The case study REST API is *Pet Store*, a service to manage a collection of pets. The specification of the API describes the following 7 operations, arranged in 4 endpoints:

*POST /pet* to enter data of a new pet;

*GET /pet* to retrieve a list of all stored pets;

*PUT /pet* to update data of a pet;

*GET /pet/findByStatus* to search pets by status;

*GET /pet/findByTags* to search pets by tag;

*GET /pet/{id}* to retrieve a specific pet;

*DELETE /pet/{id}* to delete all the data for a specific pet.

We have manually prepared an example test suite for the Pet Store API consisting of five test cases based on the business logic of the API, one of which is reported in Listing 1. They are five HTTP requests that can be executed, for instance, by means of Postman [7].

The OpenAPI specification of the Pet Store and the dump of the HTTP requests/responses generated by the test suite are available in the `example` folder of the Restats repository.

```
POST /v2/pet HTTP/1.1
Host: localhost:8080
Accept: application/json
Content-Type: application/json
Content-Length: 76

{
  "name": "doggie",
  "photoUrls": ["myphoto.com/doggie"]
}
```

Listing 1. A test case of our test suite.

The first two tests exercise the `GET /pet` operation. They obtain a successful response (200 status code). The third and fourth tests exercise the `POST /pet` operation, and they obtain a 200 and 500 status code, respectively. The 500 internal server error is caused by a request (in Listing 1) that misses the `category` mandatory parameter. This erroneous request is not properly handled by the Pet Store API, that throws an exception. Finally, a fifth request to the `PATCH /pet` operation obtains a 405 status code because such endpoint does not actually support the `PATCH` method: the test engineer would have wanted to use the `PUT` method instead, that operates similarly.

#### B. Configuring and Launching Restats

The five requests of our test suite, and the five related responses, have been recorded by the Burp proxy. Using the burp-dump plugin we exported the plain-text dumps to the `dumps` folder. Afterwards, we configured Restats with the following JSON configuration file (`config.json`):

```
{
  "modules": "all",
  "specification": "/path/to/petstore.json",
  "dumpsDir": "/path/to/dumps",
  "reportsDir": "/path/to/reports",
  "dbPath": "/path/to/database.sqlite"
}
```

We configured the tool with the modules to run (`modules`); where it can find the OpenAPI specification of the REST API under test (`specification`); where we have stored the dumps exported from Burp (`dumpsDir`<sup>1</sup>); where we want the output reports to be stored (`reportsDir`); and where to place the SQLite database (`dbPath`).

With the `modules` parameter set to `all`, we chose to run all the Restats modules. It is possible to run individually either the data collection module (with `dataCollection`), or the metrics computation module (with `statistics`).

Eventually, we have launched the tool from the folder `restats`, with the command: `python3 app.py`.

#### C. Restats Reports

Once the execution of Restats is complete, its output is available in the chosen report folder as a collection of nine report JSON files. The file `stats.json` summarizes the computed statistics for all the coverage metrics and the TCL

<sup>1</sup>Absolute paths should be preferred for higher compatibility.

```

{
  ...
  "operationCoverage": {
    "raw": {
      "documented": 7,
      "documentedAndTested": 2,
      "totalTested": 3
    },
    "rate": 0.2857142857142857
  },
  ...
  "TCL": 0
}

```

Listing 2. Part of `stats.json` regarding the operation coverage metric.

of the test suite. The other eight files, named after the eight coverage metrics, describe in detail the output for each metric.

Listing 2 shows a fragment of `stats.json` with the operation coverage metric and the computed TCL. The remaining metrics are not shown for space reasons, however, they have the same structure. Observing Listing 2, we can note that two operations, out of the seven documented in the specification, have been tested. Indeed, the computed `rate` is 0.2857, that corresponds to  $\frac{2}{7}$ . In particular, the `documented` field reports the number of operations documented in the specification. The `documentedAndTested` field reports the number of operations that have been tested among those documented in the specification. Finally, the `totalTested` field reports the total number of tested operations, including those that are *not* documented in the specification. This last field was included to reveal test cases that operate (knowingly or unknowingly) beyond what is documented.

Listing 3 shows the content of one of the eight detailed reports (specifically, the operation coverage metric report) that Restats emits for each coverage metric. The `documentedAndTested` section reports what operations have been tested, among those documented in the specification. The `documentedAndNotTested` section reports what operations, among those documented in the specification, have *not* been tested. Finally, undocumented operations that were tested, if any, will be reported in the `notDocumentedAndTested` section.

#### D. How Are These Reports Useful?

The reports computed by Restats come in handy in a number of ways. By looking at Listing 2, we could easily identify that not all the available operations in the REST API have been properly tested. Only two operations out of seven have been tested, which prevents the test suite to be considered complete. This should motivate a developer to write more test cases to test, at least once, all the documented operations.

Moreover, the `notDocumentedAndTested` section in Listing 3 indicates that the developer inappropriately used the PATCH method with the `/pet` endpoint in a test case.

## IV. RELATED WORK

In literature, several tools have been presented to measure test coverage for specific domains, like Java programs (the

```

{
  "documentedAndTested": {
    "/pet": ["get", "post"]
  },
  "documentedAndNotTested": {
    "/pet": ["put"],
    "/pet/findByStatus": ["get"],
    "/pet/findByTags": ["get"],
    "/pet/{petId}": ["get", "delete"]
  },
  "notDocumentedAndTested": {
    "/pet": ["patch"]
  }
}

```

Listing 3. The detailed output for the operation coverage metric.

JaCoCo library [15]) or Android apps (ACVTool [16] and COSMO [17]). However, to the best of our knowledge, no tool has been presented for assessing the coverage of REST APIs test cases. This might be due to the fact that, often, coverage tools rely on *code coverage*, which is not always applicable in the context of REST APIs. Indeed, API source code analysis might be complex, e.g., when a micro-service architecture includes many dynamically allocated components possibly from different vendors; or not feasible, e.g., in the case of closed-source APIs. The only work dealing with test coverage of REST APIs is the theoretical framework of Martin-Lopez et al. [12], (presented as background in Section II), but it does not come with a usable implementation.

Concerning REST API testing, we can find semi-automatic tools (like REST-Assured [6], SoapUI [8], Postman [7]), but they do not provide a measure for assessing test case coverage. Similarly, automatic test case generation tools for REST APIs (like QuickREST [9], RESTler [10] RestTestGen [11]), report what operations could be tested and what status code have been observed in the responses, but in a tool-specific way that does not follow a common formal definition of test coverage.

## V. CONCLUSION

In conclusion, the present paper introduces Restats, a tool that implements the test coverage metrics originally proposed by Martin-Lopez et al. [12]. Restats computes the coverage metrics of the REST API under test adopting a black-box perspective, namely relying only on the OpenAPI specification of the REST API under test. In particular, Restats computes the coverage metrics based on the network log that contains the plain-text HTTP requests and responses exchanged when running test cases. This makes our tool easily adaptable to multiple contexts, supporting services written in any programming language, and any testing technology and tool.

Restats can be used to measure test case coverage in multiple adoption scenarios, either by programmers at development time or by stakeholders before deployment in production. Finally, Restats can help researchers in assessing automated test case generation techniques, highlighting strengths and weaknesses of the adopted generative approach. Indeed, Restats has been used in [18] to compare four state-of-the-art black-box REST APIs testing tools.

## REFERENCES

- [1] R. T. Fielding, *Architectural styles and the design of network-based software architectures*. University of California, Irvine Doctoral dissertation, 2000, vol. 7.
- [2] F. Petrillo, P. Merle, N. Moha, and Y.-G. Guéhéneuc, “Are rest apis for cloud computing well-designed? an exploratory study,” in *Service-Oriented Computing*, Q. Z. Sheng, E. Stroulia, S. Tata, and S. Bhiri, Eds. Cham: Springer International Publishing, 2016, pp. 157–170.
- [3] A. Belkhir, M. Abdellatif, R. Tighilt, N. Moha, Y.-G. Guéhéneuc, and E. Beaudry, “An observational study on the state of rest api uses in android mobile applications,” in *2019 IEEE/ACM 6th International Conference on Mobile Software Engineering and Systems (MOBILESoft)*, 2019, pp. 66–75.
- [4] simplecloud.info. SCIM. [Online]. Available: <http://www.simplecloud.info/>
- [5] Open Bank Project. OBP Middleware. [Online]. Available: <https://www.openbankproject.com/openbankingmiddleware/>
- [6] J. Haleby. REST Assured. [Online]. Available: <http://rest-assured.io/>
- [7] Postman Inc. Postman. [Online]. Available: <https://www.getpostman.com/>
- [8] SmartBear Software. SoapUI. [Online]. Available: <https://www.soapui.org/>
- [9] S. Karlsson, A. auevi, and D. Sundmark, “QuickREST: Property-based test generation of OpenAPI-described RESTful APIs,” in *2020 IEEE 13th International Conference on Software Testing, Validation and Verification (ICST)*, 2020, pp. 131–141.
- [10] V. Atlidakis, P. Godefroid, and M. Polishchuk, “Restler: Stateful rest api fuzzing,” *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*, pp. 748–758, 2019.
- [11] E. Viglianisi, M. Dallago, and M. Ceccato, “Resttestgen: Automated black-box testing of restful apis,” *2020 IEEE 13th International Conference on Software Testing, Validation and Verification (ICST)*, pp. 142–152, 2020.
- [12] A. Martín-Lopez, S. Segura, and A. Ruiz-Cortés, “Test coverage criteria for restful web apis,” in *Proceedings of the 10th ACM SIGSOFT International Workshop on Automating TEST Case Design, Selection, and Evaluation*, 2019, pp. 15–21.
- [13] PortSwigger. Burp suite. [Online]. Available: <https://portswigger.net/burp>
- [14] crashgrindrrips. Burp-Dump. [Online]. Available: <https://github.com/crashgrindrrips/burp-dump>
- [15] EclEmma. JaCoCo code coverage library. [Online]. Available: <https://www.eclemma.org/jacoco/>
- [16] A. Pilgun, O. Gadyatskaya, S. Dashevskiy, Y. Zhauniarovich, and A. Kushniarou, “An effective android code coverage tool,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’18. New York, NY, USA: Association for Computing Machinery, 2018, pp. 2189–2191. [Online]. Available: <https://doi.org/10.1145/3243734.3278484>
- [17] A. Romdhana, M. Ceccato, G. C. Georgiu, A. Merlo, and P. Tonella, “COSMO: code coverage made easier for android,” in *14th IEEE Conference on Software Testing, Verification and Validation, ICST 2021, Porto de Galinhas, Brazil, April 12-16, 2021*. IEEE, 2021, pp. 417–423. [Online]. Available: <https://doi.org/10.1109/ICST49551.2021.00053>
- [18] D. Corradini, A. Zampieri, M. Pasqua, and M. Ceccato, “Empirical comparison of black-box test case generation tools for RESTful APIs,” in *21st IEEE International Working Conference on Source Code Analysis and Manipulation, SCAM 2021, Luxembourg City, Luxembourg, September 27 - September 28, 2021*.