

# Sophia: a Linked Data and Semantic Web toolkit for Rust

**Author:** Pierre-Antoine Champin (LIRIS – Université de Lyon)

This presentation is about Sophia<sup>1</sup>, an open-source implementation in the Rust<sup>2</sup> programming language of Linked Data and Semantic Web technologies. It aims to provide Rust developers with basic tools and extensible APIs to take full advantage of these technologies. At the same times, it relies on Rust's strengths to provide an efficient and reliable toolkit for Linked Data and Semantic Web practitioners.

## Why Rust?

Rust is a relatively recent programming language (the first stable compiler was released in 2012), sponsored by Mozilla, and emphasizing performance, reliability and productivity. It is a statically typed compiled language, providing high-level abstractions (such as iterators and generic types) while preserving low-level control and execution speed. Its unique way of managing memory at compile time (thanks to the *borrow checker*) protects programmers against a range of nasty bugs, such as dangling pointer dereferencing or race conditions in concurrent programming. Despite a steep learning curve, Rust has been elected "most lovable programming language" in the Stack Overflow Developer Survey every year since 2016<sup>3</sup>.

## Sophia is efficient

Although it is still at an early stage, Sophia already has very competitive performances, compared to mature implementations such as Jena<sup>4</sup>, librdf<sup>5</sup>, N3.js<sup>6</sup> or RDFlib<sup>7</sup>. Figures 1, 2 and 3 show that Sophia is faster than these implementations, both for loading an RDF graph in memory and for answering simple queries over this graph. They also show that Sophia is the only one among these libraries able to load up to 10<sup>7</sup> triples in memory. Finally, Sophia has the smallest memory footprint (on par with librdf). These very encouraging first results must be credited to the optimization abilities of the Rust

<sup>1</sup> [https://github.com/pchampin/sophia\\_rs](https://github.com/pchampin/sophia_rs)

<sup>2</sup> <https://www.rust-lang.org/>

<sup>3</sup> 2016: <https://stackoverflow.com/insights/survey/2016#technology-most-loved-dreaded-and-wanted>  
2017: <https://stackoverflow.com/insights/survey/2017#most-loved-dreaded-and-wanted>  
2018: <https://insights.stackoverflow.com/survey/2018#most-loved-dreaded-and-wanted>  
2019: <https://insights.stackoverflow.com/survey/2019#most-loved-dreaded-and-wanted>

<sup>4</sup> Implemented in Java, <http://jena.apache.org/>

<sup>5</sup> Implemented in C, <http://librdf.org/>

<sup>6</sup> Implemented in Javascript, <https://github.com/rdfjs/N3.js>

<sup>7</sup> Implemented in Python, [rdflib.readthedocs.org/](https://rdflib.readthedocs.org/)

compiler (itself based on the LLVM<sup>8</sup> infrastructure), as well as to the high quality of the Rust standard library. These results, with the code to reproduce them, are available at [https://github.com/pchampin/sophia\\_benchmark/](https://github.com/pchampin/sophia_benchmark/).




Figure 1: Time (in s) to load an NT file in memory




Figure 2: Memory (in kB) used to load an NT file in memory

<sup>8</sup> <https://llvm.org/>

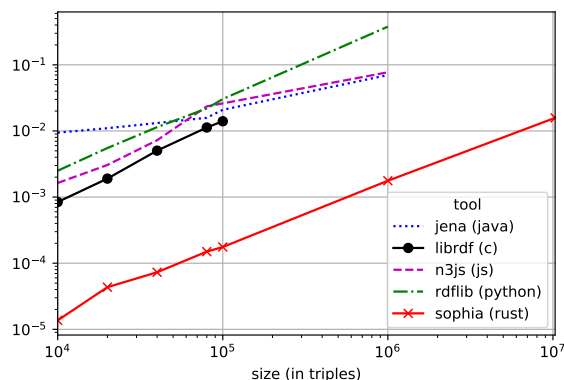


Figure 3: Time (in s) to retrieve all triples matching (\*,p,o)

## Sophia is portable

The Rust compiler can target a number of hardware/software architectures, and has several options for working with lightweight devices (e.g. not linking the standard library, or using a custom memory allocator). This makes Rust a good programming language for the Web of Things, and Sophia could prove valuable in this context.

Furthermore, Rust is the language of choice for compiling to Web Assembly<sup>9</sup>. A suite of tools, maintained by a very active community, makes it very easy to compile and deploy Web Assembly modules written in Rust and able to interoperate with Javascript code. One of our goal is to use Sophia for building efficient in-browser RDF applications.

## Perspectives

Currently, Sophia is able to parse a few RDF syntaxes (including RDF/XML, Turtle and TriG) and to serialize to N-Triples. Triples are stored in memory, and can be queried against simple triple patterns. In the future, we want to support more formats (in particular JSON-LD), persistent triple storages, and other RDF-related standards such as SPARQL, RDFS, OWL, and SHACL.

There is obviously a long way to go, but this can be achieved as a community effort. First, as mentioned above, Sophia is an open-source project, which has already raised interest from a few contributors. Second, Sophia’s goal is not to become a big monolithic project implementing all of these technologies at once. It aims to lay the foundations of a diverse ecosystem of interoperable components. It defines a number of *generic and reusable* interfaces (*traits* in Rust parlance) for all the basic concepts underlying RDF (triple, graph, dataset, etc.), that others can implement. In that sense, it follows the philosophy of the RDF-JS initiative<sup>10</sup>. We have already started to interact with another RDF-related Rust project<sup>11</sup> in order to determine how to best collaborate and complement each other.

<sup>9</sup> <https://webassembly.org/>

<sup>10</sup> <https://rdf.js.org/>

<sup>11</sup> <https://github.com/Tpt/rio>