



Offer #2024-07078

Post-Doctoral Research Visit F/M Automatic porting of vulnerability-fixing patches using Coccinelle

Contract type : Fixed-term contract

Level of qualifications required : PhD or equivalent

Fonction : Post-Doctoral Research Visit

Level of experience : Recently graduated

Context

The position is part of the project "SWHSec: Leveraging Software Heritage to Enhance Cybersecurity", funded by the Programme de Transfert au CampusCyber

Assignment

Objectives: Software projects today commonly co-exist in multiple versions. Some end-users may prefer to stay with older versions for stability, while others adopt the latest versions to take advantage of the newest features. These older versions of the software may be forked, and then customized for specific requirements. The co-existence of these many forks, however, means that when a vulnerability is detected and fixed, many variants of the fix are needed, to protect all known affected forks. The objective of this task is to develop and evaluate techniques for automating the creation of such fix variants. We will use Software Heritage as a source of case studies, as it contains many co-existing forks of software projects, including their complete development history.

Work description: Over the last 20 years, the Whisper team has been developing tools for automating the processing of C (and to a lesser extent C++) software projects. These tools include: (i) Coccinelle, which provides a patch-like domain-specific language SmPL for matching and transforming code, (ii) Prequel, which adapts SmPL to the task of searching for changes in a git history, and (iii) Spinfer], which generalizes change examples (as could be identified using Prequel) into SmPL transformation rules. In this project, we will investigate how to bring together and extend Coccinelle, Prequel, and Spinfer to automate the porting of vulnerability-fixing patches across software forks. Specifically, the main challenge is how to adapt the fix to the specific APIs, data structures, etc. of the target fork. We envision that it will be possible to use Prequel to collect information about the changes that have led to the differences between forks, Spinfer to translate these collected changes into transformation rules, and Coccinelle to apply these rules to the vulnerability fix, to adapt it to the target fork. At the same time, we would like to exploit the collected information to create an explanation for the adaptations, to provide to the user, to help give confidence that the adapted fix will not itself introduce new vulnerabilities. We will initially focus on the Linux kernel, the traditional target of Coccinelle, and then scale up the approach to the wide variety of software available on Software Heritage.

References:

Coccinelle: Julia Lawall, Gilles Muller: Coccinelle: 10 Years of Automated Evolution in the Linux Kernel. USENIX Annual Technical Conference 2018: 601-614

Prequel: Julia Lawall, Derek Palinski, Lukas Gnirke, Gilles Muller: Fast and Precise Retrieval of Forward and Back Porting Information for Linux Device Drivers. USENIX Annual Technical Conference 2017: 15-26

Spinfer: Lucas Serrano, Van-Anh Nguyen, Ferdian Thung, Lingxiao Jiang, David Lo, Julia Lawall, Gilles Muller: SPINFER: Inferring Semantic Patches for the Linux Kernel. USENIX Annual Technical Conference 2020: 235-248

Software Heritage: <https://www.softwareheritage.org/>

Main activities

- Studying recent backports in the Linux kernel to identify the issues that arise, and in particular any errors that have occurred in the backporting process.
- Manually simulating the collection of the information required to perform backporting.
- Automating the identified information collection strategies.
- Designing strategies for converting the collected information into transformation rules.
- Developing tools to support the evaluation of the proposed approach at a large scale.

- Assessing the precision and recall of the proposed approach, on the Linux kernel and on several projects available in Software Heritage
- Writing papers and giving presentations describing the approach and its results.

Skills

Technical skills and level required : The project requires a strong background in program analysis. Some familiarity with algorithms for processing source code, such as clone detection, differencing, clustering, etc. would be appreciated.

Languages : The position requires reading a large amount of C code. Some reading familiarity with C++ could also be beneficial. All of the existing software infrastructure related to the project is written in OCaml.

Other appreciated skills : Good communication skills (spoken and written). The ability to plan a project and work independently based on the plan.

Benefits package

- Subsidized meals
- Partial reimbursement of public transport costs
- Leave: 7 weeks of annual leave + 10 extra days off due to RTT (statutory reduction in working hours) + possibility of exceptional leave (sick children, moving home, etc.)
- Possibility of teleworking and flexible organization of working hours
- Professional equipment available (videoconferencing, loan of computer equipment, etc.)
- Social, cultural and sports events and activities
- Access to vocational training
- Social security coverage

General Information

- **Theme/Domain** : Distributed Systems and middleware
Software engineering (BAP E)
- **Town/city** : Paris
- **Inria Center** : [Centre Inria de Paris](#)
- **Starting date** : 2024-07-01
- **Duration of contract** : 2 years
- **Deadline to apply** : 2024-05-31

Contacts

- **Inria Team** : [WHISPER](#)
- **Recruiter** :
Lawall Julia / Julia.Lawall@inria.fr

About Inria

Inria is the French national research institute dedicated to digital science and technology. It employs 2,600 people. Its 200 agile project teams, generally run jointly with academic partners, include more than 3,500 scientists and engineers working to meet the challenges of digital technology, often at the interface with other disciplines. The Institute also employs numerous talents in over forty different professions. 900 research support staff contribute to the preparation and development of scientific and entrepreneurial projects that have a worldwide impact.

The keys to success

The position is well suited to someone who enjoys looking at a lot of code and code changes, and who is passionate about improving software development and maintenance processes, with a potential impact on the real world.

Warning : you must enter your e-mail address in order to save your application to Inria. Applications must be submitted online on the Inria website. Processing of applications sent from other channels is not guaranteed.

Instruction to apply

Defence Security :

This position is likely to be situated in a restricted area (ZRR), as defined in Decree No. 2011-1425 relating to the protection of national scientific and technical potential (PPST). Authorisation to enter an area is granted by the director of the unit, following a favourable Ministerial decision, as defined in the decree of 3 July 2012 relating to the PPST. An unfavourable Ministerial decision in respect of a position situated in a ZRR would result in the cancellation of the appointment.

Recruitment Policy :

As part of its diversity policy, all Inria positions are accessible to people with disabilities.