Coq Community Survey 2022: Summary of Results

Ana de Almeida Borges, Jean-Rémy Falleri, Jim Fehrle, Emilio Jesús Gallego
Arias, Érik Martin-Dorel, Karl Palmskog, Alexander Serebrenik, Théo
Zimmermann

HAL Id: hal-03914602
https://hal.inria.fr/hal-03914602

Submitted on 28 Dec 2022

Coq Community Survey 2022: Summary of Results

Ana de Almeida Borges, Jean-Rémy Falleri, Jim Fehrle,
Emilio Jesús Gallego Arias, Erik Martin-Dorel, Karl Palmskog,
Alexander Serebrenik, and Théo Zimmermann

The Coq Community Survey 2022 was an online public survey conducted during February 2022. Its main goal was to obtain an updated picture of the Coq user community and inform future decisions taken by the Coq team. In particular, the survey aimed to enable the Coq team to make effective decisions about the development of the Coq software, and also about matters that pertain to the ecosystem maintained by Coq users in academia and industry.

In this presentation abstract, we outline how the survey was designed, its content, and some initial data analysis and directions. Not least due to free-text answers to some questions requiring a more lengthy summary, the full presentation includes additional data and conclusions.

Survey Design, Deployment, and Analysis Process

The survey was designed, deployed, managed, and analyzed by members of a working group (WG), which was formed following a public call for participation by the Coq core team.

Survey questions were created by WG members collaboratively. The WG took inspiration from the previous Coq Community Survey in 2014, and similar more recent surveys for OCaml and other programming languages. We made efforts to adhere to survey best practices, e.g., on ordering demographic questions late to avoid their answers affecting other questions.

The deployed survey was advertised to Coq users primarily via the following Coq community forums: Discourse, Zulip chat, Coq-club mailing list, Coq Twitter, and the Coq website. We also advertised in more general venues, e.g., the Types-announce mailing list. After the survey closed, WG members collaboratively analyzed the response data, primarily using Jupyter. For GDPR conformance, we decided to keep the raw survey data within the EU. Therefore, we only used Inria-hosted tools and restricted raw data access to WG members located in the EU.
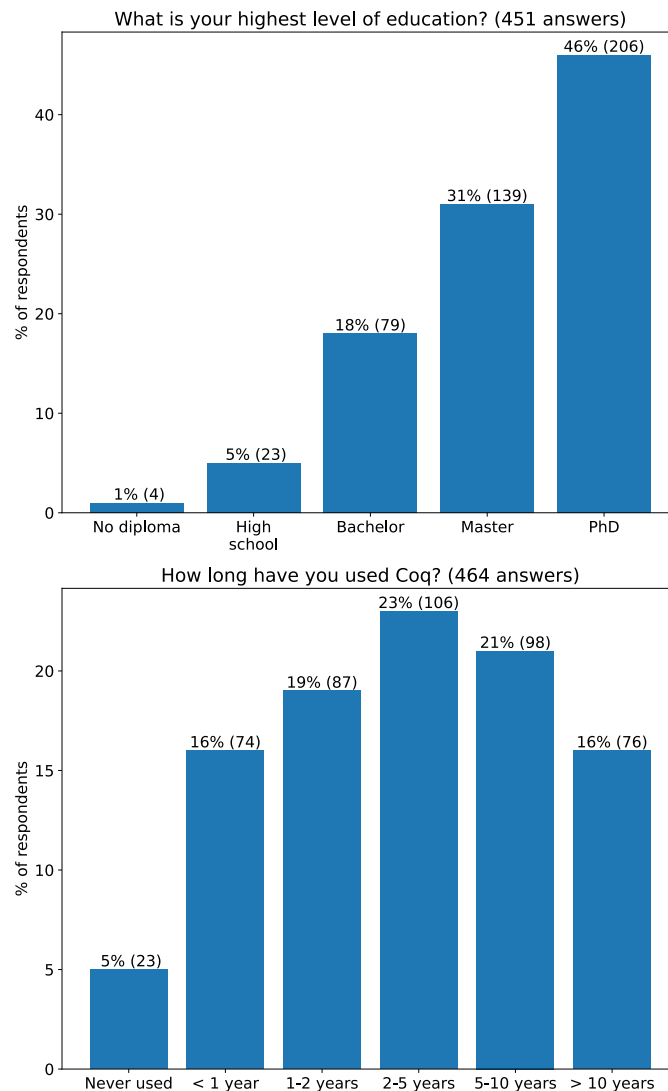
Scope of the Survey

The survey (available in English and Chinese) has 109 questions, some of which were only visible depending on answers to previous questions. Broadly, the survey asked about (i) use of Coq features, IDEs, libraries, plugins, and tools, (ii) views on renaming Coq, on Coq improvements, and issues such as inclusion, and (iii) demographic data, such as age, gender, and experience with Coq and other proof assistants or programming languages. Answering was optional for all survey questions, and most questions allowed free-text answers. In particular, it was possible to answer the survey without having used Coq, albeit with many questions hidden. For example, the following question was hidden unless the respondent had used Coq within the last 6 months:

If you have projects that are not compatible with recent versions of Coq (8.13 or more recent), why is that? Please choose all that apply:

□ It would take too much effort to port them.
□ I don't care that they aren't on the latest version.
□ A dependency of the project is not compatible with recent versions of Coq.
□ Other: _____

Sample of Survey Results

The survey received 466 submitted responses. In the initial analysis of survey responses, the WG looked at the demographics of the Coq users who took the survey, and in what context they are using Coq. *Context* here means users' experience with Coq and similar tools and languages, and their reasons for using Coq. We provide graphs summarizing two such questions below; a full presentation will show many additional graphs. The first graph shows that, as one may expect of a project with many academic users, a large majority of respondents who answered the question on education have master's or PhD degrees. The second graph shows that respondents are fairly evenly distributed into ranges on how long they have used Coq.

Beyond our initial result presentation and analysis, we plan to publish detailed results and conclusions from the survey in academic venues and provide summaries in relevant public forums for the benefit of both Coq community members and Coq core developers. We would like to thank Annalí Casanueva and all members of the WG for their contributions.