The Outlook for Formal Proof

Michael Nahas*

Sept. 2013

1 Introduction

A year ago, I fell in love with the idea of formal proof. I had seen the networked computer invade industry after industry: music is now sold as MP3s and newspapers are mostly read on websites. It seemed inevitable that all mathematics would be done on the computer and I left my career in industry to help make it happen.

It's now a year later and I'm better informed about formal proof. Before investing more of my time and my money into this field, I thought it would be good to look at where formal proof is going.

This document is a "market analysis", a document that is commonly written in industry before a company invests. It looks at the motives of the customers, the alternative technologies that the customers might use, etc.. A good market analysis identifies the people who are most likely to use a new technology and what features they care about most.

2 What is Formal Proof?

A formal proof is a proof made using the fundamental rules and axioms that define mathematics.

Since the days of Euclid, math has be based around proofs. Mathematicians weren't very precise until the late 1800's and early 1900's. Cantor made imprecise concepts concrete with set theory. Frege formalized logic. Hilbert and Tarski both produced new axioms for geometry.

Russell and Whitehead made the extraordinary effort to try formal mathematics. They were not able to prove much — it took them hundreds of pages to prove "$1 + 1 = 2$" — but along they way they were able to clarify the foundations of mathematics. Their work highlighted the role of powerful unusual axioms such as "the axiom of infinity", "the axiom of reducibility" and "the axiom of choice".

*michael@nahas.com

But human mathematicians could not work at that low-level of detail while still creating high-level concepts. Thus, the majority of proofs are "informal", where it is assumed that, with an extraordinary amount of effort, the proof could be converted into a formal proof. Mathematicians were willing to accept a certain chance of mistakenly accepting an unprovable theorem in order to make faster progress.

3 Formal Proof for Hardware and Software

One application of formal proof is to prove software or hardware correct. Everyone wants their calculations to be correct, but in any human endeavor, there is the chance of mistakes. We can never be 100% sure there are no bugs. Thus, the penalty for being incorrect is just one of the costs associated with a piece of software or hardware. There are many techniques for reducing bugs, each with their own costs.

And economics will drive which techniques are used.

To see the economic factors at work, let's look at an example: computer hardware. Computer chips cannot be changed once they are manufactured. If there's a bug, fixing it usually takes a lot of time and money. Most hardware has a simple well-known specification, so if a bug exists, customers can determine that it is the hardware at fault. So, the manufacture can be held accountable. With the high cost for even a single bug and high accountability, economics says that hardware manufactures will be willing to spend a lot to decrease the chance of a bug.

The result is that hardware is often built using formal proof. But economics also drives the manufacturer to choose the lowest cost method for generating those formal proofs. The result is heavy use of automated theorem proving (ATP). Specifically, model checking and equivalence checking, which are efficient for certain classes of problems. Only when ATP cannot be used, manufacturer use the more general — and more expensive — technique of interactive theorem proving (ITP).

Hardware has possibly the strongest economic motive for correctness. They also have a problem that fits a mathematical specification.

The best way to view the problem economically is:

$$cost\ of\ bug \cdot probability\ of\ bug - cost\ of\ technique$$

For hardware, the cost of a bug is huge, so the manufacturer will accept very expensive techniques to lower the probability of a bug. Other fields with high costs for mistakes: medical devices, satellites, or aircraft, may also support formal proof.

A more enlightening way to write the economic equation is:

$$cost\ of\ bug \cdot probability\ of\ bug - cost\ of\ technique \cdot 100\%$$

This shows the psychology of the economic decision. The cost of a bug-reducing technique is *always* paid. And paid upfront. The decrease in the probability of a bug is never seen, only inferred.

Companies that make airplanes have an existing technique (or set of techniques) for reducing the probability of a bug. And currently the planes built using that technique fly. So, for those companies, the inferred probability of a bug with the current technique is effectively 0%.

Such a company will only change to a new technique, like formal proof, if it is cheaper than their current technique, which they find effective. Moreover, the company knows that any change of technique will (1) increase the probability of a bug and (2) incur some start-up cost, so any new technique will have to be significantly cheaper than their current technique. Thus, ATP (possibly domain-specific ATP) is likely to be a big part of any usage.

If formal proof is not significantly cheaper, the only way it will get adopted is if there is a catastrophic event that makes companies believe that their current techniques are not effective. For example, Intel started using ITP after the expensive FDIV bug. At that point, the managers of the company knew both the cost of a bug *and* that the inferred probability of a bug with their current tools was not 0%. So, change is possible, but it will probably take a similar kind of event for another industry to switch to formal proof.

Tom Hales[1] noted that last week's revelations that the NSA can hack many implementations of encryption may be another industry catastrophe that creates demand for formal proof. There may be new demand for formal specifications of protocols and proven implementations of those specifications.

I feel I must also mention that there currently exists a significant mismatch between the languages used in medical devices and aeronautics — C, for example — and the languages use in most proof systems. This makes them far less likely to be used because of application constraints (running without garbage collection).

4 Formal Proof for Mathematics

I fell in love with the idea of remaking mathematics — where every proof would be a formal proof. But new technology isn't always adopted or may not be adopted universally. So, like with hardware and software, let's look at the economic motivations of the people who might use formal proof: mathematicians.

[1]Comments attributed to Tom Hales come from my notes and remembrances of a conversation on 9 Sept. 2013 (4 days ago). I have emailed Hales to check that what I've written correctly reflects his opinions. I have not heard back from him yet.

You might say "if a mathematician was economically motivated, they'd be in a different field!", but what that statement really means is: mathematicians are not motivated by *money*. Mathematicians are motivated by other things. Their biggest motivation is being respected as "smart" by their peers. With that "smart"ness comes other rewards such as prestige, awards, and jobs.

Currently, "smart"ness is generally demonstrated by proving a new theorem and publishing the informal proof, vetted by a few peers.

Why informal proofs? One reason is that formal proofs require lots of tedious details and doing those by hand is annoying. When communicating with another expert in the field, it is much faster and easier to write "obvious". Another reason not to use formal proofs is that the lengthy details get in the way of communication. Authors want to communicate more than just correctness, they want to demonstrate their "smart"ness by showing off the techniques they used to prove the theorem. Lastly, they write informal proofs because that's all that is required. An informal proof is a lower hurdle and since that's all that's necessary for the community's stamp of "smart", that's all they do.

So, is formal proof likely to replace informal proof? The economic version of this questions is: Will a mathematician be acknowledged as "smart"er by using formal proof instead of informal?

Let's start by reviewing the situation around the major existing formal proof: Thomas Hales's forthcoming proof of the Kepler Conjecture. It took Hales 10 years to come up with his initial proof. That proof came in two parts: an informal proof that broken the problem into a few thousand sub-problems and then a few computer programs that solved the sub-problems. After 4 years of inspection, the referees decided to accept Hales's informal proof but not the computer programs that solved the sub-problems. Given the significance of the theorem (it was open for 400 years) and that the software was certain to contain bugs (given its size), the referees couldn't be confident of the result.

Without the peer approval, Hales decided to do a formal proof — not just of the software but of the informal proof as well. The informal proof was written precisely by Hales, with most of the formalization done by a group of formalizers in Vietnam. The software part of the proof was formalized by Hales's students, generating 3 PhDs. The formal proof and proven software were slightly different from the original versions, to ease formalization and exchange cheap compute time for expensive formalization time. The formalization has taken slightly more than 20 person-years of work.

Some of the conclusions to draw from Hales's work are:

- formal proof can work at the scale of "real" math
- formal proof requires different skills than informal proof
- formal proof is expensive
- formal proof can be bought

Now let me revive the economic equation from proving hardware or software correct:

$$cost\ of\ bug \cdot probability\ of\ bug - cost\ of\ technique \cdot 100\%$$

Here, a bug is a publishing a faulty proof. The two techniques for reducing bugs are "informal proof checked by referees" and "formal proof checked by a computer".

Right now, the total cost of a formal proof to the author is more than 3 times that of an informal proof. Mathematicians will steer clear of formal proof, because they would rather publish 3 informal proofs rather than 1 formal proof. The only time a mathematician will do a formal proof is if (1) an informal proof would not be accepted and (2) the credit for the formal proof would be more than 3 times the value of the average informal proof.

As for "probability of a bug", it seems clear from Hales's referees that mathematicians believe that referees are good enough for human-written informal proof of extreme complexity, but not for proofs based on software.

So, for the near term, formal proofs will be rare. Formal proof will be used only on very important theorems that are proven using computers. And those proofs will have to be ones where the software part cannot be rewritten as informal human-written proof.

Because the skill set for formal proof is specialized, I expect that the people writing the formal proofs will probably be trained in that, and not in general mathematics.

Also affecting the near term is that Hales will get credit for being first. That uniqueness has value amongst his peers. Future mathematicians who write formal proof will not get that distinction. So, it may be a while until another mathematician steps up to do another formal proof.

Hales commented each field of mathematics has its own standards. When he publishes his proof, his field will probably require formalization of any programs, but other fields might not. Thus, because his field is using formal proof, it will encourage new computer-savvy mathematicians into *other fields* because they won't have the costly requirement of formalizing their software. So, formal proof can actually have a negative effect on the number of mathematicians in a research field.

What about the long term? Will every future proof be formal?

The near-term market is a "niche market" — formal proof will be used but only by a few people. Basically, the question is: will formal proof become a "mass market" product, used by everyone. Cellphones are an example of a technology that started in a niche market — used by traveling salespersons — and later became mass market.

Cellphones grew for many reasons. First, the "niche market" may have been small but it was rich enough to keep the industry going. Over time, cellphones got cheaper and smaller, and expanded into other niche markets that needed mobile communications: groundskeepers, repairmen, executives, etc..

Remember, cellphones were still expensive and had worse quality than land-lines, but the value of being always connected was worth it for these customers. Eventually, the price dropped enough to be close to the price of land-lines. Then cellphones entered the mass market.

So, will formal proof follow this kind of path? I think that between research funding and projects like Hales's, the field will keep going. I also think that the price of formal proof will drop over time - immediate reasons include: better libraries, improved automation, and better technique. So, I think it is probable that formal proof expands beyond its initial niche market. Low-level proofs (ones close to the foundations of mathematics) will be proven formally. When a person outside academia wants to publish their proof of P = NP, they will be asked to prove it formally first.

But becoming "mass market" is a different kind of jump. Color printers, supersonic aircraft, mechanical pencils ... many useful technologies get used without taking over the whole market.

Cellphones became mass market by offering more value for roughly the same price. A person with a cellphone can almost always be reached, no matter where they are. The call quality may be worse than a land-line, but being able to communicate at any time was seen as more valuable. So as the price of a cellphone approached the price of a land-line, it took over.

Will formal proof ever be more valuable than its cost?

I don't know.

We need to keep two things in mind. First, a proof is not what a mathematician is after — it's being viewed as "smart" by their peers. Second, the write up of a proof for publication is only part of the process. The major steps might be:

- picking a theorem to prove (or a general research area)
- learning about the theorem and research area
- coming up with the "ah-ha" of the proof
- writing up the proof
- submitting it for approval by referees
- evaluation by the referees

Even with a formal proof (where its correctness is a given), the mathematician must write up a human language summary of the proof and its place in mathematics, so that referees can evaluate the significance of the proof and how much public acknowledgment of "smart"ness to award to its creator. So even if a formal proof could be generated for free, it doesn't mean that a mathematician can publish that many more papers. Mathematicians will spend less time refereeing, but I don't think that that is a significant portion of a mathematician's time.

For a system with interactions, like the evaluation of math proofs, you'd expect a positive feedback loop where a participant would get more value out of the system by using the new technology. Upon reaching "critical mass", that excess value would build exponentially, until the whole community adopted the technology. Right now, I don't see a strong feedback signal. It may lessen the time each mathematician spends as a referee checking the correctness of a proof, but I don't think that is a significant portion of a mathematician's time. Formal proof does not improve the work of a mathematician such that they will get more papers published, that is, get more accolades of "smart"ness by their peers.

There is somewhat of a network effect. If everyone in a field of mathematics was publishing formal proofs, it would lower the costs for a mathematician to switch from informal to formal proofs. However, even if everyone else in the field was writing formal proofs, if a mathematician could publish informal proofs, I think they would. Because they did not have to spend the time learning the formal proof tools, they would probably publish more papers in their career than their peers and get more approval of being "smart".

There are two lurking questions here.

The first: Will formal proof affect other areas of mathematics besides publication and certification of correctness? The first three steps of creating a proof were:

- picking a theorem to prove (or a general research area)
- learning about the theorem and research area
- coming up with the "ah-ha" of the proof

Formal proof may help a little in picking a theorem to prove, but I don't think by much. Learning about a field is an area for human communication — books written in English, not formal statements of lemmas. A powerful method for searching of existing lemma would be useful, but the lemmas don't need to be proved formally for that to exist. As for "ah-ha"'s, I don't know enough about how those happen to make a statement.

Hales commented that he didn't believe formal proof, by an individual or by the community, would improve productivity in any of the other steps in the process.

The second lurking question is: What if we change how the mathematician's peers award "smartness"? The answer here is, yes, that would have a significant effect on the usage of formal proofs. And, given some anger in the math community at the current publication regime (c.f. protests of the publisher "Elsevier"), change is conceivable. Also, the existing publication process — writing LaTeX with the target of publishing on paper — may change with the prevalence of ebooks.

Right now, peer approval is for publication of proofs of unproven theorems. Occasionally, peer approval is given for a simpler or dramatically different proof

of a theorem that has already been proved. If that could change so that mathematicians got approval for rewriting a proof formally and adding it to a library, that would definitely change things. However, I don't see that happening until formal math is so developed that the rewriting of proof is just a way to fill-out the library.

There does exist a new community (Homotopy Type Theory) that seems to have adopted the social process of approval only for formal proofs. They are working with an unfamiliar set of axioms and, while they have intuitions about what should be true, their confidence is only justified when a formal proof is presented. They are working very close to the axioms, so the cost of a formal proof vs. an informal proof is low. Additionally, they are a new community so everything known has been formalized. I believe this social structure is stable and will continue, but I doubt other communities could be convinced to change to it.

Another possibility is an ideological change. In the early 1900's, "axiomization" and its rigor took over mathematics. Proofs changed from being "arguments" to "being able to be reduced to rules and axioms". Existing proofs had to be revisited to see if they could be reduced to axioms, or to see what axioms they required. If a new foundation for mathematics emerges because of formal proof, it's possible that new proofs will need to be written anyway and those proofs might be formal proof.

Even if every proof is not a formal proof, it's possible that mathematicians will adopt a "semi-formal" style for proof. Features might include:

- formal proof of difficult lemmas (with most remaining informal)
- formal statement of the theorem and its assumptions (but informal proof)
- informal proof with a structure close to that of a formal proof (c.f. Lamport's "How to Write a Proof")

5 "Pure Math" Benefits of Formal Proof

So far, I've focuses on the use of formal proof, as a way of communicating the correctness of a proof. Predicting the future value there is murky. However, it clear that mathematics has gained already and will continue to gain by the study of how to express proofs formally.

First, in exploration of the foundations of mathematics. I'm a huge fan of the Curry-Howard isomorphism and I think the experimentation with type theory has produced amazing results. Even with other foundations, building "real math" in a computer makes sure we are not overlooking any necessary steps. (Recall that mathematicians missed necessary pieces of Euclid's geometry for a thousand years and didn't know about the existance the Axiom of Choice until "real math" was studied formally.)

Second, in what I will call "mathematical data structures". There are many equivalent ways to build mathematical concepts like "group" or "graph" out of

simple concepts like "set" or "pair". I find Gonthier's choices for the 4 Color Theorem fascinating - both how he defined planarity and his choice of "hypermaps" as his representation of a graph. I think there is interesting research to be done exploring the best data structures for formal mathematics and their relationship to the usual definition used by mathematicians. My investigations into the corner cases of Delaunay Triangulations made me generate a vastly different definition than is used in most of the literature.

Lastly, I think a library of formal proofs enables data mining and large-scale analytics of mathematics. The "simplicity" of a proof has been proposed as being related to algorithmic information theory and, hence, compression. Having a body of formal proofs allows experimentation and measurement. How much simpler is one proof over another? How are multiple proofs of the same theorem related? Are they somehow all "the same" or are some fundamentally different? Likewise, if there are multiple ways of stating the same theorem or concept, is one inherently simpler or better? Can we classify fields of mathematics by the patterns in the definitions they use? Are two fields of mathematics related and we don't know it because the definitions look dramatically different?

These are all great direct mathematical contributions available to the field. And I'm sure this list isn't exhaustive.

6 Conclusions

Formal proof is a useful technology with a lot of promise, both in the practical world and in mathematics. It offers the highest level of correctness, but, when it cannot be automated, currently has a very high cost. That cost is both in learning the tools and using them. For the short run, because of the high costs, formal proof will be used rarely, only where the highest assurances are needed.

But formal proof will get used. And research will continue. The cost of learning and using the tools will go down. However, even with the price dropping, formal proof faces existing processes in industry and academia.

In industry, only computer chip manufacturers currently use formal proof. Other "safety-critical" industries use other methods for eliminating bugs: static code checks, testing, code reviews, design requirements, etc.. These currently work. To unseat the current process, formal proof has to show itself as dramatically better or substantially cheaper. For the foreseeable future, I think it will continue to be used in chip design but will not migrate to other industries unless a catastrophe occurs under the existing methods. (As Hales mentioned, that may have happened in cryptographic security since I started writing this document!)

In mathematics, formal proof faces the existing process of informal proof with referees. I think formal proof will take up a permanent role in mathematics where the current method does not work. Thus, when then theorem is important and the proofs involving software, or are poorly written, or come from outside academia, formal proof will be used. Because formal proofs require special lengthy training and different techniques, the proofs will be written by

specialists, not regular mathematicians. Formal proofs will be the standard in some subfields of mathematics - especially in new areas and near the foundations. Without a significant cultural change, I don't see everyday mathematicians using formal mathematics.

The pursuit of formal mathematics has and will continue to make contributions to mathematics. In better foundations, in new ways to phrase proofs, and in meta-analysis of formal proofs.

7 Recommendations

If this is a correct view of the future of formal math, what should we be working on to speed things up?

7.1 Code More Than Math

Right now, mathematicians seem okay accepting almost any amount of complexity in an informal proof. But they won't accept sizable computer programs. So, for now, it looks like formalization of code will be more important than "pure" mathematics.

Since more proofs will involve code, it would be good to have a standard mathematical language for code. A standard would help in publication and preservation.

7.2 Code without Garbage Collection

Coq, Isabelle and other tools generate code for functional languages that require a garbage collector. Many safety-critical industries use embedded devices and have real-time requirements that prohibit using a garbage collector. This is significant barrier to adoption.

7.3 Chase Ambulances

Cheap lawyers in America are jokingly called "ambulance chasers", as if they found their clients by following ambulances to the scenes of accidents. It may sound horrible to say, but formal mathematics will be used after an industry or a field of mathematics suffers a catastrophe. An experience like FDIV or like the collapse in confidence surrounding the Italian school of Algebraic Geometry. When controversy strikes, the field needs to pounce on it and show that there is an alternative that provides the highest level of confidence.

Hales has pointed out that the current crisis in security may be an opportunity for the field.

7.4 Social Hacking

If we believe formal proof is a good things for mathematics and we believe the culture of mathematicians prevents its adoption, the answer is to change the

culture of mathematics. This means examining the economic motivations of mathematicians and making them work in your favor.

If a mathematician does a formal proof, give them an award: a line for their C.V. and a plaque for their office. Make them an invited talk at a conference to discuss how they did it. Make it known that if a grad student does a formal proof, they have a guaranteed postdoc somewhere. Celebrate them as pioneers.

If referees are spending years trying to understand a proof, encourage them to reject it and demand a formal proof instead.

7.5 Explore Semi-formal Mathematics

An informal proof is defined as one where it is assumed, with an extraordinary amount of effort, to be convertable into a formal proof. The goal of semi-formal mathematics is to shrink the distance between the formal and informal proof — make the assumed differences smaller and better understood.

Lamport's "How to Write a Proof" is an example, although I think we can do better now. Another example is the detailed informal proof that Hales wrote before others formalized it. One motive would be to have all conclusions and assumptions had to be defined formally, which would allow us to integrate informal proofs with formal.

Machine learning is currently studying how to convert existing informal proofs into formal proofs. That area would have an easier job if the informal proof was highly structured.

7.6 Non-research Tools

Hales has shown that formal math can be used. We need to make the tools "production ready", which means having a stable foundation and features that users need. Many tools are driven by research funding, where a new feature only gets added if someone can write an academic paper about it.

Most notably, both Hales and Gonthier have mentioned that search would be very useful for users, but it's not work that will get someone published. Running a real user study - tracking them and seeing what operations they use and where they spend their time - would be a valuable guide for what features they need most.

We should consider starting a non-for-profit to support non-research tools and to maintain formal proofs.

References

[1] L. Lamport, *How to Write a Proof*, 1993.

[2] The Guardian, *Revealed: how US and UK spy agencies defeat internet privacy and security*, Sept. 5, 2013.