



HAL
open science

jsCoq: Towards Hybrid Theorem Proving Interfaces

Emilio Jesús Gallego Arias, Benoît Pin, Pierre Jouvelot

► **To cite this version:**

Emilio Jesús Gallego Arias, Benoît Pin, Pierre Jouvelot. jsCoq: Towards Hybrid Theorem Proving Interfaces. 12th International Workshop on User Interfaces for Theorem Provers (UITP '16), Jul 2016, Coimbra, Portugal. pp.15-27, 10.4204/EPTCS.239.2 . hal-01425752v1

HAL Id: hal-01425752

<https://minesparis-psl.hal.science/hal-01425752v1>

Submitted on 3 Jan 2017 (v1), last revised 7 Feb 2017 (v2)

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.

jsCoq: Towards Hybrid Theorem Proving Interfaces

Emilio Jesús Gallego Arias

MINES ParisTech
PSL Research University, France
e@x80.org

Benoît Pin

MINES ParisTech
PSL Research University, France
benoit.pin@mines-paristech.fr

Pierre Jouvelot

MINES ParisTech
PSL Research University, France
pierre.jouvelot@mines-paristech.fr

We describe jsCoq, a new platform and user environment for the Coq interactive proof assistant. The jsCoq system targets the HTML5–ECMAScript 2015 specification, and it is typically run inside a standards-compliant browser, without the need of external servers or services.

Targeting educational use, jsCoq allows the user to start interaction with proof scripts right away, thanks to its self-contained nature. Indeed, a full Coq environment is packed along the proof scripts, easing distribution and installation. Starting to use jsCoq is as easy as clicking on a link. The current release ships more than 10 popular Coq libraries, and supports popular books such as *Software Foundations* or *Certified Programming with Dependent Types*.

The new target platform has opened up new interaction and display possibilities. It has also fostered the development of some new Coq-related technology. In particular, we have implemented a new serialization-based protocol for interaction with the proof assistant, as well as a new package format for library distribution.

1 Introduction

Interactive Theorem Proving (ITP) relies on mutual human-machine feedback to build proofs by refinement. Typically, the user first requests the proof assistant to validate or guess some proof step. Depending on the output of the tool, she will continue the proof or correct the last step.

Teaching ITP is usually practice-led; students are encouraged to interact with the tools from the start, simultaneously discovering their way into the particularities of the implementations and the logical theories behind them. Popular ITP teaching material, such as [24, 7], is written in a *literate programming style*, pioneered by Donald Knuth. In literate programming, source code and comments meld to form a coherent, self-documenting book and program. An additional design objective in the context of proof assistants is that the user *must be able to interact with the book*. The same principles apply to documentation, which may be hard to understand at first without the ability to run and play with examples.

Narrowing down our attention to the Coq proof assistant [31], most teaching material is consumed by students with the help of specific Integrated Development Environments (IDE). Popular options are CoqIde or ProofGeneral [1, 25], but alternatives exist. IDEs usually contain or are based on text-editing programs to provide specialized support for proof development. This approach works well in practice; however, IDEs are mostly focused on proof development and provide varying degrees of support for document-like features. For instance, not all IDEs support images, advanced formatting, hyperlinks, or dynamic content creation. This hinders readability and it is not uncommon to see students with the same document doubly opened in a document reader and the IDE, as it is difficult to make sense of the

document without the possibility of interaction. Additionally, the specialized nature of these tools makes installation sometimes heavy, especially if third-party add-ons are involved, posing a barrier to the casual learner.

The jsCoq system intends to tackle both of these problems (readability and installation) from a document point of view: instead of equipping an IDE with document-like features, we extend documents to provide IDE-like capabilities. Our approach is to profit from modern browsers and web standards, embedding Coq as an application inside the browser. In this setting, Coq scripts are plain HTML documents, and a full Coq instance is run locally in the browser for every document. A document manager — written in JavaScript — manages the communication between the browser and the Coq instance. This setup provides a working Coq environment in a transparent way for the user. She will just click on a link, and have an interactive Coq document ready without any other special action or installation, even if the script depends on exotic libraries or add-ons. The document and proof assistant are distributed together and maintained in sync; moreover, the standards-based source code of jsCoq should ensure a good amount of forward compatibility.

There are many goals to consider when porting a system of the sheer size of Coq (currently more than 200.000 lines of source code), which we address here.

- Completeness The full Coq system should run in the browser platform, since non-trivial Coq courses require large helper libraries, from real numbers to complex decision procedures.
- Relevance We intend to support from the very start existing teaching material;
- Performance To be practical, jsCoq should be able to handle large real-world developments in a usable way, and it would be disappointing to find in jsCoq a crippled Coq system.
- Usability We rely for the front end on modern web technologies, which carry a certain *coolness* factor that may motivate the students to use and contribute to the tool.
- Maintainability The architecture of jsCoq should be such that keeping up with upstream¹ called upstream changes in Coq is low-effort, allowing it to run over an unpatched, up-to-date Coq version.

In fact, we believe we started the project at the exact moment when achieving all these goals became possible, as it is the combination of very recent improvements in the Coq API, OCaml-to-JavaScript technology, and web standardization that made the project successful.

The jsCoq development is managed in an open-source way. Our project page [11] provides information on source code, user mailing lists, builds, settings, packages, compatibility, and installation instructions. Note that the intrinsic nature of this project will unavoidably make obsolete some of the information contained in this paper, we recommend users to check our project page for up-to-date information.

The structure of the paper is as follows. We start with a brief introduction to our Web-Based Hybrid Document Model in Section 2. Section 3 gives an overview of the main components and architecture of jsCoq. We report on some preliminary practical use of jsCoq for education purposes in Section 4. We discuss related work in Section 5 and address possible future work in Section 6. We conclude in Section 7.

¹Upstream tools and technologies are those on which a given system relies, such as Coq in our case.

2 The Web-Based Hybrid Document Model

We use the term *hybrid document*² to denote documents capable of containing objects whose interpretation is given by an external program. A typical example is a word processor file embedding a spreadsheet or, in our case, an HTML page containing a Coq proof script. The hybrid-document platform of choice for jsCoq is the so-called *modern web platform* specified by the ECMAScript[®] 2015 [9] standard. In it, base documents are defined using the HTML markup language. JavaScript — a Turing-complete, object-oriented, and functional language — is used to manipulate HTML documents by means of the *Document Object Model* API.

The modern web platform enjoys wide adoption and support, with increasing compatibility and standardization making it ever more future-proof. Additionally, the ubiquity of JavaScript code in the web has led to huge improvements in execution performance: browsers can run now applications of size and complexity thought unmanageable some years ago. The process known as *transpiling*³ thus allows developers to port native applications to JavaScript without a full rewrite. The number of libraries available is staggering; a document can easily gain advanced 3D mathematical plotting [37], editing capabilities [17], or even \LaTeX support [6], with just a few lines of code.

There are many interesting examples of web-based hybrid documents. Jupyter⁴ allows to create and share documents that contain live code, equations, visualizations, and explanatory text; Eloquent JavaScript [16] is an interactive book where all the examples can be run and edited online.

Web-Based Coq Documents For jsCoq documents, we have chosen a simple solution: the user provides an arbitrary HTML document, tagging Coq code as appropriate. Generally, this means that non-editable code will be wrapped in `code` tags, whereas editable content will be placed inside `textarea` elements. At the end of the document, the user asks for jsCoq to be loaded, providing an ordered list of identifiers of the elements that jsCoq will understand as the Coq document:

```

1 <script src="js/jscoq-loader.js" type="text/javascript"></script>
2 <script type="text/javascript">
3   loadJsCoq( './' ).then( () => new CoqManager (list_of_ids , [options]));
4 </script>

```

The `CoqManager` will scan the corresponding elements, initializing the pertinent editor components, etc. Finally, a CoqIDE-style control panel will be attached to the document, containing the goals window, logs, toolbars, etc. The results can be seen in Figure 1. For a live demonstration, jsCoq landing page provides a simple example; the best way to get a feeling of the current system is to try it out. We list here more examples that we developed:

- DFT, a small development of the theory of the Fourier Transform, following [28];
- Mtac, the Mtac [41] tutorial;
- STLC, the "Simply Typed Lambda Calculus" chapter from [24];
- StackMachine, the first chapter of [7].

²or *rich document*, *interactive document*, or any of the many names used in the literature for this concept

³the process of compiling a language to another, i.e., in this case, to JavaScript

⁴previously known as IPython

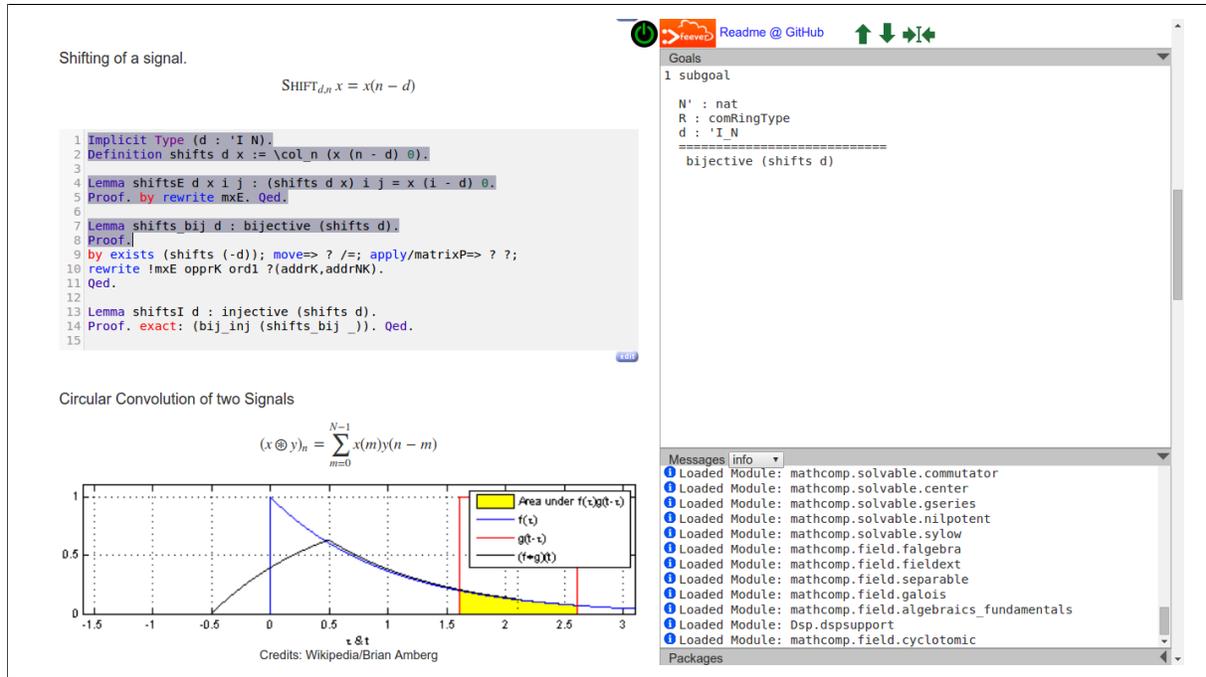


Figure 1: jsCoq 0.6 running a DFT formalization

Web Runtime Platforms We briefly comment on some aspects of our platform of choice. Particular issues of interest are compatibility and future-proofness. Regarding compatibility, we try to be standards-compliant, and jsCoq is known to work with different versions of Microsoft Edge, Mozilla Firefox, and Safari. However, we recommend using Google Chrome, as it is the browser we internally use for testing. While maybe an unpopular decision, we lack the resources to perform proper multi-browser testing; however we would be very happy to include fixes that help support in other browsers, but so far those have not been necessary. We should also note that we have run jsCoq in the Node JS [8] platform successfully. We also believe that our approach should be acceptably resistant to future changes in browsers and platforms. We have taken care to restrict our development to standard constructions that should be respected by the upstream developers; however note that due to the technology used, jsCoq usually requires a pretty recent browser.

3 System Overview

The high-level architecture of jsCoq is the typical one for IDEs: the Coq proof assistant runs in its own, separate process, whereas the user interface runs in the browser thread. The two communicate using message passing following the Web Worker Specification standard [18]. Additionally, a library manager takes care of downloading Coq packages and registering them with the browser virtual filesystem.

The Coq worker Coq offers an XML-based communication protocol — used by CoqIDE and others — which allows to incrementally build proof scripts. However, the protocol relies on Unix features that are not available in a browser environment, and does not provide direct access to Coq’s internal document API which we find valuable. We have thus chosen to implement a simple RPC protocol on top of the OCaml API of Coq. This choice has saved a considerable amount of work

on interfacing with XML objects and provided us with somewhat more flexibility and freedom, without interfering with the established XML-protocol.

The worker relies on a new domain-specific language (DSL) for communication with Coq. Together, the language interpreter and definition are about 200 lines of code. Automatic serialization from/to JSON is performed by the PPX mechanism of OCaml. The worker is then linked with Coq and compiled by `js_of_ocaml` [33] to a monolithic JavaScript file. When instantiated, the worker will constitute a full Coq instance, waiting for control messages.

The Coq package manager The Coq library system is designed under the assumption of an underlying Unix filesystem. However, synchronous file I/O is not available in the browser; thus `jsCoq` needs to manage available libraries and their exact location. For historical reasons, the manager is currently written in OCaml and linked with the worker. However, there is no reason it could not be rewritten in JavaScript and moved out of the Coq worker.

The user interface Written in ECMAScript, the main responsibility of this part is to maintain the user-side document model, taking care of navigation, and synchronizing it with the Coq worker.

Before describing these parts in more detail below, we would like to take a quick look at the (admittedly short) evolution of `jsCoq`. The first versions of `jsCoq` did not use a web worker thread; instead, Coq ran in the main thread. This was a conscious design choice taken to keep things simple and speed up the development. In this implementation, Coq exported a synchronous `jscoq` object with method calls. This worked pretty well and did not require any kind of protocol, but, in the long term, the worker solution is desirable for UI latency purposes. Experience with that first design led to the development of the *SerAPI* protocol (see Section 6), which forms the basis for the current protocol, presented below.

3.1 Coq Protocol

The `jsCoq` protocol is still in evolution as detailed in Section 6. We thus present a snapshot that should provide the reader with the general ideas behind it. We recommend consulting the actual source code to see all the low-level details. Also, for space reasons, we omit some non-interesting calls used to gather version information, etc.

The protocol is defined by the following DSL:

```

1  type jscoq_cmd =
2  | Init      of string list list * string list list
3  | Add      of stateid * int * string
4  | Cancel   of stateid
5  | Observe  of stateid
6  | Goals    of stateid
7  | SetOpt   of bool option * string list * gvalue
8  | GetOpt   of string list
9
10 type jscoq_answer =
11 | Added      of stateid
12 | Cancelled  of stateid list
13 | GoalInfo   of stateid * richpp
14 | Feedback   of feedback
15 | CoqOpt     of gvalue
16 | CoqExn    of loc option * (stateid * stateid) option * richpp
17 | JsonExn   of string

```

Our message-based API is based on the State Transaction Machine (STM) API for Coq [2]. The STM API was designed for asynchronous proof processing and allows to incrementally build and edit Coq

documents. Such design has been proven invaluable in our project. We have introduced minor differences for UI-related convenience; in particular, we have a simpler error handling and we have replaced the “document editing” operation by a “document cancellation” one. This makes the mapping between the Coq document model and the typical editor document model simpler.

For commands, `Init(loadpath, init_mods)` will initialize Coq and set up the proper loadpath and initial libraries; `Add(sid,eid,cmd)` will add `cmd` to the Coq internal document on top of state `sid`, with `eid` used to report parsing errors. `Cancel(sid)` will cancel a previously defined state and their dependent states; `Observe(sid)` will commit to the given state. Finally, `Goals(sid)` asks for the Coq goals to be proved in state `sid`.

The first answers mostly correspond to the first commands; for instance, `GoalInfo(sid, s)` provides a pretty-printed string `s` for the goal in state `sid`. `Feedback` notifies the UI of changes in the Coq state machine, such as states becoming ready. Fatal errors are signaled by `CoqExn(loc,(sid, fid),msg)`, with `sid` the first good state identifier and `fid` the state identifier that produced the exception.

A typical editing session will thus consist in several `Add` commands as the user writes a document followed by `Observe` when execution of a state is requested. Coq will in turn asynchronously acknowledge parsing and execution begin and end points, information that is used by the front end to color the sentence appropriately. When the user edits an already submitted sentence, a `Cancel` event will be sent to Coq, which will in turn respond with a list of sentences to cancel. We thus free the UI of having to take care of sentence dependencies. `GoalInfo` will be performed by the UI whenever it wants to display goals. For instance, when adding many sentences at once, jsCoq will only request goals for the last successful one. “Document backtracking” can be understood in two different ways in this protocol. If the users want to return to a previous document state, observing that document state will move the current state to that area. If, additionally, the user wants to change the document, a cancel must be issued.

In addition, to allow the UI to request information about and load available Coq packages, informing the UI back about progress/completion, we designed a dedicated library manager control language. It is defined as follows:

```

1 type lib_cmd =
2   | GetInfo
3   | InfoPkg of string * string list
4   | LoadPkg of string * string
5
6 type lib_event =
7   | LibInfo      of string * Jslib.coq_bundle
8   | LibProgress of progress_info
9   | LibLoaded   of string

```

3.2 Serialization

The previous command and answer definitions are given as OCaml datatypes, and implemented by a reasonably straightforward OCaml interpreter. However, the browser side, and, in particular, the Web Worker API, requires messages to be JSON objects. What is the best way to relate our OCaml datatypes to their JSON representation? We think that the direct manipulation of JSON objects in OCaml would have consumed most of our available development time. Thus, core to our design is the use of the new PPX system for OCaml meta-programming, which supports the automatic generation of serializers.

In particular, we use the `ppx_yojson` [36] package, which will automatically generate serialization functions for the definitions of our small DSL and — more importantly — for core *Coq data types*. Indeed, this has proven to be a great choice, resulting in a very low overhead when tracking upstream

changes, exporting complex structures, or just experimenting. A Coq datatype is declared serializable using pragmas, such as in the following declaration:

```
1 type feedback =
2   [%import: Feedback.feedback]
3   [@@deriving yojson]
```

The above code will generate the following functions:

```
1 val feedback_to_yojson : feedback -> json
2 val feedback_of_yojson : json -> (feedback, string) Result.result
```

implementing the bridge between OCaml values and their JSON representation.

We can use this method to export any data type exposed in the OCaml plugin API of Coq. As the jsCoq system matures, support for querying more Coq objects is gradually introduced. For example, using this mechanism, we can reliably recognize and handle user-specified notations, printing them in an enhanced way if desired; this is hard to do without serialization support due to the highly flexible Coq pretty-printing engine.

3.3 Document Manager

The document manager relates the HTML document model shown to the user to the internal document maintained by the Coq proof engine. Thanks to careful API design, the UI state is minimized and the relation can be mostly implemented in a stateless, reactive way. The manager currently has three distinct components:

- a `CoqPanel` object, providing the user interface for the goal and query buffers;
- a `CoqProvider` abstract object that encapsulates the management of Coq statements and which, in particular, takes care of selecting the next statement, highlighting, change notifications, etc.;
- a `CoqManager` object that queries the providers and coordinates them with the panel and Coq itself, propagating errors and logs, and keeping track of the proper Coq state.

A key feature of our approach is the use of `CodeMirror` [17] as an instance for the `CoqProvider` component. `CodeMirror` is an open-source software described as “a versatile text editor implemented in JavaScript for the browser. It is specialized for editing code, and comes with a number of language modes and add-ons that implement more advanced editing functionality.” In practice, our Coq `CodeMirror` mode is able to parse and recognize Coq statements, and will notify the manager when a particular part is invalidated by the user.

3.4 Package Manager

The package manager takes care of loading the needed packages and Coq `.vo` files into the Coq instance. Packages are described as JSON files, stating their dependencies and the set of Coq logical paths the package provides. Unqualified modules are considered as deprecated by jsCoq.

The basic unit of the jsCoq package format is the logical path `pkg_id`, which consists of the module identifier, a list of strings. We assume logical and physical paths equal, thus logical path `A.B.C` will correspond to `A/B/C/`. For each logical path, the package format requires a list of `.vo` files (the `vo_files` field) and `.cma` files. For example:

```

1 {
2   "pkg_id": [ "Coq", "extraction" ],
3   "vo_files": [
4     "ExtrHaskellNatInt.vo", "ExtrOcamlString.vo", "ExtrHaskellBasic.vo",
5     "ExtrOcamlIntConv.vo", "ExtrHaskellZNum.vo", "ExtrOcamlNatBigInt.vo",
6     "ExtrOcamlBigIntConv.vo", "ExtrHaskellZInteger.vo",
7     "ExtrOcamlNatInt.vo", "ExtrHaskellNatNum.vo", "ExtrHaskellString.vo",
8     "ExtrOcamlZInt.vo", "ExtrHaskellZInt.vo", "ExtrOcamlBasic.vo",
9     "ExtrOcamlZBigInt.vo", "ExtrHaskellNatInteger.vo"
10  ],
11  "cma_files": [ "extraction_plugin.cma" ]
12 }

```

The second unit of the jsCoq package manager is the *bundle*. A bundle is a set of loadpaths, together with a list of dependencies:

```

1 {
2   "desc": "math-comp",
3   "deps": [],
4   "pkgs": [...]
5 }

```

Then, the package manager will proceed to load bundles in the background and notify the IDE when they are ready. We currently support 16 popular Coq packages, including the full Mathematical Components library [13].

Adding a new package to jsCoq is reasonably straightforward; the current process basically consists in adding a few lines to a makefile pointing where the sources to the package are. Then, our build system will compile and generate the corresponding JSON file automatically. Unfortunately, packages must be compiled with the same exact Coq and OCaml versions used to build jsCoq; thus remote distribution of packages seems hard for the moment.

3.5 Document Generation

Manually writing HTML documents with Coq proofs can be tedious in some cases. To alleviate this task, we provide `uDoc`, a fork of the CoqDoc tool that generates jsCoq documents. For the time being, the tool is aimed at achieving maximal compatibility with existing developments using CoqDoc, and the mapping from Coq files to jsCoq documents is reasonably straightforward. The `uDoc` tool has been used to generate the jsCoq versions of Software Foundations [24], Certified Programming with Dependent Types [7], the MTAC tutorial [41], etc.

4 Practical Validation

So far our jsCoq system has behaved in a stable way, filling a niche for casual Coq users, tutorials, and people willing to try Coq add-ons in an easy way. Indeed, even at this early stage, jsCoq has already been used to support some Coq courses and a tutorial, including:

- “Winter School Advanced Software Verification and Computer Proof”, Sophia Antipolis, January 18-22, 2016;
- “Mathematical Components, an Introduction” and ITP tutorial, Nancy, August 27, 2016; and
- the “Mathematical Components” book.

These two courses total around one hundred jsCoq users. Feedback was quite positive, as the tool worked well for everybody and allowed them to experience Coq; the main complaints came from the immaturity of the user interface, in particular the fact that our choice of UI panels does not adapt well to different screen sizes.

Regarding the jsCoq instance hosted in our servers, it would be safe to say that more than a thousand unique users have tried to access it. However, note that we do not track or gather any personally identifiable information, and so this estimate has been done on traffic data only and may not be very exact. There is also interest in using jsCoq to teach established Coq classes; work is underway and we believe this could happen soon, as the tool approaches its 1.0 release. Additionally, the tool has been proved valuable in question and answers sites such as <http://stackoverflow.com>, where it allowed users to share runnable Coq code snippets.

Performance-wise, our experimental findings are consistent with the predictions by the `js_of_ocaml` developers: jsCoq usually runs on par or faster than the bytecode version of Coq. Memory use is acceptable, with a jsCoq instance having loaded the Coq prelude plus the Mathematical Components library topping at 300MiB in Chrome 54.0. We have also found that the stability of a particular Coq release is quite consistent; thus, releases that tend to run well usually continue to do so in future browser versions. Note however that newer Coq versions may introduce performance problems by themselves. This is usually difficult to predict, as, of today, Coq upstream changes are not tested in our backend before integration occurs.

5 Related Work

There exist a large amount of work in the domain of hybrid document systems, starting with the WEB system [20], mathematical software [21, 38], and web-based solutions [23, 16, 27, 37].

In the realm of theorem proving, we can highlight Proof General [1] as a popular tool for Emacs users, providing a very comprehensive feature set and significantly easing script editing for several interactive proof assistants. Along with CoqIDE [31], this is the standard choice for Coq users. Proof General communicates with Coq using a non-structured, text-based protocol, but it has been recently updated to support a more modern XML protocol [29]. Advanced coding features such as completion and improved display are provided by the Company Coq [25] Proof General add-on. The use cases provided by Company Coq have been very useful in shaping the direction of our own system.

The idea of proof-by-pointing [4, 5] — which allows the user to develop proofs by interacting with logical connectives — also provided guidance and motivation for the development of our tool.

We are obviously indebted to the work of Barras et al [2], which introduced most of the Coq-level technology allowing the development of jsCoq, work which is itself inspired by developments in the Isabelle/jEdit editing technology [35, 34]. Other interesting IDE efforts for the Coq system are [10, 34, 3].

To the best of our knowledge, jsCoq is one of the first systems to embed a full theorem prover inside a browser. A prominent web-based system that depends on a server is ProofWeb [19], which provides a web interface to a Coq server and many other theorem provers. The Logitext [39] approach is related to proof-by-pointing, but web-based and specialized for sequent calculus proofs.

PeaCoq [26] is a web-based front end for Coq. The jsCoq editor component was derived from its editor implementation. Currently, PeaCoq provides much richer proof-building capabilities than jsCoq, but it relies on a central server. This situation has recently changed, as the PeaCoq development version is based on SerAPI [12]. Thus, we hope for the two tools to share the same back end and protocol in the

near future. Given that SerAPI is capable of running as a Web Worker, this would also free PeaCoq from its central server dependency, if so desired.

6 Future Work

In its current iteration, we consider the jsCoq prototype to be a usable beta; however, quite a bit of work still remains to be done. The pending tasks can be categorized into three main areas:

- user interface and web platform;
- document generation;
- protocol and technical foundations.

We provide below more details on each specific point.

User interface and web platform Currently, we are taking little advantage of the possibilities of the web platform. Apart from a few experiments, the text-based roots and workflow of Coq are still much present in jsCoq, as it builds upon the existing facilities. Interesting features such as support for mouse interaction or active document components may be possible to support without big modifications to Coq. However, some other features will require upstream changes in order to achieve a robust implementation.

A key challenge is to improve printing; in particular, an often-requested feature is nice support for mathematical notations. So far, only some experiments have been completed, but we are optimistic that real progress will happen here.

Our user interface proper would also benefit from a rework to allow for more flexibility and functionality. Some interesting ideas are to provide a more data-centric view of the goal buffer, with modifications and goal history. A key challenge is to choose an interesting “web framework” upon which to base our work. So far, the *phosporjs* framework, developed for the Jupyter scientific programming community, looks very promising. Also, our CodeMirror mode needs some additional work.

Regarding our platform support, we are missing two important features: support for Coq’s virtual machine (VM) and support for reliable timeout/interruptions. The VM of Coq is written in C, and is thus outside of the scope of `js_of_ocaml`; however, we believe that supporting it would be possible by using the emscripten [40] transpiler. This task is however low-priority as our users will rarely choose the web platform for high performance computing. The second point is more delicate, as both timeout and interruptions are implemented in Coq by using Unix signals. Indeed, both Microsoft Windows and the browser environment lack proper signal support, and, while some workarounds exist, they are still far from complete. In the near future, we plan to use such workarounds; however, in some cases, killing the whole jsCoq process is required if the user wants to interrupt a long-running computation.

Document generation Given our intended audience, having good content-generation tools is very important. In general, all the challenges of book and web page generation do apply to our current interactive document creation workflow. There are already plans for extending the current `udoc` tool or incorporating jsCoq support to other document generators, and we would like to coordinate with the rest of the Coq ecosystem to achieve smooth integration. In particular, it would be very useful to have good support for exercise-style boxes, as well as allowing for extended interaction of the student with particular Coq features (think, for instance, “click to show the type”, or “hover to show definition” on selected terms). Longer term, we would like to provide a better document-generating experience than the one of the `CoqDoc/udoc` tool.

Also, support for interactive elements — think of a “Coq box”, i.e., an HTML `div` that is filled with the result of the execution of some Coq program — is a must. We are also very close to providing Proviola-like [30] functionality.

An interesting, but insufficiently developed, initiative is our “CollaCoq”⁵ service. At CollaCoq, users can paste arbitrary Coq programs, obtaining a link that they can share, so others can run and experiment with the script. While functional, the service would benefit from additional work.

Protocol and technical foundations As we have previously discussed, work on jsCoq has spurred the creation of a new communication protocol with Coq. Dubbed *SerAPI* [12], for “Seralization-based API”, this protocol provides a more data-centric view of the Coq system and will be at the core of future jsCoq versions. A full discussion of SerAPI is outside the scope of this paper, but we believe that it can get Coq closer to interesting scientific computing initiatives such as the OpenDreamKit project⁶, which intends to provide an open toolkit for the development of mathematics-focused Virtual Research Environments.

7 Conclusion

We have introduced and detailed jsCoq, a new web-based execution platform and user environment for the Coq interactive proof assistant. jsCoq is just a set of static HTML/JavaScript files; thus it can be easily hosted by anyone, and used locally or offline.

Even though jsCoq is still a work in progress, we believe that the possibilities are exciting, and we think it is not too presumptuous to state that this experiment has brought some fresh air to the area of Coq user interfaces. The tool already enjoys some practical impact: it has been used to give some courses, with more courses and tutorials planned in the near future.

Acknowledgments

We would like to thank Clément Pit–Claudel and Enrico Tassi for discussion over this work, the developers of `js_of_ocaml` for their support, and the anonymous reviewers for their insightful comments and careful reading of our paper. This research has been funded by the ANR FEEVER project.

References

- [1] David Aspinall (2000): *Proof General: A Generic Tool for Proof Development*. In Graf & Schwartzbach [14], pp. 38–42. Available at http://dx.doi.org/10.1007/3-540-46419-0_3.
- [2] Bruno Barras, Carst Tankink & Enrico Tassi (2015): *Asynchronous Processing of Coq Documents: From the Kernel up to the User Interface*. In Urban & Zhang [32], pp. 51–66. Available at http://dx.doi.org/10.1007/978-3-319-22102-1_4.
- [3] Jesper Bengtson & Hannes Mehnert (2013): *Kopitiam—a unified IDE for developing formally verified Java programs*. Technical Report, IT University of Copenhagen.
- [4] Yves Bertot (1999): *The Coq System: Design and Architecture*. *Formal Asp. Comput.* 11(3), pp. 225–243. Available at <http://dx.doi.org/10.1007/s001650050049>.

⁵<https://x80.org/collacoq/>

⁶<http://opendreamkit.org>

- [5] Yves Bertot, Gilles Kahn & Laurent Théry (1994): *Proof by Pointing*. In Hagiya & Mitchell [15], pp. 141–160. Available at http://dx.doi.org/10.1007/3-540-57887-0_94.
- [6] David Cervone et al.: *MathJax: A JavaScript display engine for mathematics that works in all browsers*. <https://www.mathjax.org/>.
- [7] Adam Chlipala (2011): *Certified Programming with Dependent Types*. MIT Press. Available at <http://adam.chlipala.net/cpdt/>.
- [8] Ryan Dahl et al.: *NodeJS*. <https://nodejs.org>.
- [9] Ecma International (2015): *ECMAScript 2015 Language Specification*, 6th edition. Ecma International, Geneva. Available at <http://www.ecma-international.org/ecma-262/6.0/ECMA-262.pdf>.
- [10] Alexander Faithfull, Jesper Bengtson, Enrico Tassi & Carst Tankink (2016): *Coqoon: An IDE for interactive proof development in Coq*. In: TACAS, Eindhoven, Netherlands. Available at <https://hal.inria.fr/hal-01242295>.
- [11] Emilio Jesús Gallego Arias: *jsCoq project page*. <https://github.com/ejgallego/jscoq>.
- [12] Emilio Jesús Gallego Arias (2016): *SerAPI: Machine-Friendly, Data-Centric Serialization for COQ*. Available at <https://hal-mines-paristech.archives-ouvertes.fr/hal-01384408>. Working paper or preprint.
- [13] Georges Gonthier, Assia Mahboubi & Enrico Tassi (2008): *A Small Scale Reflection Extension for the Coq system*. Research Report RR-6455. Available at <https://hal.inria.fr/inria-00258384>.
- [14] Susanne Graf & Michael I. Schwartzbach, editors (2000): *Tools and Algorithms for Construction and Analysis of Systems, 6th International Conference, TACAS 2000, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS 2000, Berlin, Germany, March 25 - April 2, 2000, Proceedings. Lecture Notes in Computer Science 1785*, Springer.
- [15] Masami Hagiya & John C. Mitchell, editors (1994): *Theoretical Aspects of Computer Software, International Conference TACS '94, Sendai, Japan, April 19-22, 1994, Proceedings. Lecture Notes in Computer Science 789*, Springer.
- [16] Marijn Haverbeke (2011): *Eloquent Javascript*, 1st edition. No Starch Press. Available at <http://eloquentjavascript.net/>.
- [17] Marijn Haverbeke et al.: *CodeMirror is a versatile text editor implemented in JavaScript for the browser*. <https://codemirror.net/>.
- [18] Ian Hickson et al. (2015): *Web Worker Specification*. <https://www.w3.org/TR/workers>.
- [19] Cezary Kaliszyk (2007): *Web Interfaces for Proof Assistants. Electronic Notes in Theoretical Computer Science 174(2)*, pp. 49 – 61. Available at <http://www.sciencedirect.com/science/article/pii/S1571066107001697>. Proceedings of the 7th Workshop on User Interfaces for Theorem Provers (UITP 2006).
- [20] Donald Ervin Knuth (1984): *Literate programming. The Computer Journal 27(2)*, pp. 97–111.
- [21] Michael B. Monagan, Keith O. Geddes, K. Michael Heal, George Labahn, Stefan M. Vorkoetter, James McCarron & Paul DeMarco (2005): *Maple 10 Programming Guide*. Maplesoft, Waterloo ON, Canada.
- [22] Greg Morrisett & Tarmo Uustalu, editors (2013): *ACM SIGPLAN International Conference on Functional Programming, ICFP'13, Boston, MA, USA - September 25 - 27, 2013*. ACM. Available at <http://dl.acm.org/citation.cfm?id=2500365>.
- [23] Fernando Pérez & Brian E. Granger (2007): *IPython: a System for Interactive Scientific Computing. Computing in Science and Engineering 9(3)*, pp. 21–29. Available at <http://ipython.org>.
- [24] Benjamin C. Pierce, Chris Casinghino, Marco Gaboardi, Michael Greenberg, Cătălin Hrițcu, Vilhelm Sjöberg & Brent Yorgey (2015): *Software Foundations*. Electronic textbook. <http://www.cis.upenn.edu/~bcpierce/sf>.

- [25] Clément Pit-Claudel & Pierre Courtieu (2016): *Company-Coq: Taking Proof General one step closer to a real IDE*. In: *CoqPL'16: The Second International Workshop on Coq for PL*, doi:10.5281/zenodo.44331. Available at <http://hdl.handle.net/1721.1/101149>.
- [26] Valentin Robert: *PeaCoq is a web-based front-end to the Coq proof assistant*. <https://github.com/Ptival/PeaCoq>.
- [27] Benoît Rognier, Guillaume Duhamel, Romain Guillot & Jérémie Maquet: *Edukera*. <http://edukera.com>.
- [28] Julius Orion Smith III (2007): *Mathematics of the Discrete Fourier Transform (DFT): with Audio Applications*, 2nd edition. W3K Publishing. Available at <https://ccrma.stanford.edu/~jos/mdft/>.
- [29] Paul Steckler: *Proof General with XML Protocol Support*. <https://github.com/psteckler/ProofGeneral>. Accessed: 2016-09-13.
- [30] Carst Tankink, Herman Geuvers, James McKinna & Freek Wiedijk (2010): *Intelligent Computer Mathematics: 10th International Conference, AISC 2010, 17th Symposium, Calculemus 2010, and 9th International Conference, MKM 2010, Paris, France, July 5-10, 2010. Proceedings*, chapter Proviola: A Tool for Proof Re-animation, pp. 440–454. Springer Berlin Heidelberg, Berlin, Heidelberg. Available at http://dx.doi.org/10.1007/978-3-642-14128-7_37.
- [31] The Coq development team (2016): *The Coq proof assistant reference manual*. LogiCal Project. Available at <http://coq.inria.fr>. Version 8.5pl1.
- [32] Christian Urban & Xingyuan Zhang, editors (2015): *Interactive Theorem Proving - 6th International Conference, ITP 2015, Nanjing, China, August 24-27, 2015, Proceedings*. Lecture Notes in Computer Science 9236, Springer. Available at <http://dx.doi.org/10.1007/978-3-319-22102-1>.
- [33] Jérôme Vouillon & Vincent Balat (2014): *From bytecode to JavaScript: the Js_of_ocaml compiler*. *Softw., Pract. Exper.* 44(8), pp. 951–972. Available at <http://dx.doi.org/10.1002/spe.2187>.
- [34] Makarius Wenzel (2013): *PIDE as front-end technology for Coq*. CoRR abs/1304.6626. Available at <http://arxiv.org/abs/1304.6626>.
- [35] Makarius Wenzel (2014): *Asynchronous User Interaction and Tool Integration in Isabelle/PIDE*. In Gerwin Klein & Ruben Gamboa, editors: *Interactive Theorem Proving - 5th International Conference, ITP 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 14-17, 2014. Proceedings*, Lecture Notes in Computer Science 8558, Springer, pp. 515–530. Available at http://dx.doi.org/10.1007/978-3-319-08970-6_33.
- [36] whitequark: *ppx_deriving_yojson*. https://github.com/whitequark/ppx_deriving_yojson.
- [37] Steven Wittens: *mathbox: Presentation-quality WebGL math graphing*. <https://gitgud.io/unconed/mathbox>.
- [38] Wolfram Research Inc. (2010): *Mathematica 8.0*. Available at <http://www.wolfram.com>.
- [39] Edward Z. Yang: *Logitext is an educational proof assistant for first-order classical logic using the sequent calculus, in the same tradition as Jape, Pandora, Panda and Yoda*. <http://logitext.mit.edu/main>.
- [40] Alon Zakai et al.: *emscripten*. <http://emscripten.org>.
- [41] Beta Ziliani, Derek Dreyer, Neelakantan R. Krishnaswami, Aleksandar Nanevski & Viktor Vafeiadis (2013): *Mtac: a monad for typed tactic programming in Coq*. In Morrisett & Uustalu [22], pp. 87–100. Available at <http://doi.acm.org/10.1145/2500365.2500579>.