PERSPECTIVES ON FORMAL VERIFICATION

by
Harvey M. Friedman
Distinguished University Professor of Mathematics, Philosophy, and Computer Science Emeritus
Ohio State University

The Fifth International Conference On Certified Programs And Proofs
St. Petersburg, Florida
January 18, 2016

This material is based upon work supported by the National Science Foundation under Grant No. CCF-1162331. Any opinions, findings, conclusions, or recommendations expressed here are those of the author and do not necessarily reflect the views of the National Science Foundation.

It's a nice idea to have an opinionated speaker from outside your field give their views before you start the serious part of your meeting.

All of you are either involved in real implementations or are far closer to them than I am.

I view formal verification - both of mathematical theorems and computer systems - as of fundamental importance, both practically and theoretically. It has a long way to go to reach ultimate goals, but has already been of practical and theoretical importance.

I am a foundationalist looking at verification from the outside. I am particularly interested in the formulation of strategic goals for fields, that speak to the general intellectual community.

This usually overlaps with, but by no means coincides with, the priorities and viewpoints of the field under discussion.

As a consequence, most of the goals that capture my imagination may not be particularly high on your list of priorities.

Also, I am not likely to be familiar with some of your efforts that make significant contributions to my favorite goals. I look forward to hearing more about these.

For whatever it is worth, I will be putting this talk on my website at

https://u.osu.edu/friedman.8/ Downloadable Lecture Notes
1. FOUNDATIONS OF MATHEMATICS

I became convinced of the theoretical importance of verification long ago. By the early 1900s our present fully rigorous formalization of mathematics clearly emerged. This is of course the usual ZFC axiomatization.

Let's take a look at the general shape of this axiomatization. We have the well known split:

1 Purely Logical Part. First order predicate calculus.
2. Proper Axioms. Most conveniently, ZFC.

A number of conceptual issues immediately arise here. For instance, what does purely logical mean? Where does logic end and set theory begin?

Should second order and higher order logic be in the purely logical part, or be viewed as set theoretic?

First order predicate calculus famously supports the great Gödel Completeness Theorem. The situation with second and higher order logic is different. Should this count as a reason to keep second and higher order logic out of 1 - the purely logical part?

Perhaps the point is that first order predicate calculus is the minimum setup that does the job. But what does "minimum" mean? For that matter, what does "setup", "job", and "do" mean?

There are bits and pieces of work related to such questions, including striking characterizations of first order predicate calculus as a semantic system. But for present purposes, the work is far from conclusive, and doesn't really get to the heart of the matter decisively.

In this murky environment, enter formal verification.

True, formal verification doesn't really address these questions. However, formal verification does directly address a related issue in a largely decisive way:

IS THE SIZE OF PURELY FORMAL PROOFS OBTAINED FROM SEMIFORMAL PROOFS OF MAJOR THEOREMS REASONABLE? E.G., AT MOST TEN THOUSAND PRINTED PAGES? OR IS THERE AN EXPONENTIAL TYPE BLOWUP INVOLVED WHEN WE MOVE FROM THE USUAL SEMIFORMAL PROOFS CREATED BY MATHEMATICIANS TO A FORMAL PROOF?

IS THE MENTAL EFFORT INVOLVED IN GOING FROM A FULLY UNDERSTOOD SEMIFORMAL PROOF TO A FORMAL PROOF REASONABLE? Of course, for this community, it is second nature that there is no kind of exponential blowup from the semi-formal proof to the formal proof, and in fact the blowup is rather controlled.

I have seen modest linear blowups proposed with some real justification, although perhaps not well justified for such major theorems as FLT that have not been formalized.

The situation is similar, though perhaps not as clear cut, for the mental effort required.

We have seen some striking computer assisted proofs.

CAN WE DEMONSTRATE THAT COMPUTER ASSISTANCE IS REQUIRED, IN THE PRACTICAL SENSE, FOR THE PROOF OF SOME MAJOR MATHEMATICAL THEOREMS?

The obvious way of demonstrating controlled blowup of semiformal proofs to formal proofs is to actually construct the formal proofs. So this raises the following question.

CAN WE PROVIDE EVIDENCE THAT GOING FROM SEMIFORMAL PROOFS TO FORMAL PROOFS IS REASONABLY CONTROLLED, THAT IS MORE CONVINCING THAN SIMPLY CREATING FORMAL PROOFS? CONTROLLED, BOTH IN THE SIZE AND MENTAL EFFORT SENSE?

Before major theorems were formalized, most of the math community tended to think that practical formalization was either

i. Impossible because of an intrinsically necessary size blowup; or at least

ii. Impractical in terms of the outrageous effort required to get a handle on an unmanageably enormous variety of details.

Most mathematicians - even top ones - are not at ease with first order predicate calculus. They do not find it natural, and they rely entirely on natural instincts in their construction of semiformal proofs.
They have spent almost no conscious time studying or thinking about the "rules of the game". They would like to think that the "rules of the game" are fluid, and so they don't instinctively "feel" - as we do - the power of predicate calculus and formalization.

In a way, I can see how they might be skeptical. For instance, consider trying to formalize the reasoning done in physical science.

Say in the construction of experiments and the interpretation of their outcomes. What factors are important and what factors should be neglected?

If I drive into the parking lot to enter an atom smasher facility while an experiment is active, does this affect the outcome of that experiment? "Obviously" not! What's involved in formally verifying that?

Consider "obvious" judgments as to the strength of evidence for scientific theories based on confirmation. Even rigorous formulation of scientific theories are lacking from the formal point of view.

Yet more challenging is the formalization of judgments about the musical expressivity in the playing of simple piano pieces.

So my point is that the actual controlled construction of perfectly formal proofs of major theorems is a striking, and arguably surprising, development - even if it is now a commonplace idea to everyone here.

But consider going much further along these lines, now that we have such a large and growing inventory of actual formalized proofs of theorems ranging from the trivial to the major. What can we say about their structure?

The crudest aspect of formalized proofs is their size.

HOW STABLE IS THE SIZE OF FORMALIZED PROOFS WHERE THE SAME SEMIFORMAL PROOF IS BEING FORMALIZED WITH THE SAME SOFTWARE BY DIFFERENT HUMANS? WHAT IF WE VARY THE SOFTWARE BEING USED?

There are obviously much more interesting features than the size of formal proofs. Some have been intensively studied in f.o.m. such as levels of constructivity, and levels of strength (interpretation power) of set theoretic axioms. But these do not take any real advantage of actual formalization.

WHAT ARE SOME IMPORTANT FEATURES OF ACTUAL FORMALIZED PROOFS? E.G., DEVELOP A STRATEGIC ORGANIZATION OF CRITICAL INFERENCE STRATEGIES, AND DOCUMENT THEIR FREQUENCY OF USE. HOW STABLE AMONG USERS AND AMONG SYSTEMS?

There is a possible game changing development here for f.o.m. Suppose that we discover a significant property shared by actual formalized proofs. Say that you never do such and such. Or that you always do such and such. Or some combination. But some theorems of ZFC don't have a proof with this property.

Let's call the proofs obeying this newly discovered condition, purple proofs.

Then maybe much of f.o.m. needs to be reworked, replacing provability by purple provability.

CAN THE PURPLE CONSISTENCY OF ZFC BE PROVED WITHIN FINITE SET THEORY?

Of course, we expect that

THE PURPLE CONSISTENCY OF ZFC CANNOT BE PROVED WITHIN PURPLE FINITE SET THEORY.

If so, any proof within finite set theory of the purple consistency of ZFC will be non purple.

This raises the prospect that metamathematical proofs may have some feature that distinguishes them from normal mathematical proofs. This would provide a formal difference between mathematical areas. More generally,

WHAT DISTINGUISHING FEATURES (QUANTITATIVE OR QUALITATIVE) DO FORMALIZED PROOFS IN THE VARIOUS AREAS OF MATHEMATICS HAVE, THAT WOULD CONTRAST ONE AREA FROM ANOTHER?

2. READABILITY

Mathematicians generally have little tolerance for even semiformal proofs that don't look simple and beautiful and easy on the eyes.

They would much rather see a proof sketch, with lots of even significant steps missing, that is readily absorbed, and for which they can instantly tell that they can easily fill in any level of detail that is demanded.

In this unforgiving environment, there is quite a high bar for the readability of completely formalized proofs. Yet I think that the challenge can be met in interesting ways.

The general problem can be formulated as a kind of Turing test for formal proofs.

CREATE FORMALLY VERIFIED TREATMENTS OF VARIOUS BASIC AREAS OF MATHEMATICS WHICH ARE SO READABLE THAT THEY CAN BE EFFECTIVELY USED AS TEXTBOOKS WITH REAL WORLD NON COMPUTER SCIENCE STUDENTS.

This will require a lot of reworking of verification systems along lines that you may find to be of low priority. Here are four.

i. The mathematicians want free logic. If you tell them that you want 1/0 = 0 for your convenience, then they will simply tell you that it is inconvenient for them to talk to you.

ii. They won't listen to any elaborate typing. Their idea of typing is, say, $(\forall x \in \Re)$ (P(x)), or $(\forall x \in f(y))$ (P(x,y)). Things get a little tricky with, say, {x+y: P(x,y)}. Is this a two dimensional sum, or is x fixed, and we are taking a one dimensional sum over y? They want to write things like

{x+y: P(x,y)}, x fixed.

iii. They want to break statements up using English into bite sized chunks: Let x,y be gadgets. Let z be a badget. Then blah blah blah.

iv. They make conventions: Until further notice, we use α,β,γ for gadgets with blah blah blah.
There are quite a number of such things that have to be addressed nicely in order to have any chance of readability.

It would be rather useful to have an understanding of levels of obviousness.

Perhaps readable texts can be constructed with various levels of obviousness.

This corresponds to the level of the steps that need to be filled in by the verifier when processing the text. But we are not really at that point where this is practical.

Our systems normally can't automatically recognize anything that isn't pretty low level obvious, unless it falls within our very limited inventory of decision procedures, or the system is very special purpose. Nevertheless,

DEVELOP AN INFORMATIVE THEORY OF LEVELS OF OBVIOUSNESS.

A candidate for the lowest interesting level of obviousness is "self proving". The idea is that you simply do the obvious unraveling, with no creativity.

However, this notion must take into account the use of prior definitions and theorems.

Generally, this prior material cannot be just plugged in, but must be manipulated, and there may be more than one item that needs such manipulation. This kind of obviousness quickly becomes beyond the software to find on its own.

So this suggests the need to use citations in the text, perhaps accompanied by indications on how they are to be manipulated. This leads to major design issues.

We would like the author to be able to create readable text based on such obviousness, but it is not at all clear if

the author can reliably foresee what happens when the software tries to fill in the steps, based on such hints.

Of course, it is reasonable to require that the author create such readable text hooked up to the software that tests for obviousness on the spot. Maybe with the proper author friendly feedback from the software, such readable text based on such obviousness can be readily constructed. I think this has to be designed with a lot of care.

3. DECISION PROCEDURES

We now have a small but growing arsenal of decision procedures for little fragments of mathematics. One of the most famous is that for the ordered field of reals. However, we rapidly run into blowups, and you can make a career developing algorithms for fragments that avoid blowups in more and more contexts.

I would like to suggest a related idea that sounds very bold. Why not try to decide absolutely any interesting class of mathematical statements, even if it is well known to be algorithmically undecidable?

Here by "decide", I simply mean: parameterize the statements in question, and start with very small choices of parameters, creating modest sized finite sets of target statements.

Develop tools to handle these efficiently, and then slowly raise the parameters, developing more tools, etc.

CAN WE EXAMINE THE CORPUS OF FORMAL PROOFS AND IDENTIFY MAYBE 200 CORE (MOSTLY) UNDECIDABLE DECISION PROBLEMS, AND "DECIDE" THEM IN THE ABOVE WAY, WITH THE INTENTION OF DRAMATICALLY ADVANCING GENERAL PURPOSE VERIFICATION?

4. CERTAINTY

A formally verified proof is supposed to be certain. The existing high confidence in verification probably makes this a fairly low priority issue for you. Especially in the context of verification of mathematical theorems, where mistakes are not catastrophic. However, there are other contexts, with catastrophic consequences, where certainty is of a higher priority.

I will stick to verification of mathematics here, and stay away from catastrophes.

Just how certain can we make a mathematical statement referring to infinitely many objects, as they normally do?

For our context, we might as well identify the certainty of a mathematical statement here with its provability in ZFC.

Of course, an issue of an entirely different kind, suitable for an entirely different talk, is whether the use of ZFC leads to certainty. For that kind of issue, one gets interested in minimizing the fragment of ZFC one is using. But, that's not an issue to be addressed here.

IS THERE A REALIZABLE "ABSOLUTE" CERTAINTY - OR AT LEAST A KIND OF CERTAINTY THAT CANNOT BE STRENGTHENED? OR IS THERE NO STRONGEST KIND OF CERTAINTY?

My instincts are that the certainty issue involves an entangled web of conceptual challenges that may defy full analysis, but there are plenty of opportunities for important deep insights.

There is already a nasty problem at the outset. Most, but not all, important mathematical theorems have fairly complicated formal statements. The statements normally sit on a hierarchy of standard developments. There may be some substantial additional structure that needs to be presented that is peculiar to the mathematical theorem being treated.

So mistakes can be made in the very statement of the theorem. This is like "there is a bug in the spec".

HOW DO WE TREAT THIS THEOREM STATEMENT ISSUE? WHAT DOES IT MEAN TO SAY THAT WE HAVE GIVEN A CORRECT STATEMENT OF A THEOREM? OR CAN WE AT LEAST SAY SOMETHING INTERESTING ABOUT THIS ISSUE?

On another note, the uninformed thinker about certainty is expected to say something like this.

How do you know that your verification system is correct? After all, it probably isn't perfect, since it is a complicated piece of software sitting on top of probably an impossibly complicated operating system running on some possibly buggy commercial hardware, etc.

We know how to evade this issue by using certificates. We simply use the verification system as a means, or major step, for creating a file. Then we subject the file to an ultimate verifier constructed simply for the purpose of verifying that the file meets certain criteria of perfection.

The plan is thus to shift the burden of certainty from the verification system to this ultimate verifier. We would want the entire ultimate verifier system, hardware and software, to be particularly simple and transparent.

A specially designed low level automaton would seem to be ideal for this purpose. The file to be verified should be a kind of enriched marked up ZFC proof, where the markings facilitate the action of the automaton, and also allow its hardware/software design particularly transparent, so that this design can itself be readily subjected to verification.

HOW SHOULD THIS ULTIMATE VERIFIER BE DESIGNED, BOTH HARDWARE AND SOFTWARE? HOW CAN WE DESIGN THE ULTIMATE VERIFIER IN ORDER TO BE BEST SUBJECT TO VERIFICATION? IS THIS A NEVER ENDING PROCESS WHICH RESULTS IN MORE AND MORE CERTAINTY? OR IS THERE A NATURAL DEMONSTRABLY MINIMUM CORE? WHAT EXACTLY IS A CORE HERE? CAN WE PROVE THAT IT IS MINIMUM? TO WHAT EXTENT CAN WE OR SHOULD WE TAKE INTO ACCOUNT PHYSICS ISSUES SUCH AS COSMIC RAYS?

We still have to argue that if a file passes the ultimate verifier, then in fact, we have provability in ZFC.

This is a piece of combinatorial mathematics that we also want to formally verify in the same way.

No matter what the details of this approach to certainty is going to be, we must have human beings examine something. We want to avoid requiring that they have any special abilities or training.

HOW CAN WE ASSURE THAT A WIDE SPECTRUM OF HUMAN BEINGS CAN BE TRUSTED, IN AGGREGRATE, TO APPROPRIATELY EXAMINE SOME APPROPRIATELY TINY BUT SUFFICIENT ULTIMATE CORE OR EVOLVING CORES? WHAT SHAPE SHOULD THIS CORE OR CORES TAKE?

HOW DO WE WANT TO INTERACT WITH THE HUMANS THAT ARE CALLED UPON TO EXAMINE THIS ULTIMATE CORE? WHAT EXACTLY IS THEIR MISSION?

5. EDUCATIONAL ASPECTS

I want to close with another kind of issue for verification.

HOW DO WE GET FORMAL VERIFICATION PROPERLY EMBEDDED IN THE COMPUTING CULTURE?

Formal verification depends on having formal specs. Not enough computer systems come with formal specs. Computer executives think that asking developers to write formal specs will unacceptably slow product development.

So it would seem that the formal spec revolution largely rests on the Universities. Students need to understand the value of formal specs and how to create them.

The truth is that formal spec construction requires a relatively significant amount of mathematical sophistication.

This line of thought suggests that the formal spec revolution needs to be furthered by strategically reforming the mathematics curriculum. It is a natural step to take on the entire math curriculum K-16 (13-16 is the undergraduate component).

I hope you have found some amusing food for thought here, and I thank you for listening!