



HAL
open science

OPAM for Coq

Guillaume Claret

► **To cite this version:**

| Guillaume Claret. OPAM for Coq. [Technical Report] Inria - Paris 7. 2015. hal-02292537

HAL Id: hal-02292537

<https://hal.science/hal-02292537>

Submitted on 20 Sep 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

OPAM for Coq

Guillaume Claret

2015

Contents

1	Abstract	1
2	Introduction	2
3	OPAM	3
4	Repositories	4
4.1	Stable packages	4
4.2	Development packages	5
4.3	Distribution mechanism	6
5	Contribs	6
6	Bench system	6
6.1	Architecture	7
6.2	Strategies	7
7	Website	8
8	Related work	10
9	Conclusion	11

1 Abstract

MORE AND MORE PEOPLE are using the Coq system to develop theories or verify theorems and softwares. Unfortunately, there are no central places to share Coq developments which offer a unified and mechanized way to install new libraries or to express dependencies between projects. The existence of such a central place would foster the Coq community, by enabling more re-usability, more visibility and more competition among the Coq projects.

We will present our work on setting up a package manager infrastructure for Coq. We reuse the package manager OPAM from the OCaml community and design a specific repository architecture, an automated bench system and a website to present the list of the Coq packages. Today, the OPAM infrastructure is usable and used by more and more people to share their developments. To the best of our knowledge, this infrastructure is the first package management system able to organize mathematical proofs in the large.

2 Introduction

IN THE PROGRAMMING LANGUAGE COMMUNITIES, the package managers have long been considered as one of the fundamental tools to create an ecosystem of re-usable libraries. Noticeable examples are the `npm`¹ package manager for JavaScript, the `Gems`² repository for Ruby or the `CRAN`³ repository for the language for statistics R. It happens that the challenges faced by the users of Coq to install external libraries, share their developments or express complex dependency chains between projects are very similar to those encountered with the programming languages. However and surprisingly, to the best of our knowledge, no modern package management systems have been proposed for a theorem prover.

When I started working on the idea of a package manager for Coq, we realized that we were many trying to reach this goal⁴. As the project started to grow, more people of the community got involved in order to add new ideas and reach a general consensus⁵. Thus, this work on a package manager for Coq is mostly a team work.

Instead of redeveloping a new package manager, we decided to use the existing OPAM package manager from the OCaml community. The Coq system or plugins being implemented in OCaml and most of the OCaml users being used to OPAM, the OPAM package manager appeared as the simplest choice for us, as much technically as socially.

In this chapter we will present what I mostly contributed to:

- an architecture of repositories⁶ which combines a repository for the stable versions, a repository for the development versions and a system of distribution (page 4). Most of the work was actually devoted to the maintenance of these repositories (accounting for more than 400 commits);
- a port of the existing `Contribs` repository⁷, containing more than 150 projects, to the OPAM system (page 6);

¹www.npmjs.com

²rubygems.org

³cran.r-project.org

⁴Mainly Thomas Braibant and Cyril Cohen.

⁵In particular Enrico Tassi and Guillaume Melquiond.

⁶Available on github.com/coq/opam-coq-archive.

⁷The `Contribs` project is hosted on GForge on coq-contribs.gforge.inria.fr.

- a bench system (page 6) to continuously monitor the status of the packages by verifying that they can all be installed without errors and that they respect the policy of the Coq repositories;
- a website (page 8) listing the descriptions of the stable packages available for Coq⁸. We generate this website using the program OpamWebsite⁹, entirely written in Coq thanks to our library Coq.io¹⁰ to implement the concurrent inputs–outputs operations.

3 OPAM

THE PACKAGE MANAGER OPAM¹¹ was introduced by the company OCaml-Pro¹² to simplify the distribution and the installation of OCaml libraries. OPAM is entirely written in OCaml and can use various dependency constraint solvers including aspcud¹³. This is a source package manager, meaning that the packages are installed by compiling the sources.

This package manager is provided to the user through the command `opam`. We first initiate an OPAM folder where we will install our packages:

```
1 mkdir opam
2 opam init --root=opam
```

and setup the environment variables:

```
1 eval 'opam config env --root=opam'
```

Then we have access to the standard OPAM operations:

```
1 opam show coq
2 opam install -v -j4 coq
3 opam list
4 opam search http
5 opam update && opam upgrade
```

By setting up OPAM in different folders, we can setup different development environments. Indeed, two different projects may need different versions of Coq or incompatible libraries. Another equivalent way to handle several development setups is to use the switch mechanism of OPAM, which creates one new folder per switch in the current OPAM folder.

⁸The list of the OPAM packages for Coq is available on coq.io/opam.

⁹The program OpamWebsite is available under MIT license on github.com/coq-io/opam-website.

¹⁰The Coq.io library is available under MIT license on coq.io.

¹¹opam.ocamlpro.com

¹²www.ocamlpro.com

¹³The constraint solver aspcud is available on www.cs.uni-potsdam.de/wv/aspcud.

The OPAM packages are available on some repositories, activated with the command:

```
1 opam repo add repo-name repo-url
```

A repository is composed of one folder per package, containing three files: `descr` for a short textual description, `url` for the download link and `opam` for all the other metadata such as the compilation commands, the license, the homepage, etc.

4 Repositories

WE ORGANIZE THE OPAM REPOSITORIES for Coq¹⁴ into three main repositories plus several distribution repositories. This complex architecture is set to fit the requirements of the users and the wills of all the Coq developers.

4.1 Stable packages

We put the stable packages in the `released` repository, which can be activated by:

```
1 opam repo add coq-released https://coq.inria.fr/opam/released
```

We consider as stable a package whose sources do not change (by opposition to a development branch for example, which is updated at every commits). The stable packages must follow some simple rules¹⁵ in order to be accepted.

Installable The installation of a package must succeed when we start from a fresh OPAM setup. The uninstallation must revert the installation process by removing all the installed files.

Version The version names must be a triple of numbers separated by dots, like for example `3.2.0`. This convention is proven to fit a wide variety of projects, as it is the convention used by every `npm` packages.

The main rationale behind this convention is that, since we express package dependencies using inequalities over version numbers, the ordering over the version names should be as clear as possible. The ordering over triples of numbers in OPAM corresponds to the natural lexicographic order. By contrast, the OPAM ordering¹⁶ over mixes of digits, letters and special characters is much

¹⁴Available on github.com/coq/opam-coq-archive.

¹⁵We continuously verify these rules thanks to the bench system.

¹⁶The OPAM ordering is the Debian ordering, described on www.debian.org/doc/debian-policy/ch-controlfields.html#s-f-Version.

more complex¹⁷. The other rationale is to follow the SemVer guidelines¹⁸, which recommend to use three numbers to identify the *major* changes, *minor* changes and *patch* level changes.

Namespace To separate the Coq packages from the OCaml packages, we introduce a concept of *namespace* for OPAM by using the colon separator ":". In particular, all Coq packages must be prefixed by "coq:" to indicate that they are in the "coq" namespace. This idea of namespaces is just a convention which we can use for other projects too. For example, we prefix all the packages of the Coq.io project by "coq:io:".

Name To ensure uniformity, the package names must be in small caps and use the dash "-" as a word separator. An example of a complete package name with a version name is the following:

```
coq:io:hello - world.1.1.0
```

Description Each package must contain a link to its homepage and provide its license information. The license information is mandatory because so many developers forget about specifying their license, making their projects *de facto* proprietary¹⁹.

Parallel compilation Since OPAM is a source package manager, we must compile each package during the installation. In order to reduce the compilation time, we enforce the use of the `%{jobs}%` option in the build commands. This option represents the number of parallel processes allowed for the compilation, as specified by the `-j` or the `-jobs=` option of OPAM. The `Makefile` generated by the `coq_makefile` command can use this option to speedup the compilation process.

Checksum A package in the stable repository must come with the checksum of its sources, to ensure that the sources cannot be silently updated.

4.2 Development packages

We put the development packages in the two repositories `extra-dev` and `core-dev`:

```
1 opam repo add coq-extra-dev https://coq.inria.fr/opam/extra-dev
2 opam repo add coq-core-dev https://coq.inria.fr/opam/core-dev
```

¹⁷This led for example to the bug github.com/coq/opam-coq-archive/issues/8.

¹⁸semver.org

¹⁹See this study on the license usage for the GitHub repositories: github.com/blog/1964-open-source-license-usage-on-github-com.

The `extra-dev` repository is for the user packages, the `core-dev` repository for the development versions of Coq. Most development packages follow a Git branch of a project, and thus evolve at each new commit. We try to apply the same rules as for stable repository (except for the checksum and the version names, which typically contain strings such as `"beta"` or `"dev"`).

4.3 Distribution mechanism

There is a project of a distribution mechanism for the Coq packages. The idea is to provide a subset of mutually compatible packages, with at most one version per package, following the policy of the Debian distribution. To implement the fact that distributions are subsets of the stable packages, the distribution repositories would be composed of symbolic links to packages of the released repository.

5 Contribs

HISTORICALLY, THE USER PACKAGES of Coq were regrouped on a single SVN repository²⁰, containing the source code of each package. To install a package, someone would download the source code and compile it using a standard:

```
1 make
2 make install
```

The source code of the Contribs was graciously maintained by some of the Coq developers and represents more than 150 projects.

We split the single SVN repository of the Contribs into individual Git repositories with one repository per project, keeping the history of the changes for each project. We created one OPAM package per project, keeping as much metadata as possible. Since the Contribs are not released with a version number, they point to Git branches (with one branch per Coq version in each project) and are hosted on the development repository `extra-dev`. We fixed many of the Contribs because we realized, during the import to OPAM, that the `make install` commands or the generation of the `Makefile` were often broken.

6 Bench system

IN ORDER TO MAKE SURE that no packages are broken, we continuously run a bench system²¹ to check the OPAM packages for Coq. We present the results in

²⁰Hosted on GForge on coq-contribs.gforge.inria.fr.

²¹The results of the bench are available on coq-bench.github.io.

a colored table with the installation times for valid packages. The *best* column contains the best score obtained by each package.

We check that the packages are well-formed, using our lint²² checker. We also check the installation of the dependencies and of the package itself, verifying that the uninstallation effectively reverts the installation process. We record the duration of each operation and compute the size of the installed files.

6.1 Architecture

The bench system is hosted on the GitHub organization `coq-bench`²³. We will detail the four projects of the organization.

Run This Ruby program runs the benchmarks. For each bench we generate a clean Docker²⁴ image. Then we install OCaml, OPAM and Coq and test each package. Either the `released` repository is activated or both the `released` and the `extra-dev` repositories. We save the results into a database.

Database The database is a textual table in the CSV²⁵ format. There is one file per bench and one row per package. The first row gives a legend for each of the 31 columns.

Make-HTML This Ruby program generates static HTML pages representing the content of the CSV database.

Coq-Bench.GitHub.io These are the static HTML pages of the bench website. GitHub provides us a nice and simple way to host web pages from a Git repository using the GitHub Pages²⁶ service.

6.2 Strategies

There are many possible strategies to check all the packages with respect to their installation order. The installation order is important because OPAM will choose to install different dependencies in different contexts. Ideally, we would like to optimize the installation order to reduce the total execution time of the bench, by always installing and testing the dependencies first, so that no packages are compiled twice. Unfortunately this is not possible in general. Here is a simple counter example. With the list of packages presented in the Table 1, we must compile B twice (once with A.1.0.0 and once with A.2.0.0) to test both C.1.0.0 and C.2.0.0. Instead, we provide the two following (non-optimal) installation strategies.

²²The sources of the lint checker are on github.com/coq-bench/run/blob/master/lint.rb.

²³github.com/coq-bench

²⁴Docker is a system simplifying the manipulation of Linux containers. It is available under Apache license on www.docker.com.

²⁵en.wikipedia.org/wiki/Comma-separated_values

²⁶The GitHub Pages services is presented on pages.github.com.

Package	Dependencies
A.1.0.0	\emptyset
A.2.0.0	\emptyset
B.1.0.0	A (any versions)
C.1.0.0	A.1.0.0 and B
C.2.0.0	A.2.0.0 and B

Table 1: Counter example to the optimal installation strategy.

Clean This is the simplest strategy. We install each package in a fresh environment with only `Coq` installed. This strategy is robust and reproducible, but not really optimal for packages with large dependencies since the dependencies are always recompiled from scratch.

Tree This is a more complex strategy. We install as many packages as possible until all new packages are incompatible with the current environment. The main source of incompatibility is the fact that we cannot install two packages with the same name but different version numbers. Once we are blocked, we roll-back until new packages are installable. This strategy is more clever but also more fragile and harder to maintain, to subtle changes of the `opam` tool for example. We use the branch mechanism of `Git` on the `OPAM` installation folder to switch efficiently between `OPAM` states. At the end of the process, we obtain a tree of all the `Git` branches used to explore the space of the packages (Figure 1).

7 Website

WE DEVELOPED A WEBSITE²⁷ to present the list of the `OPAM` packages available for `Coq`. We statically regenerate this website every hours using the program `OpamWebsite`²⁸. We entirely wrote this program using our `Coq.io` library to handle the inputs–outputs operations²⁹. The inputs–outputs operations essentially consist in calling the `opam` command with the right options to grab the metadata of each package and then creating the new `HTML` files. We formally verify the `OpamWebsite` program using our method of formal use cases.

²⁷The website of the `Coq` packages is available on coq.io/opam.

²⁸The program `OpamWebsite` is available under MIT license on github.com/coq-io/opam-website.

²⁹This process of using our own tools is called `dogfooding`.

8 Related work

A LOT OF PROGRAMMING LANGUAGES have their own package managers. Our architecture and the decision to setup a package manager for Coq is influenced by the experience of the other programming language communities. Of course, we mainly got inspiration from the OCaml community, which developed the OPAM package manager. We do not try to compete with other package management systems. Instead, we try to reuse as much existing work as possible to design a simple package platform for Coq.

Historically, the way to share developments in Coq was to use the `Contribs` mechanism (page 6). This approach had some limitations. For example, the original authors of each project could not access to the code hosted on the central SVN repository, introducing a fork of each project. Instead, OPAM separates the metadata from the source code of each package. The `Contribs` also lacked a versioning system and a dependency manager, to automatically compute and install the correct dependency of a project through a command line tool. Finally, the acceptance policy for the `Contribs` was private. By contrast, the OPAM packages are published using public pull-requests on the OPAM repository for Coq³⁰. We believe that all discussions concerning the Coq packages must be public in order to preserve the fairness of the platform.

To the best of our knowledge, no other theorem provers propose a modern package management system. The `ProofPeer`³¹ project aims to enable collaborative works at a large scale, by design a new theorem prover integrating all the tools required for collaboration. We take the opposite direction, by reusing existing tools and concepts from the open source community, such as the idea of a package manager. The `Mizar` system³² proposes the MML library³³, which is a centralized repository including the sources of more than 1200 projects. Each project is published as an article on the *Journal of Formalized Mathematics*³⁴ and peer reviewed before acceptance. Projects can express dependencies to other projects (but with no versioning). The technical architecture of the MML library is similar to the architecture of the `Contribs` repository of Coq, in the sense that all the source code is centralized, and thus suffers from the same limitations. The idea of a peer reviewed conference for the user contributions is interesting.

The OCaml community also made a bench system for the OPAM packages³⁵, but this system was not available at the time we made ours. The OCaml bench has the advantage of presenting a log of the differences between two consecutive benches, what make it easier to read when there are hundreds of packages to monitor. The website listing the OCaml packages for OPAM is written in OCaml.

³⁰ Available on github.com/coq/opam-coq-archive.

³¹ The `ProofPeer` project is hosted on www.proofpeer.net.

³² The `Mizar` system is available on mizar.uwb.edu.pl.

³³ We can explore the theorems of the MML library on mmlquery.mizar.org.

³⁴ fm.mizar.org

³⁵ The results of the OCaml bench are available on opam.ocaml.org/builds.

We rewrote this website in Coq to experiment our Coq.io library³⁶.

9 Conclusion

WE HAVE PRESENTED our work on the use of OPAM as a package manager for Coq. We especially worked on the design of a repository architecture and on the implementation of a bench system and a website for the OPAM packages for Coq. As of today, OPAM is usable for Coq and we hope that in the future more and more people will use OPAM to share their Coq developments and reuse other's contributions.

³⁶The Coq.io library is available under MIT license on coq.io.