

Where the money is

Jeremy Avigad

Department of Philosophy
Department of Mathematical Sciences
Hoskinson Center for Formal Mathematics

Carnegie Mellon University

November 13, 2023

Formal methods

Formal methods are a body of logic-based methods used in computer science to

- write specifications for hardware, software, protocols, and so on, and
- verify that artifacts meet their specifications.

The same technology is useful for mathematics.

Formal methods in mathematics

Since the early twentieth century, we have known that mathematics can be represented in formal axiomatic systems.

Computational proof assistants allow us to write mathematical definitions, theorems, and proofs in such a way that they can be

- processed,
- verified,
- shared, and
- searched

by mechanical means.

Formal methods in mathematics

The technology is useful:

- for verifying mathematics
- for building communal libraries
- for searching for mathematical results
- for collaborating
- as a gateway to AI
- as a gateway to automated reasoning tools
- for verifying mathematical and scientific software

I will focus on the last two.

Formal methods in mathematics

The image shows a Visual Studio Code editor window with two panes. The left pane displays a Lean 4 proof script for a theorem named `u256_sqrt_soundness`. The right pane shows the corresponding tactic state.

```
lean4 > Verification > LibFuncs > u256 > u256_sqrt_soundness.lean > {} u256_sqrt_soundness
10 theorem auto_sound_u256_sqrt
11   -- arguments
12   (range_check value_low value_high : F)
13   -- code is in memory at s.pc
14   (hmem : MemAt mem u256_sqrt_code o.pc)
15   -- input arguments on the stack
16   ( _ : range_check = mem (o.fp - 5))
17   (hin_value_low : value_low = mem (o.fp - 4))
18   (hin_value_high : value_high = mem (o.fp - 3))
19   -- conclusion
20 : EnsuresRet mem o fun k τ =>
21   ∃ μ ≤ k, RcEnsures mem (rcBound F) μ (mem (o.fp - 5))
22   (mem (τ.ap - 2))
23   (spec_u256_sqrt value_low value_high (mem (τ.ap - 1)))
24 := by
25 apply ensures_of_ensuresb
26 intro vbound
27 -- tempvar sqrt0
28 mkdef hl_sqrt0 : sqrt0 = mem o.ap
29 -- tempvar sqrt1
30 mkdef hl_sqrt1 : sqrt1 = mem (o.ap + 1)
31 -- tempvar remainder_low
32 mkdef hl_remainder_low : remainder_low = mem (o.ap + 2)
33 -- tempvar remainder_high
34 mkdef hl_remainder_high : remainder_high = mem (o.ap + 3)
35 -- tempvar sqrt_mul_2_minus_remainder_ge_u128
36 mkdef hl_sqrt_mul_2_minus_remainder_ge_u128 :
37   sqrt_mul_2_minus_remainder_ge_u128 = mem (o.ap + 4)
38 -- range check for sqrt0
39 step.assert_eq hmem0 hmem with rc_sqrt0
40 -- range check for sqrt1
41 step.assert_eq hmem1 hmem with rc_sqrt1
42 -- tempvar sqrt0_plus_sqrt1
43 step.assert_eq hmem2 hmem with tv_sqrt0_plus_sqrt1
44 mkdef hl_sqrt0_plus_sqrt1 : sqrt0_plus_sqrt1 = sqrt0 + sqrt1
45 have htv_sqrt0_plus_sqrt1 : sqrt0_plus_sqrt1 = mem (o.ap +
46   5) := by
47   rw [hl_sqrt0_plus_sqrt1, hl_sqrt0, hl_sqrt1]
48   exact tv_sqrt0_plus_sqrt1.symm
49 -- tempvar a
```

The right pane shows the tactic state for the goal:

```
Lean infoview
u256_sqrt_soundness.lean:36:41
Tactic state
1 goal
▼ case h.h.h'
F : Type
inst! : Field F
inst! : DecidableEq F
inst! : PreludeHyps F
mem : F → F
o : RegisterState F
range_check value_low value_high : F
hmem : MemAt mem u256_sqrt_code o.pc
x! : range_check = mem (o.fp - 5)
hin_value_low : value_low = mem (o.fp - 4)
hin_value_high : value_high = mem (o.fp - 3)
vbound : N
sqrt0 : F
hl_sqrt0 : sqrt0 = mem o.ap
sqrt1 : F
hl_sqrt1 : sqrt1 = mem (o.ap + 1)
remainder_low : F
hl_remainder_low : remainder_low = mem (o.ap + 2)
remainder_high : F
hl_remainder_high : remainder_high = mem (o.ap + 3)
sqrt_mul_2_minus_remainder_ge_u128 : F
hl_sqrt_mul_2_minus_remainder_ge_u128 : sqrt_mul_2_minus_remainder_
= mem (o.ap + 4)
rc_sqrt0 : mem o.ap = mem (mem (o.fp - 5))
├ Ensuresb vbound mem { pc := o.pc + 1, ap := o.ap + 1, fp :=
o.fp } fun k τ =>
  t.pc = mem (o.fp - 1) ∧
  t.fp = mem (o.fp - 2) ∧
  ∃ μ,
    μ ≤ k + 1 ∧
    RcEnsures mem (rcBound F) μ (mem (o.fp - 5)) (mem
(t.ap - 2))
  (spec_u256_sqrt value_low value_high (mem (τ.ap -
1)))
└ Expected hinc
```

Where the money is

Formal verification is hard and time consuming.

As the story goes, when William Sutton was asked why he robbed banks, he replied, “because that’s where the money is.”

The only reason to verify certain pieces of mathematical software and certain mathematical results is because that’s where the problems are.

Where the money is

We don't always need to verify everything down to axiomatic primitives.

Sometimes even just having a formal *specification* can be useful:

- It makes claims and models precise.
- It enables us to combine results from different sources.

We can then pick and choose which parts need to be verified.

Where the money is

I will discuss some examples where verification can make a difference.

- convex optimization
- applications of SAT solvers
- blockchain applications
- computational chemistry

Convex optimization

Numerical and symbolic methods for optimization are used extensively in engineering, industry, and finance.

Various methods are used to reduce problems to convex optimization problems.

1. Start with a problem you want to solve.
2. Reduce it to a convex optimization problem.
3. Transform it to the format required by a solver.
4. Solve it.
5. Check a certificate of correctness.
6. Report the results.

Convex optimization

Numerical and symbolic methods for optimization are used extensively in engineering, industry, and finance.

Various methods are used to reduce problems to convex optimization problems.

1. Start with a problem you want to solve.
2. Reduce it to a convex optimization problem.
3. Transform it to the format required by a solver.
4. Solve it.
5. Check a certificate of correctness.
6. Report the results.

Convex optimization

Sometimes sophisticated mathematical arguments are needed to reduce a problem. There are a variety of tricks and techniques to make problems convex.

Alexander Bentkamp, Ramon Fernández Mir, and I developed a prototype system in Lean:

- It provides a framework for defining problems and carrying out reductions.
- It has an extensible library of convex functions and their properties.
- It is situated within Lean's mathematical framework.
- It carries out “DCP transformations” automatically, and verifies them.

For now, we trust the floating point computations.

Applications of SAT solvers

Propositional satisfiability solvers are used in a variety of applications, including:

- hardware and software verification
- planning problems and AI
- cryptography
- proving combinatorial theorems.

Method:

- Start with the problem you want to solve.
- Reduce it to a SAT problem (or similar).
- Call the solver.
- Check the result (often an UNSAT proof).
- Report the result.

Applications of SAT solvers

Propositional satisfiability solvers are used in a variety of applications, including:

- hardware and software verification
- planning problems and AI
- cryptography
- proving combinatorial theorems.

Method:

- Start with the problem you want to solve.
- Reduce it to a SAT problem (or similar).
- Call the solver.
- Check the result (often an UNSAT proof).
- Report the result.

Applications of SAT solvers

Clever encodings and tricks are used to represent problems so that they can be solved efficiently.

Cayden Codel, James Gallicchio, Wojciech Nawrocki, Marijn Heule, and I are developing a Lean library to:

- carry out reductions and verify their correctness; and
- verify the results of SAT solvers and related tools.

Blockchain applications

Despite turmoil in decentralized finance, blockchain technology is here to stay. Billions of dollars are lost each year due to bugs in smart contracts.

StarkWare Industries provides a “layer two” solution:

- Write programs in a programming language, Cairo, such that successful termination guarantees a claim.
- Compile it to machine code.
- Encode the claim that the machine code runs to completion as the existence of solutions to a parametric family of polynomials.
- Publish on blockchain a short cryptographic certificate that guarantees the existence of the solutions.

Blockchain applications

Despite turmoil in decentralized finance, blockchain technology is here to stay. Billions of dollars are lost each year due to bugs in smart contracts.

StarkWare Industries provides a “layer two” solution:

- Write programs in a programming language, Cairo, such that successful termination guarantees a claim.
- Compile it to machine code.
- Encode the claim that the machine code runs to completion as the existence of solutions to a parametric family of polynomials.
- Publish on blockchain a short cryptographic certificate that guarantees the existence of the solutions.

Blockchain applications

With colleagues at StarkWare, we have:

- verified the encoding of termination claims as polynomials;
- written a proof-producing compiler from Cairo to machine code;
- verified components of the library, such as elliptic curve signature validation, down to the cryptographic certificates.

At this stage, we trust the cryptographic protocol, though others have worked on verifying those too.

Computational chemistry

I have learned from Tyler Josephson that Monte Carlo methods are fundamental to computational chemistry. Thousands of papers on simulations are published each year.

Method:

- Start with a model of a physical system and a quantity you want to compute.
- Design a Monte Carlo algorithm to compute it.
- Implement it, and run it.
- Report the results.

Computational chemistry

I have learned from Tyler Josephson that Monte Carlo methods are fundamental to computational chemistry. Thousands of papers on simulations are published each year.

Method:

- Start with a model of a physical system and a quantity you want to compute.
- Design a Monte Carlo algorithm to compute it.
- Implement it, and run it.
- Report the results.

Computational chemistry

Researchers use tricks and heuristics to make a Markov chain Monte Carlo (MCMC) algorithm converge quickly.

A *detailed balance* condition ensures that the samples converge to the correct distribution. Violations lead to discrepancies in the literature.

Tyler and I proposed to:

- Verify detailed balance conditions in Lean.
- Implement and verify probabilistic programs to sample from the relevant distributions.

Conclusions

Morals:

- Talk to people outside formal methods. There are important verification problems everywhere.
- Be flexible. Choose your battles.
- Be creative.
- Be proud. Verification matters.