

C-CoRN, the Constructive Coq Repository at Nijmegen

Luís Cruz-Filipe, Herman Geuvers, Freek Wiedijk
lcf|herman|freek@cs.kun.nl

University of Nijmegen

Abstract. We present C-CoRN, the Constructive Coq Repository at Nijmegen. It consists of a library of constructive algebra and analysis, formalized in the theorem prover Coq. In this paper we explain the structure, the contents and the use of the library. Moreover we discuss the motivation and the (possible) applications of such a library.

1 Introduction

A repository of formalized constructive mathematics [4] in the proof assistant Coq [11] has been constructed over the last five years at the University of Nijmegen. This is part of a larger goal to design a computer system in which a mathematician can do mathematics.

At this moment we don't have a complete idea of what such a system should look like and how it should be made; through foundational and experimental research we want to contribute ideas and results that support the development of such a system in the future.

One of the things that is very important for such a 'mathematical assistant' to be used is the availability of a large and usable library of basic results. But such a library shouldn't just be a huge collection of proved results (including the 'proof scripts' that are input for the proof assistant to carry out the proof). In our view, a library of formalized mathematics should be:

Accessible: one should be able to get a fairly fast overview of what's in it and where to find specific results;

Readable: once one has come down to the basic objects like definitions, lemmas and proofs, these should be presented in a reasonable way;

Coherent: results about a specific theory should be grouped together and theories extending others should be defined as such;

Extensible: it should be possible to include contributions from other researchers.

How can one make such a (large) coherent library of formalized mathematics? Ideally, this should also be independent of the Proof Assistant one is working with, but right now we don't know how to do that. Several other projects deal with this question. The Mowgli project [1] aims at devising system independent tools for presenting mathematics on the web. The OpenMath [8] and OMDoc

[23] standards aim at exchanging mathematics across different mathematical applications, which is also one of the aims of the Calculemus project [6]. This may eventually lead to ways of sharing mathematical libraries in a semantically meaningful way that preserves correctness, but that's not possible yet (an exception is NuPRL which can use HOL results [27].)

So, to experiment with creating, presenting and using such a library, one has to stick to one specific theorem prover, and already there many issues come up and possible solutions can be tested. We have chosen to use the Coq proof assistant, because we already were familiar with it and because we were specifically interested in formalizing constructive mathematics.

This paper first describes the backgrounds of C-CoRN: its history and motivation. Then we describe the structure of the repository as it is now and the methodology that we have chosen to develop it. Finally we discuss some applications and future developments.

2 History

The C-CoRN repository grew out of the FTA-project, where a constructive proof of the Fundamental Theorem of Algebra was formalized in Coq. This theorem states that every non-constant polynomial f over the complex numbers has a root, i.e., there is a complex number z such that $f(z) = 0$.

One of the main motivations for starting the FTA-project was to create a library for basic constructive algebra and analysis, to be used by others. Often, a formalization is only used by the person that created it (or is not used further at all!), whereas an important added value of formalizing mathematics – in our view – is to create a joint computer based repository of mathematics. For the FTA-project, this meant that we didn't attempt to prove the theorem as fast as possible, but that in the proving process we tried to formalize the relevant notions at an appropriate level of abstraction, so they could be reused.

An important condition for the successful use of a library of formalized mathematics is to have good documentation of the code. There are two main purposes of documentation:

1. to show to the world what has been formalized via a 'high level' presentation of the work (in our case that would be a \LaTeX document giving a mathematical description of the formalized theory);
2. to help the interested outsider to extend (or change or improve or vary on) the formalized theory.

For (1) one wants to produce a \LaTeX document that 'goes along' with the formalization. This may be generated from the formalization (but it is not quite clear whether it is at all possible to generate something reasonably, and mathematically abstract from the very low level formal proof code). Alternatively – and this is the approach followed in the FTA-project –, this \LaTeX file may be created in advance and then used as a reference for the proof to formalize. The goal of the FTA-project was to formalize an *existing* proof and not to redo the

mathematics or ‘tailor’ the mathematics toward the proof assistant. This meant that the \LaTeX document we started from was an original constructive proof of FTA, described in [20], with lots of details filled in to ease the formalization process. The same approach has been followed throughout the rest of C-CoRN: existing mathematical proofs and theories were formalized, so the (high level) mathematical content corresponds to an existing part of a book or article.

For (2), some simple scripts were created in the FTA-project to be able to extract from the Coq input files a useful documentation for outsiders interested in the technical content. However, this was pretty *ad hoc* and not very satisfactory, and it was changed in C-CoRN, as described in Section 5.

After the FTA-project was finished, i.e., after the theorem had been formally proved in Coq, it was not yet clear that it had been successful in actually creating a usable library, because all people working with the library until then were part of the project. The only way to test this would be by letting outsiders extend the library. This is not too easy: due to the fact that we have tactics implemented in ML (e.g. to do equational reasoning), one cannot use the standard image of Coq and has to build a custom image first. Therefore, the first real test only came when the first author of this paper started as a new Ph.D. student to formalize constructive calculus (leading to the Fundamental Theorem of Calculus) in Coq. The FTA library turned out to be very usable. Most importantly, there was almost no need to restructure the library or redefine notions, implying that most of the basic choices that were made in the FTA-project worked. (Of course, the basic library was extended a lot, with new results and new definitions.) Hereafter, the library was re-baptized to C-CoRN, the Constructive Coq Repository at Nijmegen, since the FTA and the work of the FTA-project had become only a (small) part of the formalized mathematics. Since then, one new Ph.D. student and a foreign master student (visiting Nijmegen) have made contributions to the repository and found the structure (including notations, automation facilities, documentation) quite useful.

3 Why C-CoRN?

Formalizing mathematics can be fun. In the process of formalizing, one discovers the fine structure of the field one is working with and one gains confidence in the correctness of the definitions and the proofs. In addition to this, formalizing mathematics can also be useful. We indicate some of its possible uses:

Correctness guaranteed: the formalized mathematics is checked and therefore the proofs are guaranteed to be correct. This can be vital in the realm of software or system correctness, where one wants to be absolutely sure that the mathematical models and the results proved about them are correct;

Exchange of ‘meaningful’ mathematics: that the mathematics is formalized means that it has a structure and a semantics within the Proof Assistant. So a mathematical formula or proof is not just a string of symbols, but it has a structure that represents the mathematical meaning and its building blocks have a definition (within the Proof Assistant). These can in principle

be exploited to generate meaningful documents or to exchange mathematics with other applications;

Finding mathematical results: based on the semantics and the structure of the formalized mathematics, it should be possible to find results easier. Querying based on the (meaningful) structure is already possible, but more semantical querying would be welcome.

The potential uses of formalized mathematics only really become available if one can share the formalization and let others profit from it, e.g. by making it possible for them to study it, extend it or use it for their own applications or further development. A key requirement for this is that the formalized mathematics be presented. Ordinary (non-computer-formalized) mathematical results are published in articles, books or lecture notes and are in that way shared with other mathematicians or users of mathematics. Giving a good presentation of formalized mathematics in practice, though, turns out to be quite hard. There are various reasons for this:

Idiosyncrasies of the Proof Assistant: when talking to a Proof Assistant, things have to be stated in a specific way, so the system understands it: definitions have to be given in a specific format, proofs have a specific form, etc. Moreover, each Assistant has its own logical foundations (e.g. set theory, type theory or higher order logic), making it easy to express some concepts (e.g. inductive definitions in type theory) and hard to express others (e.g. subsets in type theory). Because of this, mathematical theory will be defined in a specific way that best fits the idiosyncrasies of the system at hand. When presenting the formal mathematics, one would like to abstract from these ‘arbitrary’ choices;

Verbosity of formalized mathematics: to make the Proof Assistant understand (or be able to verify) what the user means, a lot of detail has to be given. By itself, this is unavoidable (and fine, because we really want the mathematics to be verified and that doesn’t come for free). But in the presentation phase, one wants to abstract from these finer details and ‘view’ the mathematics on a higher level. This is not so easy to achieve;

Access to the formalized mathematics: how does one find a result in the library, how does one get an overview of the material? One can query the library with syntactic means, searching a formula of a specific form, as described in [21]. This is helpful, but can a useful result still be found if it occurs in the library in disguise? Also, if a user knows that a specific lemma appears in the library, she will want to apply it, which in a Proof Assistant is done by using its name. But what is the name of (the proof of) a lemma? One probably doesn’t want to clutter up the names in the presentation of the math, so ‘logical’ naming conventions come in handy.

To be able to really understand and work on these points, one needs a ‘testbed’ to experiment with. This was one of the reasons to start C-CoRN: to have a library of mathematics that people can really contribute to, read and work with.

Of course, such libraries already exist. The prominent one is the Mizar Mathematical Library, MML. However, Mizar was not good for experimenting with a library of formalized mathematics, because we wouldn't have control over the library: one can't just install Mizar and formalize a large library on top of the existing one. Soon the system gets slow and one has to submit the work to the Mizar library to have it integrated. Moreover, the Mizar system itself is not open in the sense that one can program one's own tools (e.g. for automation or documentation) on top of it¹. Another important reason not to choose Mizar was that we were aiming at a library of *constructive* mathematics. For all these reasons, Coq was a good choice, with the drawback that there was only a very small initial library to begin with. (There are many 'Coq contributions', but they do not form one coherent library as discussed in Section 8.)

The main reason for working constructively is that we want to use our library, apart from studying how it conducts as a repository, to study the connections between 'conceptual' (abstract) mathematics and 'computational' (concrete) mathematics. The first (conceptual math) deals with e.g. the proof of (and theory development leading to) the Fundamental Theorem of Algebra, while the second (computational math) deals with an actual representation of the reals and complex numbers and the actual root finding algorithm that the FTA exhibits. In this paper we will not elaborate on this any further, but just point out that this was an important motivation for choosing to work with Coq. At present work is being done in program extraction from the C-CoRN library, and this relies heavily on the fact that the library is constructive.

4 Contents

C-CoRN includes at present formalizations of significant pieces of mathematics. In this section, we will give an overview of the main results which are in the library. So far everything in the library has been formalized constructively. Although we do not exclude adding non-constructive theorems to the library, working constructively has some added value, as indicated in the previous section.

The C-CoRN library is organized in a tree-like fashion. This structure agrees with the dependencies among the mathematics being formalized. In Figure 1 the directory structure of C-CoRN can be seen.

At the bottom of the library are the tactics files and the Algebraic Hierarchy, developed in the framework of the FTA-project. Most of the tactics are to make equational reasoning easier, see [19]. In the Algebraic Hierarchy, the most common algebraic structures occurring in mathematics (e.g. monoids, rings, (ordered) fields) are formalized in a cumulative way and their basic properties are proved. For reasons which will be discussed in Section 5, this formalization proceeds in an axiomatic way. Furthermore, this hierarchy is built in such a way

¹ Right now, C-CoRN is not open either: we want to have 'central control' over the library. But the point we want to make here is that the Proof Assistant Coq, on which C-CoRN is based, *is* open.

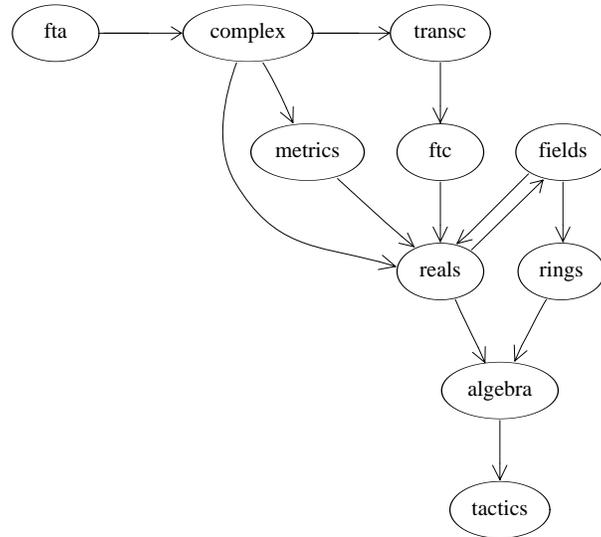


Fig. 1. Directory structure of C-CoRN

that more complex structures are instances of simpler ones, allowing us e.g. to apply lemmas which have been proved for groups to any ordered field.

Real number structures are axiomatically defined as complete Archimedean ordered fields. The C-CoRN library includes not only a concrete model of the real numbers (namely, the standard construction as Cauchy sequences of rationals) but also a formal proof that any two real number structures are equivalent, and some alternative sets of axioms that can be used to define them. Thanks to these results, it makes sense to work with a generic real number structure rather than with any concrete one. This part of the library is described in detail in [16].

Among generic results about real numbers included in the library we point out the usual properties of limits of Cauchy sequences and the Intermediate Value Theorem for polynomials, which allows us in particular to define n^{th} roots of any nonnegative number.

At this point the library branches in various independent directions.

One of the branches consists of the rest of the original FTA-library, which contains the definition of the field of complex numbers and a proof of the Fundamental Theorem of Algebra due to M. Kneser; this work is discussed in [20]. Other important results in this part of the library include the definition of important operations on the complex numbers (conjugation, absolute value and n^{th} roots) and their properties.

The second main branch deals with a development of real analysis following [3]. Here, properties of real valued functions are studied, such as continuity, differentiability and integrability. Several important results are included, among which Rolle's Theorem, Taylor's Theorem and the Fundamental Theorem of Calculus. Also, this segment of the library includes results allowing functions to be

defined via power series; as an application, the exponential and trigonometric functions are defined and their fundamental properties are proved. Logarithms and inverse trigonometric functions provide examples of function definition via indefinite integrals.

A separate branch of the library, currently in its initial stage of development, deals with topological and metric spaces. At present this part of the library is very small; it includes simple properties of metric spaces, as well as a proof that the complex numbers form a metric space.

The sizes of the different parts of the C-CoRN library are shown in Figure 2. This data does not include the files dealing with metric spaces, as these are still

Description	Size (Kb)	% of total
Algebraic Hierarchy (incl. tactics)	533	26.4
Real Numbers (incl. Models)	470	23.3
FTA (incl. Complex Numbers)	175	8.7
Real Analysis (incl. Transc. Fns.)	842	41.6
Total	2020	100

Fig. 2. Contents and size of C-CoRN (input files)

in an early stage of development, nor those dealing with applications to program extraction, which will be discussed in Section 6.

5 Methodology

In order to successfully pursue the goal of formalizing a large piece of mathematics, it is necessary to work in a systematic way. In this section we look at some of the general techniques that are used to make the development of the C-CoRN library more fluent.

We will focus on four main aspects:

Documentation: In order to be usable, a library needs to have a good documentation that allows the user to quickly find out exactly what results have been formalized, as well as understand the basic notations, definitions and tactics.

Structuring: Another important issue is the structure of the library. We feel that lemmas should be somehow grouped according to their mathematical content rather than to any other criterion; e.g. all lemmas about groups should be put together in one place, all lemmas about order relations in another, and so on. A related aspect is how to name lemmas. Experience shows that following some simple rules can make the process of looking for a particular result both easier and faster.

Axiomatic Approach: C-CoRN aims at generality. This suggests that mathematical structures (e.g. real numbers) be formalized in an axiomatic way

rather than by constructing a particular example and working on it. We will examine some of the consequences of this style of working.

Automation: Finally, any successful theorem-proving environment must have at least some automation, otherwise the proving process quickly becomes too complex. We give an overview of the specific tactics that were developed for C-CoRN and show how they help in the development of the library.

Documentation

Providing a good documentation for the formalized library in parallel with its development was a central preoccupation from the beginning of the FTA-project. In fact, having a human-readable description of what has been formalized can be very useful in communicating not only content but also ideas, notations and even some technical aspects of the formalization process.

Such a documentation should at any given moment reflect the state of the library, and as such should be intrinsically linked to the script files. At present, Coq provides a standard tool, called `coqdoc` (see [15]), that automatically generates postscript and `html` documentation from the Coq input files. Additional information can be introduced in the documentation via comments in the script file.

The C-CoRN documentation is presently produced using `coqdoc`. It includes all definitions, axioms and notation as well as the statements of all the lemmas in the library, but no proofs: being meant as *documentation*, rather than *presentation* of the library, we feel that the presence of long and incomprehensible proof scripts in the documentation would undermine its purpose. For the same reason, tactic definitions are omitted from the documentation, but not their description: although the actual code is not presented, the behavior of the existing C-CoRN specific tactics is explained as well as how and when they can be used.

In the `html` version, hyperlinks between each occurrence of a term and its definition allow the users to navigate easily through the documentation, being able to check quickly any notion they are not familiar with.

Structuring

There are several ways that the lemmas and files in a library of formalized mathematics can be organized. The current trend in most major systems seems to be adding individual files to the library as independent entities, and seldom if ever changing them afterward.

However, C-CoRN is intended as a growing system upon which new formalizations can be made. The approach above described directly conflicts with this purpose; for it typically leads to dispersion of related lemmas throughout the library and unnecessary duplication of work.

For this reason, we feel that this is not the best way to proceed. In C-CoRN, lemmas are organized in files according to their statements, and files are distributed in directories according to their subjects. Thus, different areas of mathematics appear in different directories and different subjects within one area will

be different files in the same directory. Of course, this requires central control over the repository: after an extension, the library has to be reconsidered to put the definitions and lemmas in the ‘right’ place. This may become problematic if many files are contributed within a short time.

No part of the library is, strictly speaking, immutable: new lemmas can be added at any time to existing files, if they are felt to belong there. In this way, new lemmas then become immediately available to other users. In practice, though, the lower in the tree structure of Figure 1 a file is, the less often it will be changed.

Coupled with this method of working, the documentation system described above makes looking for a particular statement a simpler process than in most of the systems the authors are acquainted with. But in addition to this, naming conventions are adopted throughout C-CoRN that allow experienced users to find a specific lemma even quicker without needing to consult the documentation. These naming conventions, however, are too specific to be explainable in a short amount of space. The interested reader can find them throughout the C-CoRN documentation.

Axiomatic Approach

One finds two approaches to formalizing algebraic operations. On the one hand one just has types for various number structures, like the natural numbers, the integers, the real numbers and the complex numbers, and for each of those one defines a separate set of arithmetical operations. On the other hand – which is the approach that is followed in C-CoRN, as described in [17] – one can have a hierarchy of the commonly appearing algebraic structures in mathematics, such as groups, rings, fields and ordered fields. and then instantiate these to specific number structures. In this approach the theory of the real numbers will not refer to a specific type of real numbers, but just to a type of ‘real number structure’, which later can be instantiated to a concrete model.

This second approach has advantages and disadvantages. An advantage is that the theory that is developed for a certain structure is maximally reusable. For example the group properties can be reused for the integers, the rational numbers, the real numbers, polynomials, vectors in a vector space, and so on. In our axiomatic approach each of these structures will be just an instance of the already existing algebraic type of groups, and the laws that were proved for it will be immediately available. In the first approach the same theory has to be developed over and over again, every time a new structure with the same algebraic properties is defined.

Another advantage is that the same notation will be used for the same algebraic operation. This is especially useful in a system that has no overloading, like Coq. For instance, in the first approach one has different additions on natural numbers, integers, real numbers, while in C-CoRN all of these are simply written as $(x[+]y)$.

A third advantage is that the development of the theory will more closely follow the development of algebra in a mathematical textbook.

The main disadvantage of the axiomatic approach is that the terms that occur in the formalization are usually much bigger, because they have to refer to the specific structure used. Also, because of the hierarchy of the types of algebraic structures, there will be functions needed in the terms to get to the right kind of algebraic structure. However this is not a problem for the user, since all these operations are implicit: the specific structure is generally an implicit argument, while the functions that map algebraic structures are coercions. On the other hand, internally these terms are still big, so it slows down the processing of the formalization by Coq.

Another slight disadvantage of this approach is that sometimes proofs can be less direct than in the case that all functions are concretely defined. This also affects program extraction. For instance, if one knows that one is dealing with the rational numbers, a proof might be possible that gives a much better extracted program. In the case that one has to give a proof from axioms, this optimization might not be available.

Automation

An important part of the C-CoRN library consists in tools designed to aid in its own development. Together with definitions, notations and lemmas, several automated tactics are defined throughout C-CoRN.

These tactics vary in complexity and in their underlying mechanism. Thus, there are several tactics based on Coq's `Auto` mechanism, which simply performs Prolog-style depth-first search on a given collection of lemmas. Each tactic is designed to work within a specific subject, such as reasoning equationally in different algebraic structures (`Algebra`) or proving continuity of real-valued functions (`Contin`).

Other tactics base themselves on the principle of reflection to tackle wider classes of problems in a more uniform and more efficient way. We mention `Rational`, described in detail in [17], which provides proofs of equalities in rings or fields, but can solve a much larger class of goals than `Algebra`; and `Deriv`, described in [12], a reflective tactic which can prove goals of the form $f' = g$ when f and g are real-valued (partial) functions. Although tactics based on reflection are usually more powerful than those based on `Auto`, they are also more time-consuming when the goals are simple, and usually cannot infer as much information from the context as the latter.

Finally, an interface for equational reasoning is also provided via the `Step` tactic. This tactic allows the user to replace a goal of the form $R(a, b)$, where R is a relation and a and b have appropriate types, by either $R(c, b)$ or $R(a, c)$, where c is a parameter given by the user. This tactic looks through a database of lemmas that state extensionality of (various types of) relations, and chooses the one which applies to R . Then it applies either `Algebra` or `Rational` to prove the equational side condition generated by the lemma.

Presently the `Step` tactic has been generalized to work in a much wider domain than that of C-CoRN, and it is expected to be included in the standard distribution of Coq in a near future.

6 Applications

Besides the direct interest of formalizing mathematics *per se*, there are some interesting applications that are either being explored at present or are planned for the near future.

One of the consequences of working constructively, and therefore without any axioms, is that, according to the Curry-Howard isomorphism, every proof is an algorithm. In particular, any proof term whose type is an existential statement is also an algorithm whose output satisfies the property at hand.

In Coq there is an extraction mechanism available which readily transforms proof terms into executable ML-programs (see [24]). Marking techniques are used to significantly reduce the size of extracted programs, as most of the information in the proofs regards *correctness* rather than *execution* of the algorithm and can thus safely be removed. In [13] it is described how this extraction mechanism was used to obtain, from the formalized proof of the Fundamental Theorem of Algebra, an algorithm that computes roots of non-constant polynomials. At the time of writing the extracted program is too complex and does not produce any output in a reasonable amount of time; but the same method has been used to produce a *correct* program that can compute 150 digits of e in little over one minute.

Of course, the performance of these extracted programs can in no way compete with that of any existing computer algebra system. However, we feel that in situations where correctness is more important than speed, program extraction may very well be successfully used.

7 Future Developments

There are presently a number of different directions in which we would like to see C-CoRN extended in a near future.

One goal is to extend the library by adding new branches of mathematics to the formalization, or by building upon existing ones. In particular, the following areas are considered important:

Complex Analysis: presently there exist a usable algebraic theory of complex numbers and a formalization of one-variable real calculus. These provide a basis upon which a formalization of complex analysis can be built;

Basic Topology: there are no general topology results yet available in C-CoRN; a development of the elementary properties of topological spaces would not only extend the library, but would probably make it possible to unify different parts of the library where instances of the same general lemmas are proved for specific structures;

Metric Spaces: similarly, several of the properties of the absolute value operation on real numbers and its correspondent on complex numbers are in fact instances of properties which can be proved for any distance function on a metric space. We hope that the small development currently existing

in C-CoRN will enable us to prove these and other similar results in a more uniform manner;

Number Theory: on a different line, number theory seems to be a subject where an attempt at formalization could be very successful, since Coq is a system where it is for example very easy to use induction techniques. Furthermore, the preexistence of a library of real analysis would make it much easier to prove results which require manipulating specific integrals;

Group Theory: this is also a subject that would be interesting to explore in C-CoRN. Although we have built an algebraic hierarchy which includes monoids, groups and Abelian groups among its inhabitants, so far most of the development has been done only when at least a ring structure is available. Formalizing important results of group theory would be an important test ground for the usability and power of the algebraic hierarchy.

On a different note, we would like to develop applications of C-CoRN. There are currently plans to do this in two different ways:

Program extraction: the new extraction mechanism of Coq, described in [24], has made it possible to extract and execute programs from the C-CoRN library, as has been explained in [14]. However, the results so far have been slightly disappointing. Recent work has shown that much improvement may be obtainable, and we hope to pursue this topic;

Education: a formalization of basic algebra and analysis should not only be useful for additional formalizations (by researchers) but also for students, who can use it as course material. This addresses a different audience, to which the material has to be explained and motivated (using lots of examples). We believe that a formalization can be useful as a starting point for an interactive set of course notes, because it gives the additional (potential) advantages that all the math is already present in a formal way (with all the structure and semantics that one would want to have) and that one can let students actually work with the proofs (varying upon them, making interactive proof exercises). In the IDA project (Interactive course notes on Algebra, see [9]), a basic course in Algebra has been turned into an interactive course, using applets to present algebraic algorithms. Other experience, especially on the presentation of proofs in interactive course notes, is reported in [5]. We want to investigate C-CoRN as a basis for such a set of course notes.

For usability it is very important to have good automation tactics and powerful (or at least helpful) user interfaces. Along these lines, we have some concrete plans:

Equational reasoning: the present-day `Rational` tactic has as its main drawback that it only works for fields, and cannot use any specific properties which happen to hold in concrete fields such as the real numbers. A variant of `Rational` which works on rings has also been built, but it is not totally satisfactory as it is simply a copy of the original code for `Rational` with some bits left out. We would like to generalize this tactic so that the same program

would work at the different levels of the algebraic hierarchy, recognizing at every stage which properties were available.

Dealing with *inequalities*: it would also be nice to have a similar tactic to reason about inequalities in ordered structures.

8 Related Work

The Coq system is distributed with a basic standard library. There is quite some duplication between what one finds there and what we offer in C-CoRN.

In particular the theory of real numbers by Micaela Mayero [25] is part of the standard library. This duplication extends to the tactics: what there is called the `Field` tactic is the `Rational` tactic in C-CoRN. However, the two theories of real numbers are quite different. Mayero's reals are classical and based on a set of axioms that constructively cannot be satisfied, while the C-CoRN reals are constructive and also have various concrete implementations. Another difference is that in the Mayero reals division is a total function which is always defined (although it is unspecified what happens when one divides by 0), which is not an option in a constructive setting. In C-CoRN, division can only be written when one knows that the denominator is apart from 0. This means that it gets three arguments, of which the third is a proof of this apartness. This difference also shows in the tactics `Field` and `Rational`. The first generates proof obligations about denominators, while the second does not need to do this, because this information already is available in the terms.

Besides the standard library, all significant Coq formalizations are collected in an archive of contributions. From the point of view of the Coq project, C-CoRN is just one of these contributions, although it is currently a considerable part of this archive. The contributions of Coq have hardly any relation to each other. There is no effort to integrate the Coq contributions into a whole, like we tried to do with the C-CoRN library. Everyone uses the standard library, but hardly anyone uses any of the other contributions.

Apart from Coq [11], there are several other systems for formalization of mathematics that have a serious library. The most important of these are: Mizar [26], HOL98 [30] and HOL Light [22], Isabelle/Isar [28], NuPRL/MetaPRL [10] and PVS [29].

The largest library of formalized mathematics in the world is the library of the Mizar system, which is called Mizar Mathematical Library or MML. To give an idea of the size of this library: the source of the C-CoRN repository is about 2 megabytes, the sources of the Coq standard library together with the Coq contributions are about 25 megabytes, while the source of MML is about 55 megabytes. (Of course these sizes are not completely meaningful, as a Coq encoding of a proof probably has a different length from a Mizar encoding of the same proof. Still it is an indication of the relative sizes of the libraries of both systems.)

Some of the theorems that are the highlights of C-CoRN are also proved in MML, like the Fundamental Theorem of Algebra and the Fundamental Theorem of Calculus.

Unlike the Coq contributions, the MML is integrated into a whole: all Mizar articles use all the other articles that are available. So our goals with C-CoRN are similar to that of MML. However, there also are some differences. First, the MML is classical, while almost all of C-CoRN currently is constructive. More importantly, although the MML is revised all the time to make it more coherent mathematically, until recently theorems were never moved. In C-CoRN related theorems are in the same file, but in MML a theorem could potentially be found anywhere. Recently the Encyclopedia of Mathematics in Mizar project (or EMM) has been started to improve this situation, but it is still in an early stage. Another difference is that in C-CoRN arithmetic is formalized in an axiomatic style, as discussed above. In the MML both the non-axiomatic and axiomatic styles are available, but the majority of the formalizations use the former.

Mizar users are encouraged to submit their formalizations to the MML. Mizar is not designed to allow large libraries that are separate from the MML: the performance of the system degrades if one tries to do this. When Mizar users submit their work to MML they sign a form to give up the copyright to their work, so that it can be revised if necessary by the Mizar developers. An approach on how to integrate work by others into C-CoRN still has to be developed. It will need to address the issues that (1) we want to discourage the development of libraries on top of C-CoRN that are not integrated into it, and (2) we want to be free to revise other people's work without getting conflicts over this.

The other systems mentioned above have a library similar to that of Coq. These libraries also have similarities to the C-CoRN library. For instance, in the HOL Light library both the Fundamental Theorem of Algebra and the Fundamental Theorem of Calculus are proved. The Isabelle, NuPRL and PVS libraries also contain proofs of the Fundamental Theorem of Calculus. A comparison between these proofs and the one in C-CoRN can be found in [12].

Of course the C-CoRN repository is the only serious constructive library available today.

Acknowledgments We are grateful to everyone who has contributed to the development of C-CoRN, either by contributing files or by contributing ideas. We especially thank Henk Barendregt, Sébastien Hinderer, Iris Loeb, Milad Niqui, Randy Pollack, Bas Spitters, Dan Synek and Jan Zwanenburg.

Support for the first author was provided by the Portuguese Fundação para a Ciência e Tecnologia, under grant SFRH / BD / 4926 / 2001, and by the FCT and FEDER via CLC.

References

1. A. Asperti and B. Wegner. MOWGLI – A New Approach for the Content Description in Digital Documents. In *Proc. of the 9th Intl. Conference on Electronic Resources and the Social Role of Libraries in the Future*, volume 1, Autonomous Republic of Crimea, Ukraine, 2002.
2. Andrea Asperti, Bruno Buchberger, and James Davenport, editors. *Mathematical Knowledge Management, 2nd International Conference, MKM 2003*, volume 2594 of *LNCS*. Springer, 2003.

3. Errett Bishop. *Foundations of Constructive Analysis*. McGraw-Hill Book Company, 1967.
4. Constructive Coq Repository at Nijmegen. <http://vacuumcleaner.cs.kun.nl/c-corn/>.
5. Paul Cairns and Jeremy Gow. A theoretical analysis of hierarchical proofs. In Asperti et al. [2], pages 175–187.
6. The CALCULEMUS Initiative. <http://www.calculemus.net/>.
7. Paul Callaghan, Zhaohui Luo, James McKinna, and Robert Pollack, editors. *Types for Proofs and Programs, Proceedings of the International Workshop TYPES 2000*, volume 2277 of LNCS. Springer, 2001.
8. O. Caprotti, Carlisle D.P., and Cohen A.M. The OpenMath Standard, version 1.1, 2002. <http://www.openmath.org/cocoon/openmath/standard/>.
9. Arjeh Cohen, Hans Cuypers, and Hans Sterk. *Algebra Interactive!* Springer, 1999.
10. Robert L. Constable et al. *Implementing Mathematics with the Nuprl Development System*. Prentice-Hall, NJ, 1986.
11. The Coq Development Team. *The Coq Proof Assistant Reference Manual, Version 7.2*, January 2002. <http://pauillac.inria.fr/coq/doc/main.html>.
12. Luís Cruz-Filipe. Formalizing real calculus in Coq. Technical report, NASA, Hampton, VA, 2002.
13. Luís Cruz-Filipe. A constructive formalization of the Fundamental Theorem of Calculus. In Geuvers and Wiedijk [18], pages 108–126.
14. Luís Cruz-Filipe and Bas Spitters. Program extraction from large proof developments. In *Theorem Proving in Higher Order Logics, 16th International Conference, TPHOLs 2000*, LNCS, pages 205–220. Springer, 2003.
15. Jean-Christophe Filliâtre. *CoqDoq: a Documentation Tool for Coq, Version 1.05*. The Coq Development Team, September 2003. <http://www.lri.fr/~filliatr/coqdoc/>.
16. Herman Geuvers and Milad Niqui. Constructive reals in Coq: Axioms and categoricity. In Callaghan et al. [7], pages 79–95.
17. Herman Geuvers, Randy Pollack, Freek Wiedijk, and Jan Zwanenburg. The algebraic hierarchy of the FTA Project. *Journal of Symbolic Computation, Special Issue on the Integration of Automated Reasoning and Computer Algebra Systems*, pages 271–286, 2002.
18. Herman Geuvers and Freek Wiedijk, editors. *Types for Proofs and Programs*, volume 2464 of LNCS. Springer-Verlag, 2003.
19. Herman Geuvers, Freek Wiedijk, and Jan Zwanenburg. Equational reasoning via partial reflection. In *Theorem Proving in Higher Order Logics, 13th International Conference, TPHOLs 2000*, volume 1869 of LNCS, pages 162–178. Springer, 2000.
20. Herman Geuvers, Freek Wiedijk, and Jan Zwanenburg. A constructive proof of the Fundamental Theorem of Algebra without using the rationals. In Callaghan et al. [7], pages 96–111.
21. Ferruccio Guidi and Irene Schena. A query language for a metadata framework about mathematical resources. In Asperti et al. [2], pages 105–118.
22. John Harrison. *The HOL Light manual (1.1)*, 2000. <http://www.cl.cam.ac.uk/users/jrh/hol-light/manual-1.1.ps.gz>.
23. M. Kohlhase. OMDoc: Towards an Internet Standard for the Administration, Distribution and Teaching of mathematical Knowledge. In *Proceedings of Artificial Intelligence and Symbolic Computation*, LNAI. Springer-Verlag, 2000.
24. Pierre Letouzey. A new extraction for Coq. In Geuvers and Wiedijk [18], pages 200–219.

25. Micaela Mayero. *Formalisation et automatisation de preuves en analyses réelle et numérique*. PhD thesis, Université Paris VI, December 2001.
26. M. Muzalewski. *An Outline of PC Mizar*. Fond. Philippe le Hodey, Brussels, 1993. <http://www.cs.kun.nl/~freek/mizar/mizarmanual.ps.gz>.
27. P. Naumov, M.-O. Stehr, and J. Meseguer. The HOL/NuPRL Proof Translator: A Practical Approach to Formal Interoperability. In R.J. Boulton and P.B. Jackson, editors, *The 14th International Conference on Theorem Proving in Higher Order Logics*, volume 2152 of *LNCS*, pages 329–345. Springer-Verlag, 2001.
28. T. Nipkow, L.C. Paulson, and M. Wenzel. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer, 2002.
29. N. Shankar, S. Owre, J. M. Rushby, and D. W. J. Stringer-Calvert. *The PVS System Guide*. SRI International, December 2001. <http://pvs.csl.sri.com/>.
30. Konrad Slind. *HOL98 Draft User's Manual*. Cambridge University Computer Laboratory, January 1999. <http://hol.sourceforge.net/>.