COMPUTER THEOREM PROVING IN MATH

CARLOS SIMPSON

Abstract. We give an overview of issues surrounding computer-verified theorem proving in the standard pure-mathematical context.

1. Introduction

When I was taking Wilfried Schmid's class on variations of Hodge structure, he came in one day and said "ok, today we're going to calculate the sign of the curvature of the classifying space in the horizontal directions". This is of course the key point in the whole theory: the negative curvature in these directions leads to all sorts of important things such as the distance decreasing property.

So Wilfried started out the calculation, and when he came to the end of the first double-blackboard, he took the answer at the bottom right and recopied it at the upper left, then stood back and said "lets verify what's written down before we erase it". Verification made (and eventually, sign changed) he erased the board and started in anew. Four or five double blackboards later, we got to the answer. It was negative.

Proof is the fundamental concept underlying mathematical research. In the exploratory mode, it is the main tool by which we percieve the mathematical world as it *really is* rather than as we would *like it to be*. Inventively, proof is used for validation of ideas: one uses them to prove something nontrivial, which is valid only if the proof is correct. Other methods of validation exist, for example showing that an idea leads to computational predictions—but they often generate proof obligations too. Unfortunately, the need to prove things is the factor which slows us down the most too.

It has recently become possible, and also necessary, to imagine a full-fledged machine-verification system for mathematical proof. This might radically change many aspects of mathematical research—for

Key words and phrases. Lambda calculus, Type theory, Theorem proving, Verification, Set theory.

better or for worse is a matter of opinion. At such a juncture it is crucial that standard pure mathematicians participate.

The present paper is intended to supply a synthetic contribution to the subject. Most, or probably all, is not anything new (and I take this opportunity to apologize in advance for any additional references which should be included). What is supposed to be useful is the process of identifying a certain collection of problems, and some suggestions for solutions, with a certain philosophy in mind. The diverse collection of workers in this field provides a diverse collection of philosophical perspectives, driving our perception of problems as well as the choice of solutions. Thus it seems like a valid and useful task to set out how things look from a given philosophical point of view. Our perspective can be summed up by saying that we would like to formalize, as quickly and easily as possible, the largest amount of standard mathematics, with the ultimate goal of getting to a point where it is concievable to formalize current research mathematics in any of the standard fields.

This is a good moment to point out that one should not wish that everybody share the same point of view, or even anything close. On the contrary, it is good to have the widest possible variety of questions and answers, and this is only obtained by starting with the widest possible variety of philosophies.

We can now discuss the difference between what we might call "standard" mathematical practice, and other currents which might be diversely labelled "intuitionist", "constructivist" or "non-standard". In the standard practice, mathematicians feel free to use whatever system of axioms they happen to have learned about, and which they think is currently not known to be contradictory. This can even involve mixing and matching from among several axiom systems, or employing reasoning whose axiomatic basis is not fully clear but which the mathematician feels could be fit into one or another axiomatic system if necessary. This practice must be viewed in light of the role of proof as validation of ideas: for the standard mathematician the ideas in question have little, if anything, to do with logical foundations, and the mathematician seeks proof results for validation—it is clear that any generally accepted framework will be adequate for this task.

A growing number of people are interested in doing mathematics within an intuitionist or constructive framework. They tend to be closest to mathematical logic and computer programming, which may be why many current computer-proving tools are to some degree explicitly designed with these concerns in mind. There are many motivations for this, the principal one being deeply philosophical. Another is just concern about consistency—you never know if somebody might come

up with an inconsistency in any given axiomatic system (see [40] for example). Common sense suggests that if you take a system with less axioms, there is less chance of it, so it is possible to feel "safer" doing mathematics with not too many axioms. A subtler motivation is the question of knowing what is the minimal axiomatic system under which a given result can be proven, see [72] for example. Computer-scientists have some wonderful toys which make it possible directly to transform a constructive proof into a computational algorithm.

The problem of computer-verified-proof in the standard framework is a distinct major goal. Of course, if you prove something in a constructive framework, you have also proven it in the standard framework but the opposite is not the case. Integrating additional axioms such as replacement or choice into the theory, gives rise to a distinct optimization problem: it is often easier to prove a given thing using the additional axioms, which means that the structure of the theory may be different. The notion of computer verification has already made important contributions in many areas outside of "standard" mathematics, and it seems reasonable to think that it should also provide an important tool in the standard world.

One could also note that, when you get down to the nitty-gritty details, most technical results concerning lambda calculus, including semantics, normalization and consistency results for various (highly intuitionistic) axiomatic systems, are themselves proven in a standard mathematical framework relying on Zermelo-Frankel set theory (see [78] for example). Thus even people who are interested in the intuitionist or constructive side of things might find it useful to have good standard tools at their disposal.

One reason why computer theorem verification in the standard pure mathematical context is not recieving enough attention is that most pure mathematicians are unaware of the subject. This lack of awareness—which I shared a few years ago—is truly colossal, given the literally *thousands of papers* concerning the subject which have appeared in recent years. Thus it seems a bit silly, but one of the goals of the present paper is to try to spread the news.

This ongoing work is being done with a number of other researchers at Nice: A. Hirschowitz, M. Maggesi, L. Chicli, L. Pottier. I would like to thank them for many stimulating interactions. Many thanks to the referee who pointed out a number of references, and struggled to improve my faulty understanding of the history of the subject. Several readers of the first version sent helpful remarks and references. I would like to thank C. Raffalli, creator of the PhoX system, who explained some of its ideas during his visit to Nice last summer—these have

been very helpful in understanding how to approach proof documents. Special thanks are extended to the several members of the Coq group at INRIA, Rocqencourt, for their continuing help during visits and on the mailing list. From much further past, I would like to thank M. Larsen for teaching me what little I know about computer programming.

I would like warmly to thank S. Gutt for organizing a great conference, for letting me give a talk which was fairly widely off-subject, and for suggesting a title. And of course the birthday-teenagers for reminding us of the Sumerian numbering system [60].

2. Basic reasons

The basic goal of mechanized theorem verification is to have a language for creating mathematical documents, such that the reasoning can be checked by computer.[1] This is of interest to pure mathematicians for reasons such as the following.

2.1—To reduce the number of mistakes we make in our mathematical work (there are lots of them! [56]).

2.2—Dealing with complicated theories—it is difficult to manage development when there are too many things to prove (and often the author doesn't even know exactly what it is he needs to check).

2.3—Because of (2.1) and (2.2), nobody really has the time or energy adequately to verify all new results, which leads to problems for publication or other timely dissemination. If articles could be validated by machine then referees could concentrate on the question of how interesting they are.

2.4—Black-box theories: it is sometimes useful, but currently dangerous, to use results without knowing how they are proven. An author can have a certain situation in mind, and might use arguments which work in that situation but not in other ones, but sometimes the explicit statements of results don't adequately reflect all the hypotheses. So future users need to understand the arguments before blindly applying the statements. An advantage of machine verification is that the language has to be totally precise: utilization of a "black box" via cut-and-paste rather than axiomatically leaves no room for misinterpretation.

2.5—Integration of many theories or points of view: it is hard for humans because of the steep learning curve for each new theory. Even now the importation of *techniques*, without a full understanding of

[1]We ask that the machine verify the proofs supplied by the mathematician; obviously it can help if the machine finds the proof on its own but that constitutes a difficult, and somewhat distinct, research problem [12].

how they work, is a little-heralded part of mathematical practice, generally glossed over because of the problems of accuracy that are posed. With machine-verified proving these problems should in large part be resolved and should lead to a much more diverse fabric in which different workers develop different techniques and import the techniques of others without needing to spend such a long time learning about the details, leaving everyone with more energy to spend looking for new ideas.

2.6—New forms of reasoning: now, calculations or proofs with little motivation are ruled out by the fact that the mathematician needs to be motivated in order to understand what is going on. Locally unmotivated reasoning might lead to globally interesting results. A striking example in logic is work on minimal length axioms such as [52] [23], where much of the reasoning was found using theorem-proving software.

2.7—The advent of more and more proofs making use of computer calculation (non-existence of a projective plane of order 10 [42] [43], the 4-color theorem, more recently a theorem on hyperbolicity of 3-manifolds [25] and the Kepler conjecture [31]) represents *a priori* a somewhat different use of the computer than what we are considering here. However, these proofs pose acutely the problem of verifying logical details, and a formalized framework promises to present a significant improvement. This issue was addressed by B. Werner in a talk at Nice about the 4-color theorem, and has motivated Hales' "flyspeck project" [31].

2.8—The possibility of computer-assisted proof verification would significantly extend the age period in which mathematicians are optimally productive: older mathematicians tend to avert getting bogged down in details, which often prevents full development of the continuing stream of valid ideas.

On the whole, it is a mistake to think that computer theorem-proving will always lag behind regular math. At some point, the advantages will so much more than counterbalance the difficulties that mathematicians using these tools will jump out ahead.

3. History

3.1—Cantor, Hilbert, Russell; then Gödel and so on, worked on providing a strict logical foundation for mathematics.

Church, Rosser, Martin-Löf [48], Girard [29] worked more specifically on the logical foundations of type theory.

John McCarthy laid out the idea of automated proving in [49] [50] [51].

Robinson introduced the resolution technique [63], variants of which are still central elements of some current systems, and pointed out the importance of interactivity [62].

3.2—Some of the first explicit system projects were Automath, coordinated by de Bruijn, in which notably he pointed out how to deal with variables (via de Bruijn indices); and the first Boyer-Moore theorem provers [9], the Edinburgh Pure Lisp Theorem Prover and ACL, which grew out of work by Bledsoe [10].

In 1977, Jutting entered all of Landau's analysis book in the AUTOMATH system [38].

The Boyer-Moore systems have been several times updated, now becoming Nqthm and ACL2 among others. These have been used for a variety of mathematical work, as well as for verification of computer software. Also the HOL project sprouted from LCF [54] [30] which originated in Edinburgh.

Throughout this period and to the present, theoretical investigations into the design, structure, logic and presentation of theorem-provers have continued; the present history doesn't reflect further on this aspect.

3.3—The Mizar project represents a milestone in its intended scope. It led to the *Journal of Formalized Mathematics* which has been publishing articles verified in the MIZAR system since at least 1989. This includes much basic theory but also research articles. The goal of having a system open to any and all aspects of mathematics, used by a geographically diversified community, supported with archives published in a journal and run by a group dedicated to updating the system while maintaining backward-compatibility, has inspired the more recent generation of theorem-proving systems. In practice, the mathematical orientation of the research articles is sometimes distinguished from the mainstream of what we call standard mathematics (this is an observation but not a criticism).

3.4—A number of systems based on some kind of type theory have been developed: Lego, Isabelle, Nuprl, as well as HOL mentionned above. Of these, Coq is a project which seems relatively successful with respect to the considerations discussed below.

An overview of very recent systems would include: Elf, Plastic, Phox, PVS, IMPS, Ghilbert, Metamath. These and many more can be found

using an internet search engine—a little perseverance turns up an astounding amount of reference material which we couldn't begin to include here (and which is impossible to read in its entirety). A good place to start is with meta-pages such as [97], [98].

3.5—Industrial demand, spurred in part by some well-known problems that might have been avoided with use of the appropriate proof technology continues to be a major impetus for the development of proof engines. In this case, interest is oriented toward proving the correctness of algorithms, integrated circuits or other wiring diagrams, and so forth.

3.6—More and more conferences, journals, and mailing lists bear witness to the fact that fields related to computer proof verification are really booming. Browsing the archives of mailing lists such as [99] or [100] is a fascinating homework assignment.

3.7—In contrast we might mention that the notion of computer calcuation systems, which is perhaps much better known in the world of pure mathematics, is not quite the same thing as computer proof verification. It is surely related, though, and is also experiencing fast growth. *Idem* for the closely related field of *automated deduction* (see [12] for a survey).

3.8—Some examples of proof developments which are of interest to the pure mathematician are: Shankar's 1986 proof of Goedel's incompleteness theorem [67], Gauss' law of quadratic reciprocity (Russinoff [64]), a Paris-Harrington Ramsey theorem (Kunen [41]), constructive category theory of Huet and Saïbi [37], the fundamental theorem of algebra, by Geuvers *et al* [28]; R. O'Connor's formalization of Gödel's theory [58]; the `Reals` and other recent libraries for Coq [81]; L. Chicli's thesis about algebraic geometry [14]; the logical libraries in Isabelle [88]; Metamath's work on Hilbert spaces [87], ...

A glance at this history leads to the conclusion that it now seems like all of the ingredients are available so that we could start doing computer theorem verification in a serious way. We just have to find the right combination. It doesn't seem unreasonable to think that this moment will be seen, in retrospect, as an important juncture in the history of mathematics.

4. Extra-mathematical considerations

It is useful to elucidate some criteria explaining what would make the difference between a system which would or would not be used and accepted by mathematicians. The model of mathematical typesetting and TEXsuggests some criteria.

4.1—FREE: the program shouldn't cost anything (and should be free of any other constraints) for scientific users.
4.2—PLAIN TEXT: the main format of the proof document should be a plain text file; this is to insure that nothing is hidden to the user between the document creation and its compilation or verification, and also to insure easy exchange of documents.
4.3—OPEN SOURCE: the basic programming code used to specify the syntax of the language, and to verify the proof document, should be available to everybody. In the best case it would actually be attached as a trailer to the proof document itself, since the program is an integral part of the proof specification (see the discussion of archiving in the introduction of [25]). Open source for this key part of the system brings the guarantee that your document will be readable or compilable at any time in the future.
4.4—REFERENCE MANUAL: it should be possible to get started using only the reference manual, and at the same time the manual should answer even the most arcane questions (first-time users tend to stumble upon arcane questions more often than one might expect, cf [100]).
4.5—CROSS-PLATFORM: it should be *very easy* to install and use on all platforms. One of the most difficult aspects for a mathematician getting started in this subject is just to get a system up and running usably on whatever computer is available. The average mathematician should be able to install the system and type out a first text file which is then successfully verified, proving a simple mathematical lemma such as uniqueness of the identity element in a monoid, in no more than a couple of hours.
4.6—MODIFIABLE: the system should be modifiable by the user, in particular the user should choose (or invent, or re-invent) his own axiomatic framework, his notational system, his proof tactics in a fully programmable language, and so on—these things shouldn't be hard-wired in. In William Gibson lingo, we want something that would give the "Turing police" nightmares. This information would be like the notion of a macro header to a TeX file. People could exchange and copy their header material, but could also modify it or redo it entirely. This is very important for enabling creative research.
4.7—ENCOURAGE MATHEMATICIANS: the people who maintain the system should encourage standard mathematicians to write math in the system. This includes maintaining an archive, and publicizing the existence of the system within the standard pure mathematical community.
4.8—ATTENTIVENESS TO USER SUGGESTIONS: pure mathematicians are the scientists who work the most with proof, and it seems logical

that their viewpoints would be important to development of a fully functional system. Input from pure mathematicians is likely to be useful even in more far-flung areas of the subject such as industrial applications. This could be seen as compensation for the requirements that the system be free and open-source.

5. Starting tasks

The hardest part of the subject is the first basic task which needs to be addressed: you need to form an opinion about what you are going to want your documents to look like [75].

This task is of course first and foremost treated by the system designers. However, even within the framework of a given system, the mathematician has to address this task too, in determining the style he will use.

It is subject to a number of constraints and hazards.

5.1—Everything you write down has to have an absolutely well-defined and machine-determinable meaning.

5.2—One must write things in a reasonably economical way, in particular we need a reasonable system of notation for the standard mathematical objects which are going to be manipulated.

5.3—It is better if any difficulties which are encountered, have actual mathematical significance. An example of a difficulty which doesn't seem to have much real mathematical significance arises in type theory when we have two types `A` and `B`, with an element `a:A` and a known (or supposed) equality `A=B`. It often (depending on the type system being used) requires additional work in order to be able to say that `a:B`. To the extent possible, the pure mathematician will want to choose a format in which any such difficulty is seamlessly overcome.

5.4—At least a part of what is written should be comprehensible, preferably upon simple inspection of the pure text file.

5.5—potential problem: it is quite possible (and maybe even generally the case) to invent notation which suggests a meaning different from the actual meaning. When you prove a given statement, did you really prove what you thought you did?

Unfortunately there is no way for the computer to verify this part. The conclusion is that we require human comprehension of the definitions, and of the statements of the theorems (the main ones, at least). This imperious requirement is stronger than the natural desire also to understand the proofs.

5.6—Experience suggests that the proofs, in their explicit and verifiable form, actually contain too much information and the reader will

not really want completely to understand them. Indeed the origin of the whole project is the fact that we aren't able to follow the details of our proofs. Rather one might ask that two particular aspects stand out: the general strategy, and the unexpected main steps.

6. How it works

We describe in brief some of the salient points of the `Coq v8.0` system [81] currently used by the author. A number of other recent systems are similar. Aside from the geographical reason for my choice of system, one can point out that `Coq` scores pretty well (but naturally with room for improvement) on the above lists of *desiderata*. [2]

6.1—Some natural considerations lead to the idea of using *type theory*[3] as a basis for mathematical verification [48]. To start with, notice that we want to manipulate various sorts of objects, first and foremost the mathematical statements themselves. Although perhaps not strictly necessary, it is also very important to be able directly to manipulate the mathematical entities appearing in the statements, be they numbers, sets, or other things. To deal with our entities (which will be denoted by expressions entered into the text file), we need function application `(f x)` and function-definition

```
fun x => (... x ... x ....).
```

This combination of operators yields the *lambda calculus*. The notion of *β-reduction* is that if

```
f:= fun x =>(... x ... x ....)
```

then for any expression `y`, the application `(f y)` should reduce to `(... y ... y ....)`. It is natural to ask that the computer take into account this reduction automatically.

6.2—The first famous difficulty is that if `f:= fun x =>x x` then `(f f)= (fun x =>x x) f` $\xrightarrow{\beta}$ `(f f)` $\xrightarrow{\beta}$ `(f f)` .... Thus the β-reduction operation can loop.

The solution is the notion of *type* introduced by Russell: every "term" (i.e. allowable expression, even a hypothetical or variable one) has a "type". The meta-relation that a term `x` has type `X` is written `x:X`. The λ construction (construction of functions) is abstraction over a given type: we have to write

[2] If the makers of other systems feel impelled at this point to say "hey, our system does just as well or better", the main point I am trying to make is: That's great! Try to get us working on it!

[3] Although it should be pointed out that this is not a strictly necessary ingredient, and indeed some systems notably in the Boyer-Moore family are not based on type theory.

```
f:= fun (x:A)=>(... x ... x ....)
```

where `A` is supposed to designate a type. Furthermore the expression on the right is required to be allowable, i.e. well-typed, under the constraint that `x` has type `A`.

6.3—The Church-Rosser theorem states that:
—*In* typed λ*-calculus,* β*-reduction terminates; also calculation of the typing relation terminates.*

The verification program (such as Coq) basically implements this calculation. The fact that it is sure to terminate is a considerable advantage for the end-user: you are sure that the computer will stop and say either *ok* or *no*. And if "no" you were the one who made the mistake.

6.4—**The product operation** Π: in the above expression, what is the type of `f`? We need a new operator Π written as `f:forall (x:A), B x` where

```
B:=fun (x:A) => "the type of (...x ... x ....) when x has type A"
```

(recall that the allowability restriction on `f` requires that this expression `(B x)` exist, and in fact it is calculable).

6.5—**Sorts:** What is the type of a type? We are manipulating the types themselves as terms, so we are required to specify the type of a type. This is known as a "sort" (or in some literature, a "kind"). In current theories there are just a few, for example:
`Prop`, the type of propositions or statements; and
`Type_i`, the type of types which are thought of as sets (with a "universe index" `i`, see below), so an `X:Type_i` corresponds to a collection of objects, and an `x:X` corresponds to an element of the collection.

6.6—**Curry-Howard:** In many (but not all) type-theory based systems, the notion of mathematical proof comes out for free from the type-checking algorithm. This is known as the "Curry-Howard principle": propositions are realized as types, which the reader can think of as being sets with either zero or one element. Proving a proposition amounts to giving an element of the corresponding type, so the proof-checking mechanism is the same as the mechanism implementing the Church-Rosser algorithm refered to above, determining whether a given expression is indeed a term in the propositional type claimed by the mathematician. As suggested by the notation, the product operation Π corresponds to the logical "forall" in propositional calculus (a product of zero-or-one-element sets is nonempty if and only if all of the component sets are nonempty).

The above system corresponds to Martin-Löf's type theory [48]. One can imagine theories with a more complicated diagram of sorts, which is probably an interesting research topic for computer scientists.

6.7—INDUCTIVE OBJECTS: An interesting feature of a number of systems including `Coq` is the notion of inductive objects, which allows us automatically to create recursive datatypes. Since understanding and calculating with recursive datatypes is what the computer does best, it can be good to make use of this feature. It is often the only viable way of getting the computer hitched onto the calculational aspects of a problem. (This might change, notably with the advent of connectivity between theorem-proving systems and calculation systems.) We won't go into any further detail on the basic structure or functionality of inductive definitions; the reader is referred to the reference manual of [81].

7. Universes

In order to get a coherent system (i.e. one without any known proofs of `False : Prop`) the sorts `Type_i` have to be distinguished by universe levels `i`, with `Type_i : Type_(i+1)` whereas `Type_i` $\subset$ `Type_(i+1)` also. In the set-theoretical semantics, these universe levels can be thought of as being Grothendieck universes, see [76]. The incoherency of the declaration `Type:Type` without universe indices was shown in [29]. Nonetheless, in some type-theoretical programs such as `Coq`, a vestige of `Type:Type` is preserved by a neat, but quirky, mechanism known as *typical ambiguity* (see [21]). The user just writes `Type` and the computer makes sure that it is possible to assign indices to all occurences of this word in such a way that the required inequalities hold. If this is not possible, then a `Universe Inconsistency` error is raised. Thus, the user *writes his text file as if* he were using the inconsistent `Type:Type` system, and the verification program then verifies that no unallowable advantage was taken.

Curiously enough, this leads to a phenomenon that looks for all the world to me like quantum mechanics.[4] Suppose you prove `A -> B` (i.e. "`A` implies `B`") in one text file, and you separately prove `B -> C` in another. You might think you have proven `A -> C`. However, when you merge the two files together to get a proof of `A->C`, the verification program might come up with a `Universe Inconsistency` error if the constraints on the universe indices which are generated by the proofs of `A->B` and `B->C`, while separately satisfiable, are not satisfiable by any common assignment.

[4]This might be related to [65].

This phenomenon can be seen as a "disturbance due to measurement". If we think of `A`, `B` and `C` as measuring posts along a lab bench (such as polarized glass panes, say `A` is horizontal, `B` is at 45 degrees and `C` is vertical) then the result of starting with a photon polarized as `A`, doing the proof of `A-> B` and then "measuring" (i.e. stopping the text file and saying "wow, we just proved `B`"), then inputting the photon polarized as `B` into the next proof of `B -> C`, we get a photon polarized as `C` in the output; whereas if we do the whole proof starting from `A` and trying to get out `C`, no proof-photon comes through!

In [71] `qua.v`, a simple example of this phenomenon is given using Russell's paradox, assuming the excluded middle and using `eqT_rect` via Maggesi's proof of `JMeq_eq` [47] [18]. As pointed out by H. Boom [8] it is possible (although more complicated) to do the same thing without axioms other than the basic `Coq` system, using the Burali-Forti paradox [29] [6].

It is important to note, then, that as long as the typical ambiguity mechanism is in place, it is not allowable to prove a theorem in one file and import it as an axiom in another file. It is actually *required* that the proofs of component pieces of the argument be imported by cut-and-paste (2.4). Luckily the system has a `Require` command that does this automatically without having to copy the files.

Questions about "universe polymorphism" lead to some of the most difficult notational problems. The following discussion comes from work of Marco Maggesi, who is also working on solutions to the problem possibly by implementing the thesis of Judicaël Courant [17]. As was pointed out in [21], the basic example is the distinction between small categories and big categories, in other words assigning universe levels in the definition of category. Ideally, polymorphism should be extendable to definitions such as this. In the absence of such a mechanism, we seem to have to rewrite `Category_i` once for each level `i`. This can lead to a nightmare of notation, for example do we need to distinguish between a `comp_sm` for composition in a small category, and `comp_bg` for a big category? If this seems annoying but doable, notice that we then need `Functor_ij` for functors from a `Category_i` to a `Category_j`; and composing them, and so forth!

One of my main motivations for going back to a purely set-theoretic approach (cf 10.3, 10.4 below) is to try to avoid the above problems as much as possible. However, it seems that there will always be a threshold between the objects we are willing to manipulate and the bigger ones we would prefer not to manipulate, and looking at categories of the smaller objects inevitably bumps up against the bigger ones. Until someone comes up with an adequate general solution (perhaps the

ideas in [21] are relevant) we are left with little choice but to live with the situation.

8. Exotic versus regular math

The system described above applies to a wide variety of situations, problems and logics.

8.1—Intuitionism: One of the main utilizations of this type of system is to verify intuitionistic logic. In fact it is difficult to imagine taking a complicated argument and trying to insure, as a human reader, that the proof *didn't* use the excluded middle in some hidden place in the middle. In this sense, computer verification seems crucially important and allows a much more serious approach to the field.

8.2—Impredicative sort: in previous versions of `Coq` the sort `Set` was declared as *impredicative*,[5] in other words if `X:Set` then the product type

```
Set -> X := (Y:Set)X
```

is classified as having type `Set`. Nonetheless elements `X:Set` behave somewhat like sets: strong elimination is allowed, so that constructors in inductive sets are distinct. The intuition behind this situation is that the constructive functions on a constructive set themselves form a constructive set too. Thus these considerations are closely related to the constructive or algorithmic aspects of mathematics. Hugo Herbelin points out that this can lead to some unexpected and fascinating things, such as a proof which extracts to an algorithm where nobody can understand how it works [33].

8.3—**Caution:** the existence of an impredicative but strongly eliminatory sort is somewhat "orthogonal" to standard mathematics, in that it contradicts most choice axioms you can imagine, even the much weaker dependent choice axioms [26] [15]. Perhaps not all, though: David Nowak in [57] proposed doing the axiom of choice for the `Ensembles` library, which means looking at a function associating to any predicate `P:X->Prop` another predicate `(choice P):X->Prop` with the property that if `P` is nonempty then `(choice P)` is a singleton. One conjectures that this version is consistent even with an impredicative `Set`. Nowak's version is difficult to use in practice, and leads to a situation where the only objects which are manipulated are predicates. This nullifies much of the advantage of using a type-theory framework. So, for implementing classical mathematics we prefer working without the impredicativity.

[5] In v8.0 the impredicativity is turned off and can apparently be turned back on by an option.

8.4—Logic in a topos: recall from Moerdijk-MacLane [55] that one way of providing a model for certain intuitionistic logical systems is by looking at logic in a topos other than the topos of sets. One can say this differently (as they do in the book) by saying that any logical reasoning which is expressed solely in a limited intuitionist formulation, applies internally to logic within any topos. One can formulate as a conjecture, that this principle extends to the full intuitionist logical system as implemented in `Coq`:

Conjecture *The $\lambda - \Pi$ calculus with universes and inductive objects as described above and implemented in* `Coq`*, applies to logic within a Grothendieck topos.*

It would be a generalization of the principles set out in [55]. Some part of it may already be known, see Streicher [73] as well as the references therein (and indeed the formulation of the question appears to have been known to workers in Synthetic Domain Theory for some time [22] [70]). It is not immediate, for example, to see how to implement the sorts `Type_i`. A *caveat* might also be in order because recent versions of `Coq` include some inductive principles such as `False_rect` or `eqT_rect` which might not be realizable in the internal logic of a topos. So, the above statement is really just intended to point out that it is an interesting research problem to see how far we can extend the internal-logic principle from [55].

As an example of how this type of thing might actually lead to useful results for standard mathematics, one can use the full flexibility of the `Coq` implementation of type theory to define internal cohomology within the logic (this is done for H^1 in [71] `h1.v`). If this could be applied to a Grothendieck topos, it would show that cohomology in the topos was an internal logical construction. Internal algebraic geometry inside the topos of sheaves on the big fpqc site takes on a particularly nice form very close to the original Italian geometry, for example affine n-space is quite literally just the cartesian power of the ground field k^n.

8.5—The Hodge conjecture: The statement of the Hodge conjecture implies that the problem of saying whether a given topological cycle is algebraic, is decidable. Same for the question of whether a given cycle is Hodge (i.e. vanishing of the Hodge integrals). This might be considered as a logical analogue of Deligne's observation that the Hodge conjecture implies that Hodge cycles are absolute. The pertinance of calculatory methods showed up recently in [69]. It could be interesting to investigate further, for example does Deligne's proof of [20] imply the decidability for abelian varieties? And in exactly which level of intuitionist logic would the decidability hold? What are the

consequences of this decidability given already by the Lefschetz $(1,1)$-theorem? What happens internally in a topos? Could one obtain a counterexample to the Hodge conjecture by showing that in some cases the question is undecideable?

8.6—So there are undoubtedly lots of fun things to look at in these exotic reaches, but they don't really have anything to do with our original goal of implementing, as quickly as possible, a large amount of standard mathematics. More generally, my experience is that any "brilliant gadget" seems predestined to be a bad idea. Boring as it may seem, we have just to plod along in a mundane but systematic way.

9. Typed versus untyped style

One of the central questions in the theory has been whether to impose a strict type system, or whether to relax the typing constraints. There are some very interesting threads of discussion about the issue in the mailing lists, particularly [99] following the postings [44] and [34]. It is interesting to note that in the first thread of discussion, Lamport (forwarded by Rudnicki [44]) offered the viewpoint that the untyped style was better than the typed one. He was promptly refuted by almost all the subsequent messages. A year later when Holmes [34] asked what options people generally used to answer the question, almost everybody replied that the untyped style was obviously the way to go. Unfortunately, this seems to have marked the end of interesting discussion on the QED mailing list [99]. At about the same time the Coq system (where typed language was systematic) started picking up speed: the announcement of Coq Version 5.10 is among the last group of messages in Volume 3 of [99].

The published version of [44] is [45], issued from a combination of Lamport's original note and a rebuttal by L. Paulson. This discussion and the discussion in J. Harrison [32] are excellent references for the multiple facets of the question which we don't attempt further to reproduce here.

The distinction between the typed and the untyped style is generally seen as related to the axiom of choice. In the absence of the axiom of choice (in particular, in any sort of constructive environment) it is much more natural to keep a lot of type information, looking at the values of functions only when they are well defined and make sense. On the other hand, the axiom of choice provides a nice way to implement aspects characteristic of an untyped environment such as function evaluation which is defined everywhere [45].

Some authors make a distinction between choice and dependent choice, the latter stating only that we can choose a function whose values lie in a family of single-element types. By [26] [15] this weaker form poses the same problems vis-a-vis a constructive axiom such as impredicativity of `Set`. Furthermore it generates a lot more proof obligations since one has to prove unicity whenever it is used. For this reason we don't pay too much attention to the distinction, and look directly at the full axiom, even a strong Hilbertian form involving an ϵ-function.

10. Options for doing set theory

Sets are the most basic building blocks entering into the definition of various mathematical objects. On the other hand, as we have seen above, natural considerations about general mathematical theories lead to type theory as the framework for the system as a whole. The most critical decision about document style is therefore the way the notion of a set is encoded in the ambient type theory.

We will review some of the most prominent options. These follow the contours of the "typed versus untyped" discussion.

10.1—Interpret sets as types. Benjamin Werner's paper providing the justification for the $\lambda - \Pi$ calculus with inductive definitions [76] rests on a *semantics*, i.e. assignation of meaning to the formal symbols of the theory, where `X:Type_i` corresponds to a set in the ith Grothendieck universe U_i, and functions such as constructed by the λ operation `fun x:A =>( ... x ... )` are represented by sets which are the graphs of the corresponding set-theoretic function. We can use this semantics as our method for interpreting sets, namely when we want to talk about a set we just take any `X:Type`. In this option we make full use of the notion of dependent types, see C. McBride [5]. The following example shows how we would define the type of groups. It uses the `Record` construction which is a special inductive definition corresponding to an explicit list (of objects possibly depending on the previous ones).

```
Record Group : Type_(i+1) := {
elt:Type_i;
op : elt -> elt -> elt;
id : elt;
inv : elt -> elt;
assoc : forall x y z:elt,(op (op x y) z) = (op x (op y z));
left_id : forall x:elt, (op id x) = x;
right_id : forall x:elt, (op x id) = x;
left_inv : forall x:elt, (op (inv x) x) = id;
right_inv : forall x:elt, (op x (inv x)) = id
```

```
}.
```

With this definition, if `G:Group` then `(elt G):Type_i` is the "set" of elements of `G`.

The advantage of this viewpoint is that it is generally quite easy to understand what is going on. You the uninitiated reader didn't have any trouble reading the previous definition, and could by cut and paste give the analogous definition of, say, `Ring`, right?

One of the drawbacks with this point of view is that if `H:Group` too, and if for some reason we know `G=H`, we have problems interpreting elements of `G` as being the same as elements of `H`. For example it is not legal to write

```
forall x y z:(elt G), (op H (op H x y) z) = (op G x (op G y z))
```

(if rare at the outset, this sort of thing tends to come up in the middle of proofs). The problem is alleviated by the recent inclusion in `Coq` of the "transport" function `eqT_rect`, the utilisation of which is best accompanied by a systematic exploitation of the `JMeq` equality proposed by McBride [5]. This doesn't make the problem go away entirely, though.

This point of view adheres pretty closely to the "typed" side of the discussion in the previous section. An appropriate choice axiom could allow us to extend function definitions to larger types, whenever a default is available. Unfortunately this implies a distinction between empty and nonempty types. Some have made more exotic proposals to include a default element $\perp$ in every type (this is based on work of Scott [66], and was mentionned in the replies to [34]). I haven't looked closely at it, mostly for fear of a large number of proof obligations of the form $x \neq \perp$, and because it obviates the advantage that we can easily understand what is going on. Also because the guy in the office next door talks about it so much.

To close this subsection we also mention another more subtle drawback, which might really be the main problem. There is a very slight shift in universe indexing, compared with what we are used to in set theory. In the above `Record` if we want the underlying set `(elt g)` to be in the universe `Type_i`, the type `Group` is forced to be in the universe `Type_(i+1)`. In general, a set at a given universe level will also have elements at that level. Thus, *a priori*, the elements `g:Group` are to be considered (even though this statement doesn't have any precise technical meaning) as being in the $(i+1)$-st level. This contrasts with set theory where, if `g` is a group whose underlying set `(elt g)` is in the i-th universe, then `g` is represented by an ordered uple which is also in the i-th universe.

10.2—Setoids. From a constructive point of view, the notion of a quotient can be problematic because an equivalence relation need not

be decidable, so when defining an object as a carrier modulo a relation, the carrier might be constructive or computational but the quotient less so. Another problem is that the classical construction of a quotient as a set of equivalence classes can pose problems in certain intuitionist frameworks (imagine doing that internally in a topos, for example).

For these (or perhaps other) reasons a much-favored implementation of sets is as *setoids* (cf the `Setoid` libraries in `Coq`). A setoid is a carrier type together with an equivalence relation. We can think of them as quotient sets without adding additional axioms. One major drawback is that there is a lot of extra stuff to prove, basically that all operations defined on the level of the carrier type are compatible with the equivalence relation.

For standard mathematics, we don't really need the setoid viewpoint since we are willing to add in all of the axiomatic baggage which makes the usual theory of quotients go through. What is useful is to know that such structures exist, are natural, and have been studied, because they occasionally show up uninvited.

10.3—Introduce a type `E` of sets. In this viewpoint, the type-level is considered as the meta-language of mathematical expressions, and we introduce maybe some parameters

```
Parameter E : Type.
Parameter inc : E -> E -> Prop.
```

Here `E` is for "Ensemble" and `inc` is the elementhood relation. One can then go on to suppose all of the usual axioms in set theory, and proceed from there. For example the definition of the subset relation becomes

```
Definition sub := [a,b:E](c:E)(inc c a) -> (inc c b).
```

A development of set theory using this type of viewpoint was done by Guillaume Alexandre in his Coq contribution "An axiomatisation of intuitionistic Zermelo-Fraenkel set theory" (1995) [3]. This version is also quite analogous to the foundations of the MIZAR system [91], and has been developped rather deeply in the "ZF" logic in the Isabelle system [88].

One can even *create* such an `E` using an inductive definition and a quotient, see Aczel [1] and Werner [76]:

```
Inductive Ens : Type := sup : (A:Type)(A->Ens)->Ens.
```

This is a carrier type, bigger than what we actually think of as *Ens* because an equivalence relation of extensional equality is defined (thus, `E` should actually be thought of as a setoid and to get a type we would have to take the quotient). Once we have the basic axioms, which

in these constructions may require supposing their counterpart axioms on the type-theory side [76], the further development of the theory doesn't require knowing about the construction, so in this sense it is conceptually easier just to take `E` as a parameter.

In theories of this option, sets are terms of type `E`, but not themselves types, and elementhood is just a propositional relation on `E`. This is contrasted with the type-theoretical interpretation (10.1) where the relation `x:X` is meta to the theory, not actually accessible from within, and is decided by the computer.

This would seem at first glance to be much less powerful than version (10.1) where sets are types. In a certain sense that's true, but in some other ways it is actually more powerful. The fact that the elementhood relation is a propositional function in the theory allows us to manipulate it mathematically. The problem referred to above about elements of sets which are equal, goes away: if `a=b` and if `x:E` then the propositions `(inc x a)` and `(inc x b)` are equal since they are functions of `a=b`. Inside proofs, the fact that a certain term has a certain type is replaced by the statement that a certain set is an element of another set; in the first case the knowledge is required at the start because otherwise the statement will be rejected by the type-checker, but in the second case the proof can be deferred into a separate branch of the proof tree.

If we further add an appropriately strong version of the axiom of choice (basically Hilbert's ϵ) then we get an even better situation [44] [45] where function evaluation can be defined as the operation which to any `f,x:E` assigns a choice of `y:E` such that the pair `(pair x y)` is contained in `f`, if such a `y` exists (or a default value if no such exists). In fact, if we write down all of the axioms in this general spirit of everywhere-definedness, then everything that follows will be everywhere-defined too, and we are no longer bothered with objects depending on proofs of propositions.

The main disadvantage of the purely set-theoretical approach is that we have lost any possibility for the computer directly to compute on sets, because computation requires specific knowledge about the inductive structure of the type containing a given term and our `E` no longer has any inductive structure. (This might well be alleviated using the direct constructions of [1] and [76], though, a thought-experiment that starts off the next option.)

10.4—Combine 10.1 and 10.3. In the Aczel-Werner type constructions [76] one can construct a set i.e. term `x:E` whenever one has a type `A:Type` together with a realization function for its elements `A -> E`. Imagine that we want to realize as a set an inductive type say

```
Inductive nat : Type :=  O : nat | S : nat -> nat.
```

To do this we would need to construct an injection `nat -> E`. In this particular case it is natural to do so because the construction

$$\emptyset,\ \ \{\emptyset\},\ \ \{\emptyset,\{\emptyset\}\},\ldots$$

played an important historical role in the beginning. However, if we try to do the same for, say, a free monoid, few people ever really worry about the details of how to do it. This requirement would therefore seem to qualify as something that doesn't have any real mathematical significance: we only care about the universal properties of an inductive type, not about how it is actually realized.

In Werner's paper [76], he describes briefly how to realize inductive types as sets, and consequently how to realize every type as a set. So, let's just abstract out this process and suppose, with a parameter, that we know how to do it but we don't care what it is. We get a situation where, combining (10.1) and (10.3), sets are types, but also the elements of sets are realized as sets themselves—which after all is the bedrock assumption of classical set theory.

This is the viewpoint which is developped in [71]. We start by putting

```
Definition E:= Type.
Parameter R : (x:E)x->E.
Definition inc := [a,x:E](exists [b:x](R b) = a).
```

Thus, we identify the notion of set with the notion of type. The parameter `R` is the realization function which says that the elements of a set/type are themselves sets. The elementhood relation `inc` is defined using `R` by saying that $a \in x$ iff there exists a term b of type x whose realization is a.

We assume a certain number of axioms such as choice (in the form of a Hilber ϵ-function) and replacement. We can then go on to develop set theory for `E` just as would be done in option (10.3) above. On the other hand, we can calculate with inductive types and other type-theory features, useful among other things to establish a notational system.

11. Inheritance

Consider the notion of Lie group. This fits into at least 4 distinct (plus one overlapping) inheritance threads as shown in Figure 1. The three rightmost threads are formally similar, however the ringed spaces in question are different, and the topological spaces for the two middle threads are the same but different from the (Zariski) topological space

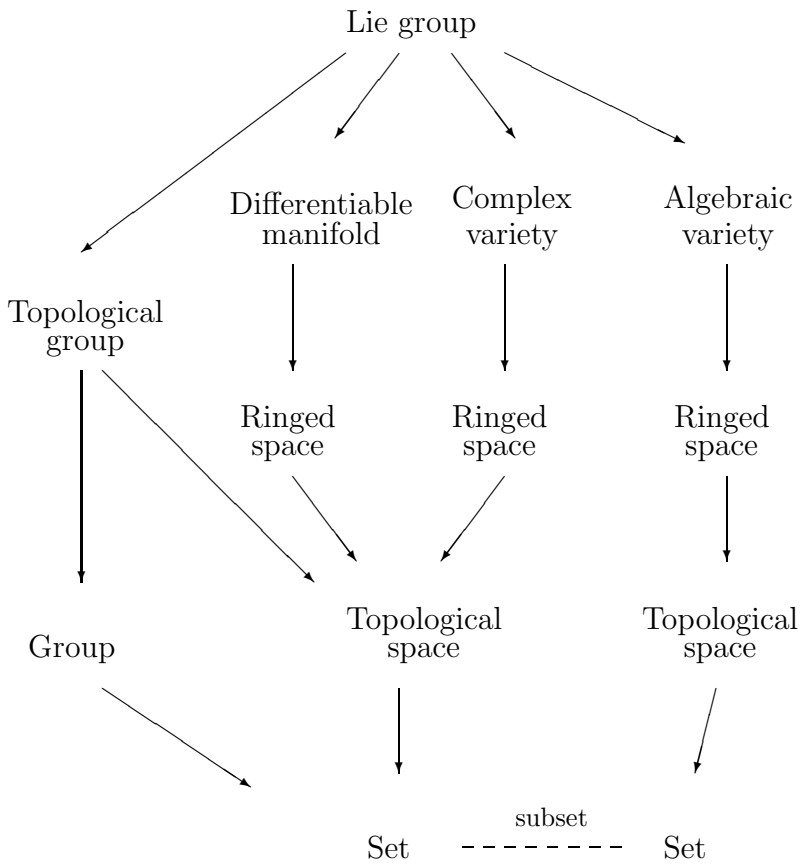

FIGURE 1. Inheritance for the notion of Lie group

in the "algebraic variety" thread. One could identify other inheritance strings, and note that there are intricate relationships between them, for example the structure sheaf of the ringed space underlying the complex variety is a subsheaf of the structure sheaf of the ringed space underlying the differentiable manifold, and the set underlying the usual topological space in the first threads is the subset of closed points in the set underlying the Zariski topological space.

If G is a Lie group, informally we tend to use the letter "G" for any and all of the above things and more.

On the other hand, the details of these inheritance relations would have precisely to be spelled out in any computer verification system.

Letting the same letter "G" mean any of an amorphous mass of notions, represents a huge collection of notational shortcuts. While

clearly impossible to keep all of them, it should be possible to keep some. This could lead to some difficult choices: how to decide which parts to keep and which parts to relegate to a more complicated notation. And in any case, how do we minimize the amount of additional notation necessary to make these shortcuts precise?

The example of a Lie group is deliberately complex in order to illustrate the problem. However, it comes up to varying degrees in pretty much all of modern mathematical practice.

12. Perspectives

We now outline a sketch of what some part of the further development of machine-verified mathematics might look like. This could start with any of a number of basic supports such as [88], [91], [81], [87], [79], ..., some of which are already quite far along.

Attached to the original preprint version of the present paper was the start of a development along the lines of option (10.4) above. The idea of notative sets as being functions on strings, first brought up in [45], was implemented. We get a primitive kind of inheritance, in which different structures corresponding to different strings can coexist in the same object, and the functions extracting these operations from a given object are uniform throughout. Thus for example the theory of groups and the theory of topological spaces would both apply immediately to topological groups. The explanatory section and source code have now been spun off as a separate preprint [71].

The proof development accompanying [71] now treats the notions of ordinal and cardinal. The next part which needs to be addressed would contain the notion of finiteness, various facts about finite sets (notably that any automorphism of a finite set decomposes as a product of transpositions); also various notions of collection or list of objects; etc. The theory of transpositions should help for defining multi-term summation in an abelian monoid.

Most objects are constructed with a single common variable, denoted `(U a)`, thought of as the "underlying set" of `a`. Morphisms between such objects can be considered as some of the maps between underlying sets. This makes it possible to develop a certain amount of categorical machinery at once for all these types of objects.

With this in mind, we can define monoids, groups, rings, algebras, modules; categories, eventually with Grothendieck topologies; presheaves (and sheaves if the category has a g.t.); ordered sets; topological spaces also eventually with other structures so that one could define topological groups and so on.

The counting system has to be extended to the numbering system. We can import much of the beautiful development of Peano arithmetic already in the `Coq` system, as well as the `ZArith` library containing a logarithmically fast implementation of integers. Then we need to define the rationals as a localization of the integers (and for this we might as well call upon the general notion of localization of a ring). Then we can define the reals as the set of Dedekind cuts of the rationals. It is not clear whether it is better to import the real number results from the existing `Coq` library, or to redo a substantial portion, since our approach using things like the axiom of choice, and an untyped point of view, might give substantial simplification here (as opposed to the case of elementary arithmetic where we wouldn't expect any such benefit).

Once we have the reals (and therefore the complex numbers too) we need to establish the basic facts about the topology on these spaces, and along the way note that this opens up the definition of metric space, so we can do the basic facts about metric spaces. This could eventually lead to stuff for analysis such as Hilbert or Banach spaces, and all of the basic theorems of analysis. This is not necessarily urgent although it would be required for the implementation of Hodge theory.

A parallel project is to develop abstract commutative and linear algebra as well as group theory and field theory. For large parts, these theories don't need the real numbers. As pointed out above, the notion of localization would be convenient to have for the construction of the rational numbers. In general, it would be good to start by developing the basic commutative algebra such as is assumed in the beginning of Hartshorne. Once we have a reasonable control over the basic notions of ring, algebra, module etc. then much of commutative algebra can be done without any new notational worries.

Linear algebra could lead to a rather extensive development of the theory of matrices, matrix groups, and group representations.

Again in parallel, we need the notions of presheaf and sheaf over a topological space. Here it seems economical to take the abstract point of view and develop directly the notion of sheaf on a Grothendieck site. Along the way the notion of presheaf on a category yields the notion of simplicial set which forms the basis for homotopy theory. The notion which is directly useful for starting out is that of sheaf on a topological space, so we need a construction going between a space and its corresponding Grothendieck site of open sets. It is unclear whether we would prefer using the full abstract notion of point for a site (or maybe a topos?—but this may lead to universe problems) or whether it is better just to look at the germ for a presheaf on a topological space. In any case, here some *ad hoc* argumentation is necessary, because we

would like the germ of a presheaf to be endowed with the same type of structure as the presheaf; the variety of different types of structures we want to preserve means that it is probably more difficult to come up with a single general theory than just to do each one separately.

Up to here, most of what is listed above has already been treated in one way or another, in at least some user-contributions for at least some of the systems listed in the bibliography. Left open is the problem of coalescing everything together.

Once we have these various theories, we can put them together to obtain the notion of locally ringed space modelled on a certain type of model space. Depending on which model spaces we look at, we can obtain the notions of differentiable manifold, complex analytic manifold, or algebraic variety (scheme). Laurent Chicli has done a formalization up to the notion of (affine) scheme in his thesis [14], using setoids (option (10.2) above).

At this point, perhaps one of the first things to do would be the comparison constructions going from schemes (over $\mathbf{C}$) to complex manifolds to differentiable manifolds. We should then be able to construct projective space, projective subvarieties, various differentiable manifolds, and so on. Another thing to do will be exterior algebra and differential calculus over a manifold.

It would be interesting to try things out by applying this to the case of Lie groups, and seeing how far into the theory one can go.

Integration is probably the next place where we will run into a major notational problem. This is because the verb "to integrate" spans a highly diverse and notationally wide-ranging collection of ideas: there is simple Riemann (or Lebesgue) integration on the real line (already implemented in [19]), but also more general integration of a measure, or integration of a differential form over various things such as a manifold, or a smooth chain; or finally integration of currents and the like. These different notions have a very high degree of overlap, however they are not quite the same, and they require differing notations and even differing levels of notational complexity.

With the theory of integration would come the possibility of getting at the fundamental properties of analysis on manifolds. At this point we will need the theory of Hilbert and Banach spaces (already implemented in [87]), and we should be able to do the basic theorems such as Stokes' theorem, the Poincaré lemma and de Rham's theorem, Hartogs' theorem for complex manifolds, the mean value theorem for harmonic functions, which gives the maximum principle and unicity of analytic continuation, the representation (and regularity) of cohomology by harmonic forms, leading up to the Hodge decomposition for the

cohomology of a complex Kähler manifold (to take an example from my own research interests). At this stage we will also want to have the basic notions of homotopy theory available.

Once all of the above is well in hand, we would be able to state and prove a first theorem which is a really nontrivial application of this wide swath of mathematics: the fact that the first (or any odd) Betti numbers of complex algebraic varieties are even.

We should then be able to start looking at more advanced things going toward current research problems. One area where we can expect important gains is homological or homotopical algebra, for example J. Lipman points out that it would be good to formalize the computations in duality theory such as [4]. More geometrical directions could include for example the construction of moduli spaces of vector bundles, connexions and so forth. Or the investigation of existence problems for connexions over higher dimensional varieties which could at first be viewed just as matrix problems in linear algebra.

Some day in the not-too-distant future we will have a computer-verified proof that the curvature of the classifying space for variations of Hodge structure in the horizontal directions is negative, and all that it entails.

References

[1] P. Aczel. The type theoretic interpretation of constructive set theory. *Logic colloquium 77*, Macintyre, Pacholsky and Paris eds., Springer (1977).

[2] S. Agerholm, M. Gordon. Experiments with ZF set theory in HOL and Isabelle. *Higher order logic theorem proving and its applications (Aspen Grove, UT, 1995)*, Springer L. N. C. S. **971** (1995), 32-45.

[3] G. Alexandre. An axiomatisation of intuitionistic Zermelo-Fraenkel set theory. See
`http://coq.inria.fr/contribs-eng.html`.

[4] L. Alonso Tarrío, A. Jeremías López, J. Lipman. *Studies in duality on Noetherian formal schemes and non-Noetherian ordinary schemes.* Contemporary Mathematics, **244**, AMS (1999).

[5] T. Altenkirch, C. McBride. Generic Programming Within Dependently Typed Programming. To appear, WCGP 2002.
`http://www.dur.ac.uk/c.t.mcbride/generic/`

[6] B. Barras, T. Coquand, B. Werner. Paradoxes in Set Theory and Type Theory. User-contribution of Coq, INRIA Rocquencourt.

[7] Y. Bertot, P. Castéran. Coq'Art. Book in press.

[8] H. Boom. Message on [100], Feb. 16, 2001.

[9] R. Boyer, J. S. Moore. *A computational logic*, ACM Monograph Series, Academic Press (1979).

[10] W. W. Bledsoe, Robert S. Boyer, William H. Henneman. Computer Proofs of Limit Theorems. *IJCAI* (1971) 586-600.

[11] Robert S. Boyer (Ed.): Automated Reasoning: Essays in Honor of Woody Bledsoe. Automated Reasoning Series Kluwer Academic Publishers (1991)

[12] A. Bundy. A Survey of Automated Deduction. *Artificial Intelligence Today. Recent Trends and Developments* (Michael J. Wooldridge and Manuela Veloso eds.) Springer Lecture Notes in Computer Science **1600** (1999), 153-174.

[13] V. Capretta. Universal Algebra in Coq.
`http://www-sop.inria.fr/Venanzio.Capretta/universal_algebra.html`

[14] L. Chicli. Sur la formalisation des mathématiques dans le Calcul des Constructions Inductives. Thesis, Université de Nice-Sophia Antipolis (Nov. 2003).
`http://www-sop.inria.fr/lemme/Laurent.Chicli/these_chicli.ps`

[15] L. Chicli, L. Pottier, C. Simpson. Mathematical quotients and quotient types in Coq. H. Geuvers, F. Wiedijk, eds. *Types for Proofs and Programs*, Springer LNCS **2646** (2003), 95-107.

[16] T. Coquand. An analysis of Girard's paradox. *Proceedings of LICS*, IEEE press (1985).

[17] J. Courant. Explicit universes for the calculus of constructions. In V. Carreño, C. Muñoz, and S. Tahar, eds., *Theorem Proving in Higher Order Logics 2002*, Springer LNCS **2410**, 115-130.

[18] A. Cuihtlauac, Reflexion pour la reecriture dans le calcul de constructions inductives. Thesis, Dec. 18, 2002.

[19] L. Cruz-Filipe. A constructive formalization of the fundamental theorem of calculus. H. Geuvers, F. Wiedijk, eds. *Types for Proofs and Programs*, Springer LNCS **2646** (2003), 108-126.

[20] P. Deligne, J. S. Milne, A. Ogus, K. Shih. *Hodge Cycles, Motives, and Shimura Varieties.* Springer LNM **900** (1982).

[21] S. Feferman. Typical ambiguity: trying to have your cake and eat it too. To appear in the proceedings of the conference Russell 2001, Munich, June 2-5, 2001.

[22] M. Fiore, G. Rosolini. Two models of synthetic domain theory. *J. Pure Appl. Alg.* **116** (1997), 151-162.

[23] B. Fitelson, D. Ulrich, L. Wos. XCB, the last of the shortest single axioms for the classical equivalential calculus, `cs.LO/0211015`, and
Vanquishing the XCB question: the methodology discovery of the last shortest single axiom for the equivalential calculus, `cs.LO/0211014`.

[24] H. Friedman. Finite functions and the necessary use of large cardinals. *Ann. of Math.* **148** n.3 (1998), 803-893.

[25] D. Gabai, G. R. Meyerhoff, N. Thurston. Homotopy hyperbolic 3-manifolds are hyperbolic. *Ann. of Math.* **157** (2003), 335-431.

[26] H. Geuvers. Inconsistency of classical logic in type theory.
`http://www.cs.kun.nl/~herman/note.ps.gz`
See also other publications at `http://www.cs.kun.nl/~herman/pubs.html`

[27] H. Geuvers, H. Barendregt. Proof Assistants using Dependent Type Systems, Chapter 18 of the Handbook of Automated Reasoning (Vol 2), eds. A. Robinson and A. Voronkov, Elsevier (2001), 1149 - 1238.

[28] H. Geuvers , F. Wiedijk , J. Zwanenburg , R. Pollack , H. Barendregt. A formalized proof of the Fundamental Theorem of Algebra in the theorem prover Coq, Contributions to Coq V.7, April 2001,
`http://coq.inria.fr/contribs/fta.html`

[29] J.-Y. Girard. Interpretation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieure. Thèse d'Etat, Université Paris 7 (1972).
[30] M. Gordon. From LCF to HOL: a short history. *Proof, language, and interaction* Found. Comput. Ser., MIT Press,(2000), 169-185,
[31] T. Hales. The Flyspeck Project Fact Sheet. `http://www.math.pitt.edu/~thales/flyspeck/index.html`
[32] J. Harrison. Formalized mathematics. (1996). `http://www.cl.cam.ac.uk/users/jrh/papers/form-math3.html`, see also an `html` version at `http://www.rbjones.com/rbjpub/logic/jrh0100.htm`
[33] H. Herbelin. A program from an A-translated impredicative proof of Higman's Lemma. User-contribution in [81], see `http://coq.inria.fr/contribs/higman.html`
[34] R. Holmes. Undefined terms, and the thread of messages following it (in particular J. Harrison's reply), [99] Volume 3, May 1995.
[35] A. Hohti. Recursive synthesis and the foundations of mathematics. `math.HO/0208184`.
[36] G. Huet. Unification in typed lambda calculus. *λ-calculus and computer science theory (Proc. Sympos., Rome, 1975)* Springer L. N. C. S. **37** (1975), 192-212.
[37] G. Huet, A. Saïbi. Constructive category theory. *Proof, language, and interaction*, Found. Comput. Ser., MIT Press (2000), 239-275,
[38] L. S. van Bentham Jutting. Checking Landau's "Grundlagen" in the AUTOMATH system. Thesis, Eindhoven University of Technology (1977).
[39] M. Kaufmann, J. Strother Moore. A Computational Logic for Applicative Common LISP. *A Companion to Philosophical Logic*, D. Jacquette (ed), Blackwell Publishers (2002), 724-741.
[40] H. Kitoda. Is mathematics consistent? `math.GM/0306007`
[41] K. Kunen. A Ramsey theorem in Boyer-Moore logic. *J. Automat. Reason.* **15** (1995), 217-235.
[42] C. Lam, L. Thiel, S. Swiercz. The nonexistence of finite projective planes of order 10. *Canad. J. Math.* **41** (1989), 1117-1123.
[43] C. Lam, L. Thiel, S. Swiercz, J. McKay. The nonexistence of ovals in a projective plane of order 10. *Discrete Math.* **45** (1983), 319-321.
[44] L. Lamport. Types considered harmful, or Types are not harmless. This appeared under the first title in a posting by P. Rudnicki on [99], Volume 2, Aug. 1994. A revised version with the second title appeared as a technical report. A balanced discussion presenting both points of view is in the next reference.
[45] L. Lamport. L. Paulson. Should Your Specification Language Be Typed? *ACM Transactions on Programming Languages and Systems* **21**, 3 (May 1999) 502-526. See `http://research.microsoft.com/users/lamport/pubs/pubs.html`
[46] Z. Luo. An extended calculus of constructions. Thesis, University of Edinburgh (1990).
[47] M. Maggesi. Proof of `JMeq_eq`, see posting on [100], Oct. 17th 2002, `http://www.math.unifi.it/~maggesi/coq/jmeq.v`
[48] P. Martin-Löf. *Intuitionistic Type Theory*, Studies in Proof Theory, Bibliopolis (1984).
[49] J. McCarthy. Recursive Functions of Symbolic Expressions and their Computation by Machine (Part I), *CACM* **3** (1960), 184-195.

[50] J. McCarthy. Towards a mathematical science of computation, *Proceedings of the Information Processing Cong.* **62**, North-Holland (1962), 21-28.
[51] J. McCarthy. A Basis for a Mathematical Theory of Computation, *Computer Programming and Formal Systems*, North-Holland (1963).
[52] W. McCune, R. Veroff. A short Sheffer axiom for Boolean algebra. `http://www.cs.unm.edu/~moore/tr/00-07/veroffmccune.ps.gz` `http://www.cs.unm.edu/~veroff/`, `http://www.mcs.anl.gov/~mccune`
[53] J. McKinna. Reply to thread "How to prove two constructors are different", on [100], Oct. 6th 2003.
[54] R. Milner. LCF: a way of doing proofs with a machine. *Mathematical foundations of computer science (Proc. Eighth Sympos., Olomouc, 1979)* Springer L. N. C. S. **74** (1979), 146-159,
[55] I. Moerdijk, S. MacLane. *Sheaves in Geometry and Logic*, Springer (1992).
[56] A. Neeman. A counterexample to a 1961 "theorem" in homological algebra. *Inventiones* **148** (2002), 397-420.
[57] D. Nowak. Ensembles and the axiom of choice, on [100], Nov. 25th 1998.
[58] R. O'Connor. Proof of Gödel's first incompleteness theorem `http://math.berkeley.edu/~roconnor/godel.html`
[59] B. Plotkin. Algebraic geometry in First Order Logic. `math.GM/0312485`.
[60] D. J. Melville. Sumerian metrological numeration systems. [6] `http://it.stlawu.edu/~dmelvill/mesomath/sumerian.html`
[61] QED manifesto. `http://www-unix.mcs.anl.gov/qed/manifesto.html`
[62] J. A. Robinson. Theorem-proving on the computer. *J. Assoc. Comput. Mach.* **10** (1963) 163-174.
[63] J. A. Robinson. A machine-oriented logic based on the resolution principle. *J. Assoc. Comput. Mach.* **12** (1965) 23-41.
[64] D. Russinoff. A mechanical proof of quadratic reciprocity. *J. Automat. Reason.* **8** (1992), 3-21.
[65] C. Schmidhuber. Strings from logic. CERN-TH/2000-316, `hep-th/0011065`.
[66] D. Scott. Domains for denotational semantics. *Automata, languages and programming (Aarhus, 1982)*, Springer LNCS **140** (1982), 577-613.
[67] N. Shankar. A mechanical proof of the Church-Rosser theorem. *J. Assoc. Comput. Mach.* **35** (1988), 475-522.
[68] N. Shankar. *Metamathematics, machines, and Gödel's proof.* Cambridge Tracts in Theoretical Computer Science **38**, Cambridge University Press (1994).
[69] I. Shimada. Vanishing cycles, the generalized Hodge conjecture, and Gröbner bases. `math.AG/0311180`.
[70] A. Simpson. Computational adequacy in an elementary topos. *Proceedings CSL '98*, Springer LNCS **1584** (1998), 323-342.
[71] C. Simpson. Set-theoretical mathematics in Coq. Preprint with attached proof files, `math.LO/0402336`.
[72] S. Simpson, ed. *Reverse Mathematics 2001*, to appear.
[73] T. Streicher. Lifting Grothendieck universes (with M. Hofmann); and Universes in toposes. Preprints available at `http://www.mathematik.tu-darmstadt.de/~streicher/`

[6] The relevance of ancient numbering systems as one of the origins of formalized mathematics was mentionned in [32].

[74] A. Trybulec. The Mizar-QC/6000 Logic Information Language, *ALLC Bulletin* Vol.6, No 2 (1978).
[75] M. Wenzel, F. Wiedijk. A comparison of Mizar and Isar. *J. Automated Reasoning* **29** (2002), 389-411.
[76] B. Werner. Sets in Types, Types in Sets. *Theoretical aspects of computer software (Sendai 1997)*, Springer LNCS **1281** (1999), 530-546. http://pauillac.inria.fr/~werner/publis/zfc.ps.gz
[77] B. Werner. An encoding of Zermolo-Fraenkel Set Theory in Coq: see http://coq.inria.fr/contribs-eng.html.
[78] B. Werner. Une théorie des constructions inductives. Thèse d'Etat, Univ. Paris 7 (1994).
[79] ACL2 system: http://www.cs.utexas.edu/users/moore/acl2/acl2-doc.html.
[80] Alfa system (formerly ALF): http://www.math.chalmers.se/~hallgren/Alfa/
[81] Coq system: http://coq.inria.fr/, especially the reference manual : http://coq.inria.fr/doc/main.html.
[82] Ghilbert system: http://www.ghilbert.org/.
[83] Helm Coq-on-line library (University of Bologna): http://www.cs.unibo.it/helm/library.html
[84] HOL system: http://www.afm.sbu.ac.uk/archive/formal-methods/hol.html.
[85] HOL-Light: http://www.cl.cam.ac.uk/users/jrh/hol-light/
[86] IMPS system, see particularly the theory library: http://imps.mcmaster.ca/.
[87] Metamath system: http://www.metamath.org/.
[88] Isabelle system : http://www.cl.cam.ac.uk/Research/HVG/Isabelle/.
[89] The Elf meta-language. http://www-2.cs.cmu.edu/~fp/elf.html
[90] LEGO system : http://www.dcs.ed.ac.uk/home/lego/.
[91] Mizar system : http://mizar.uw.bialystok.pl/.
[92] Nuprl system: http://www.cs.cornell.edu/Info/Projects/NuPrl/nuprl.html.
[93] The PhoX proof assistant: http://www.lama.univ-savoie.fr/sitelama/Membres/pages_web/RAFFALLI/af2.html
[94] PVS (Proof Verification System): http://pvs.csl.sri.com/
[95] TPS (Theorem Proving System): http://gtps.math.cmu.edu/tps.html
[96] Z/EVES system, see particularly the Mathematical Toolkit: http://www.ora.on.ca/z-eves/.
[97] Formal Methods web page at Oxford (this has a very complete listing of items on the web concerning formal methods): http://www.afm.sbu.ac.uk/.
[98] F. Pfenning. Bibliography on logical frameworks (449 entries!) http://www-2.cs.cmu.edu/~fp/lfs-bib.html
[99] The QED project (see particularly the archives of the QED mailing list, volumes 1-3): http://www-unix.mcs.anl.gov/qed/
[100] Coq-club mailing list archives: http://pauillac.inria.fr/coq/.
[101] *Journal of Formalized Mathematics* : http://mizar.uw.bialystok.pl/JFM/.
[102] arXiv e-Print archive http://arXiv.org/

[103] MathSci Net. http://www.ams.org/mathscinet (by subscription)
[104] The Google search engine. http://www.google.com

CNRS, Laboratoire J. A. Dieudonné, Université de Nice-Sophia Antipolis, 06108 Nice, Cedex 2, France
E-mail address: carlos@math.unice.fr
URL: http://math.unice.fr/∼carlos/