



Provably Trustworthy Systems

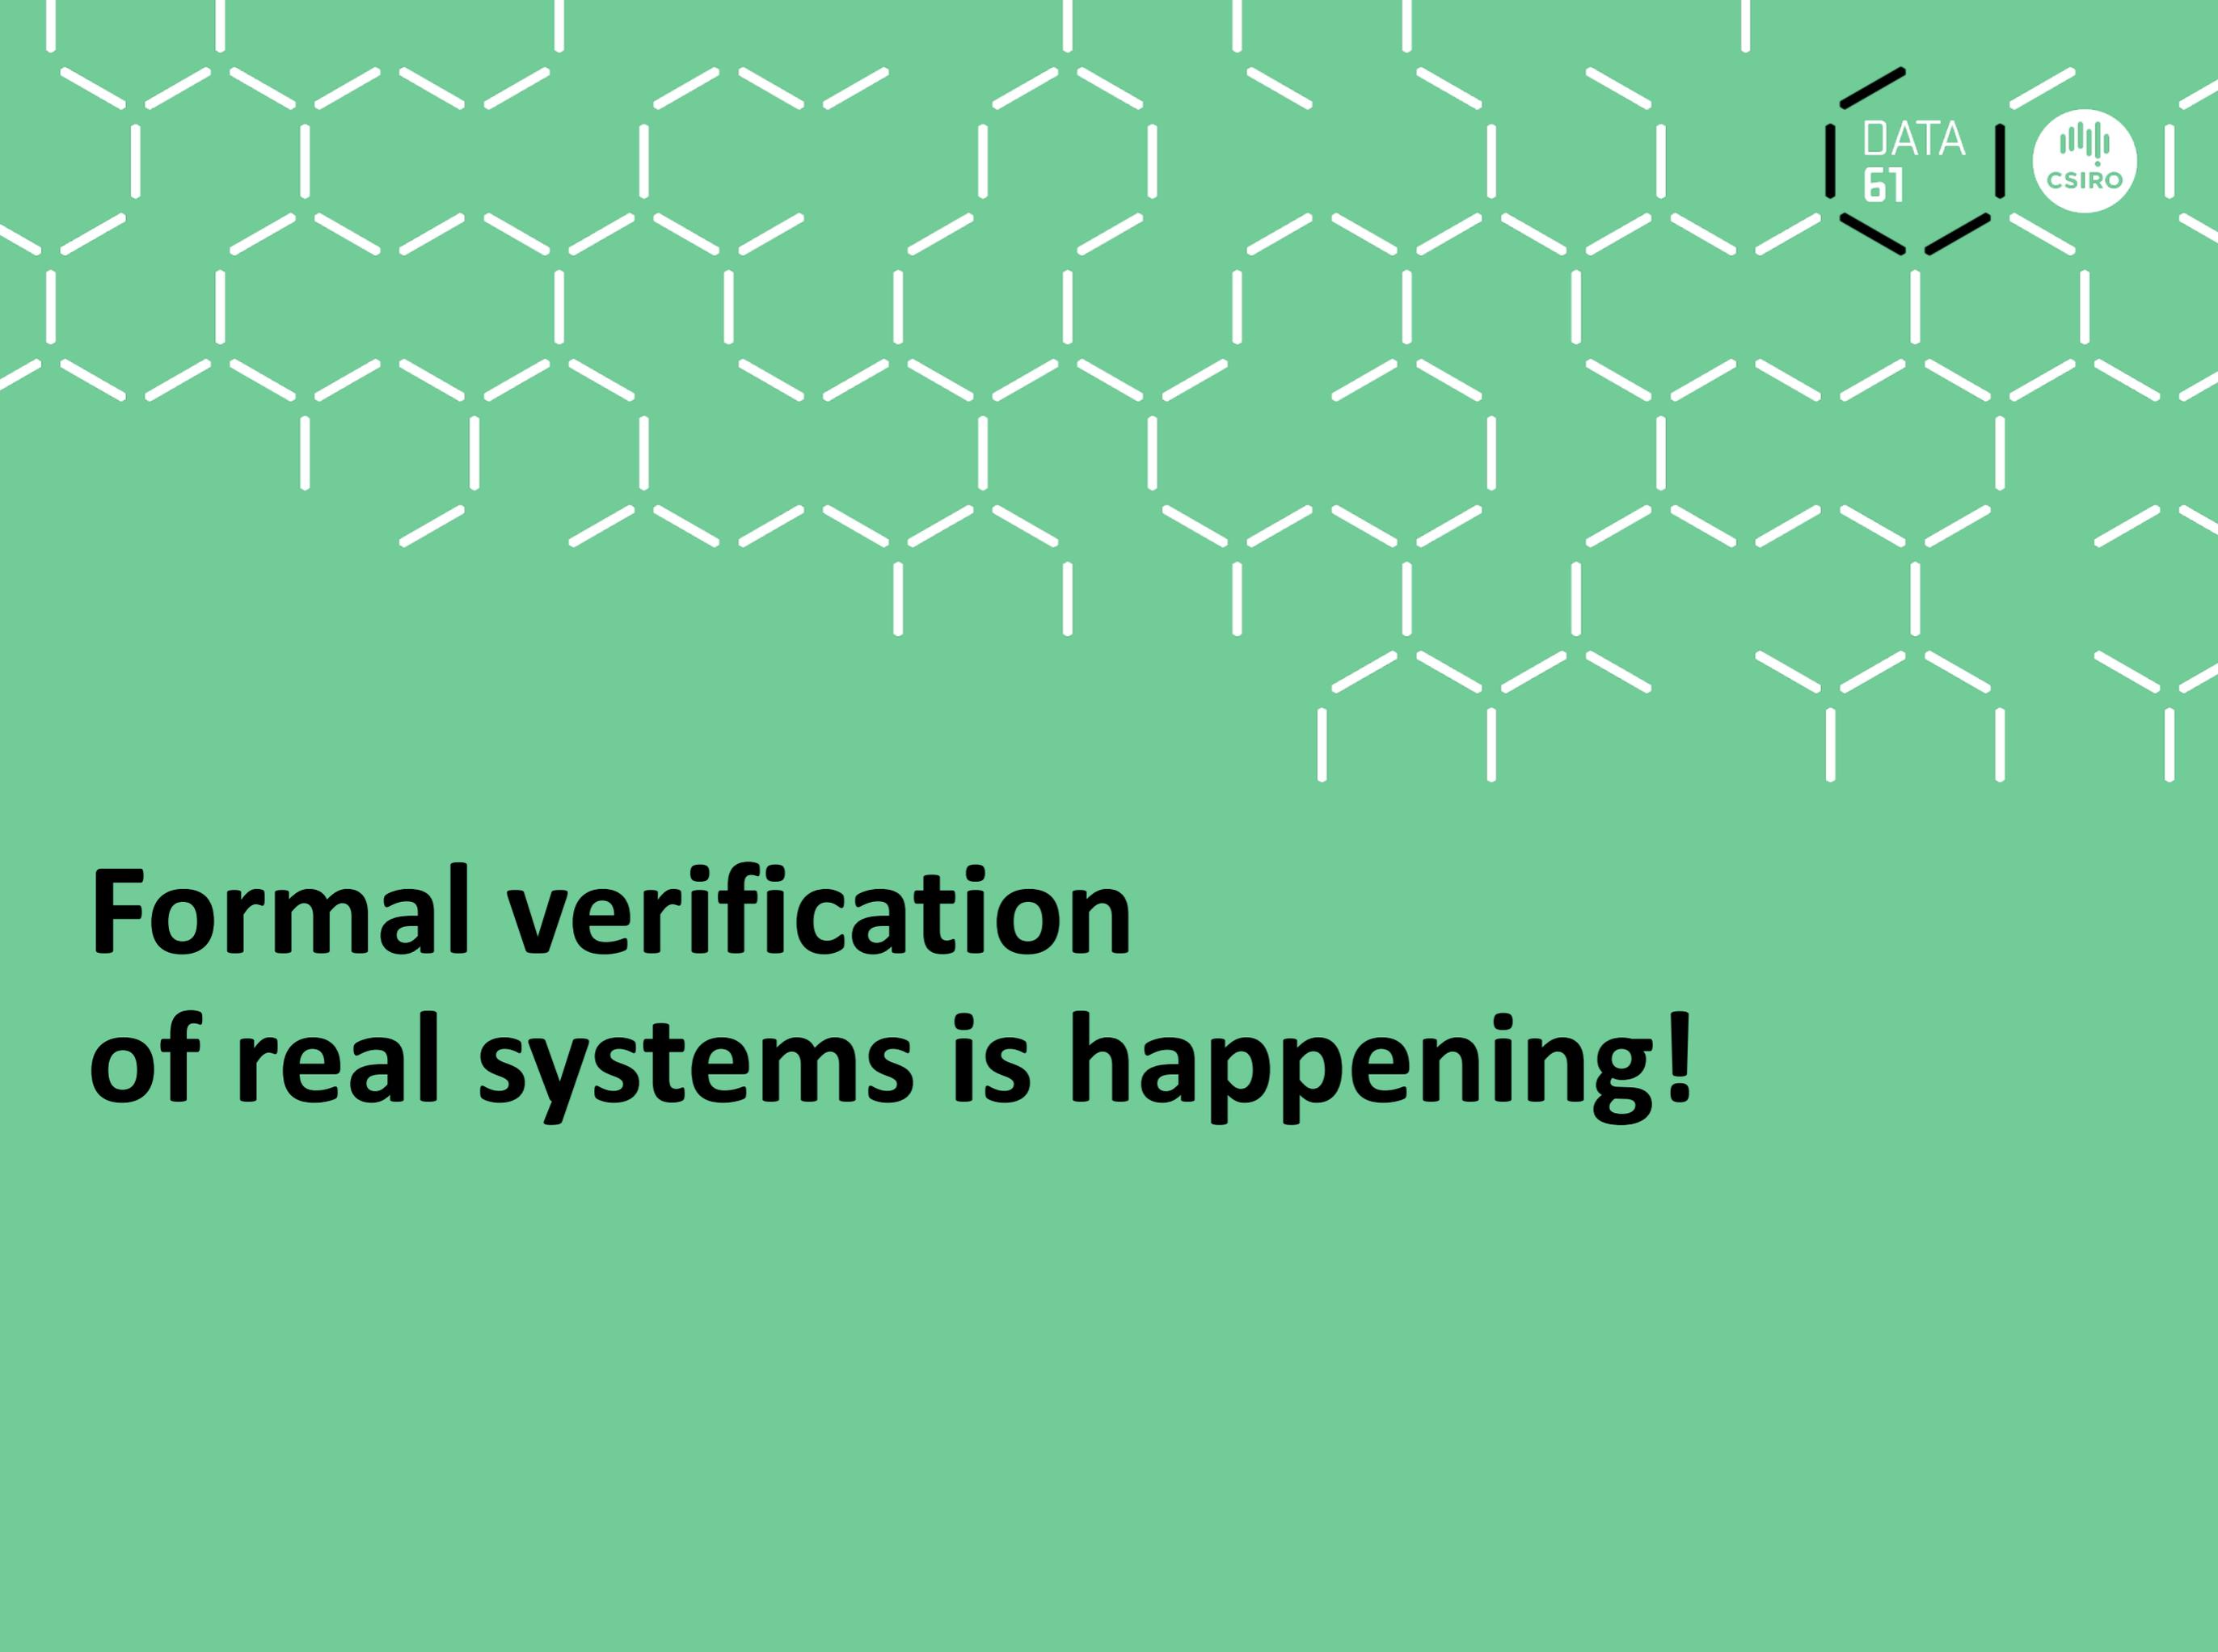
seL4 and beyond

Gerwin Klein

Royal Society Meeting on Verified trustworthy software systems
April 2016

data61.csiro.au





DATA
61



**Formal verification
of real systems is happening!**

Formal verification of real systems



- ▶ Increasingly many examples:

Formal verification of real systems



► Increasingly many examples:

- seL4
 - verified OS kernel implementation



Formal verification of real systems



► Increasingly many examples:

- **CompCert**

- verified compiler implementation



- **seL4**

- verified OS kernel implementation



Security. Performance. Proof.

Formal verification of real systems



► Increasingly many examples:

- **CompCert**

- verified compiler implementation



- **seL4**

- verified OS kernel implementation



- **Ironfleet and Ironclad**

- verified distributed system



Formal verification of real systems



► Increasingly many examples:

- **CompCert**

- verified compiler implementation



- **seL4**

- verified OS kernel implementation



- **Ironfleet and Ironclad**



- **CakeML**



CAKEML

A Verified Implementation of ML

Formal verification of real systems



► Increasingly many examples:

- **CompCert**

- verified compiler implementation



- **seL4**

- verified OS kernel implementation



- **Ironfleet and Ironclad**



- **CakeML**



- **Candle**

- verified interactive HOL theorem prover implementation

Formal verification of real systems



► Increasingly many examples:

- **CompCert**

- verified compiler implementation



- **seL4**

- verified OS kernel implementation



- **Ironfleet and Ironclad**



- **PolarSSL**

- verified SSL implementation



- **CAKEML**

- A Verified Implementation of ML

- **le**

- verified interactive HOL theorem prover implementation

Formal verification of real systems



► Increasingly many examples:

- **CompCert**

- verified compiler implementation



- **seL4**

- verified OS kernel implementation



- **Ironfleet and Ironclad**



- **PolarSSL**

- verified



- **CoCon**

- verified conference system



- HOL theorem prover

Formal verification of real systems



► Increasingly many examples:

- **CompCert**

- verified compiler implementation



- **seL4**

- verified OS kernel implementation



- **Ironfleet and Ironclad**



- **PolarSSL**

- verified



- **CoCon**

- verified conference system



- HOL theorem prover

Formal verification of real systems



► Increasingly many examples:

- **CompCert**

- verified compiler implementation



- **seL4**

- verified OS kernel implementation



- **Ironfleet and Ironclad**



- **PolarSSL**

- verified



- **CoCon**

- verified conference system



- ve HOL theorem prover

- **OpenSSL HMAC**

- verified crypto implementation



Formal verification of real systems



► Increasingly many examples:

- **CompCert**

- verified compiler implementation



- **seL4**

- verified OS kernel implementation



- **Ironfleet and Ironclad**



- **PolarSSL**

- verified



- **CoCon**



- ve HOL theorem prover

- **FSCQ**

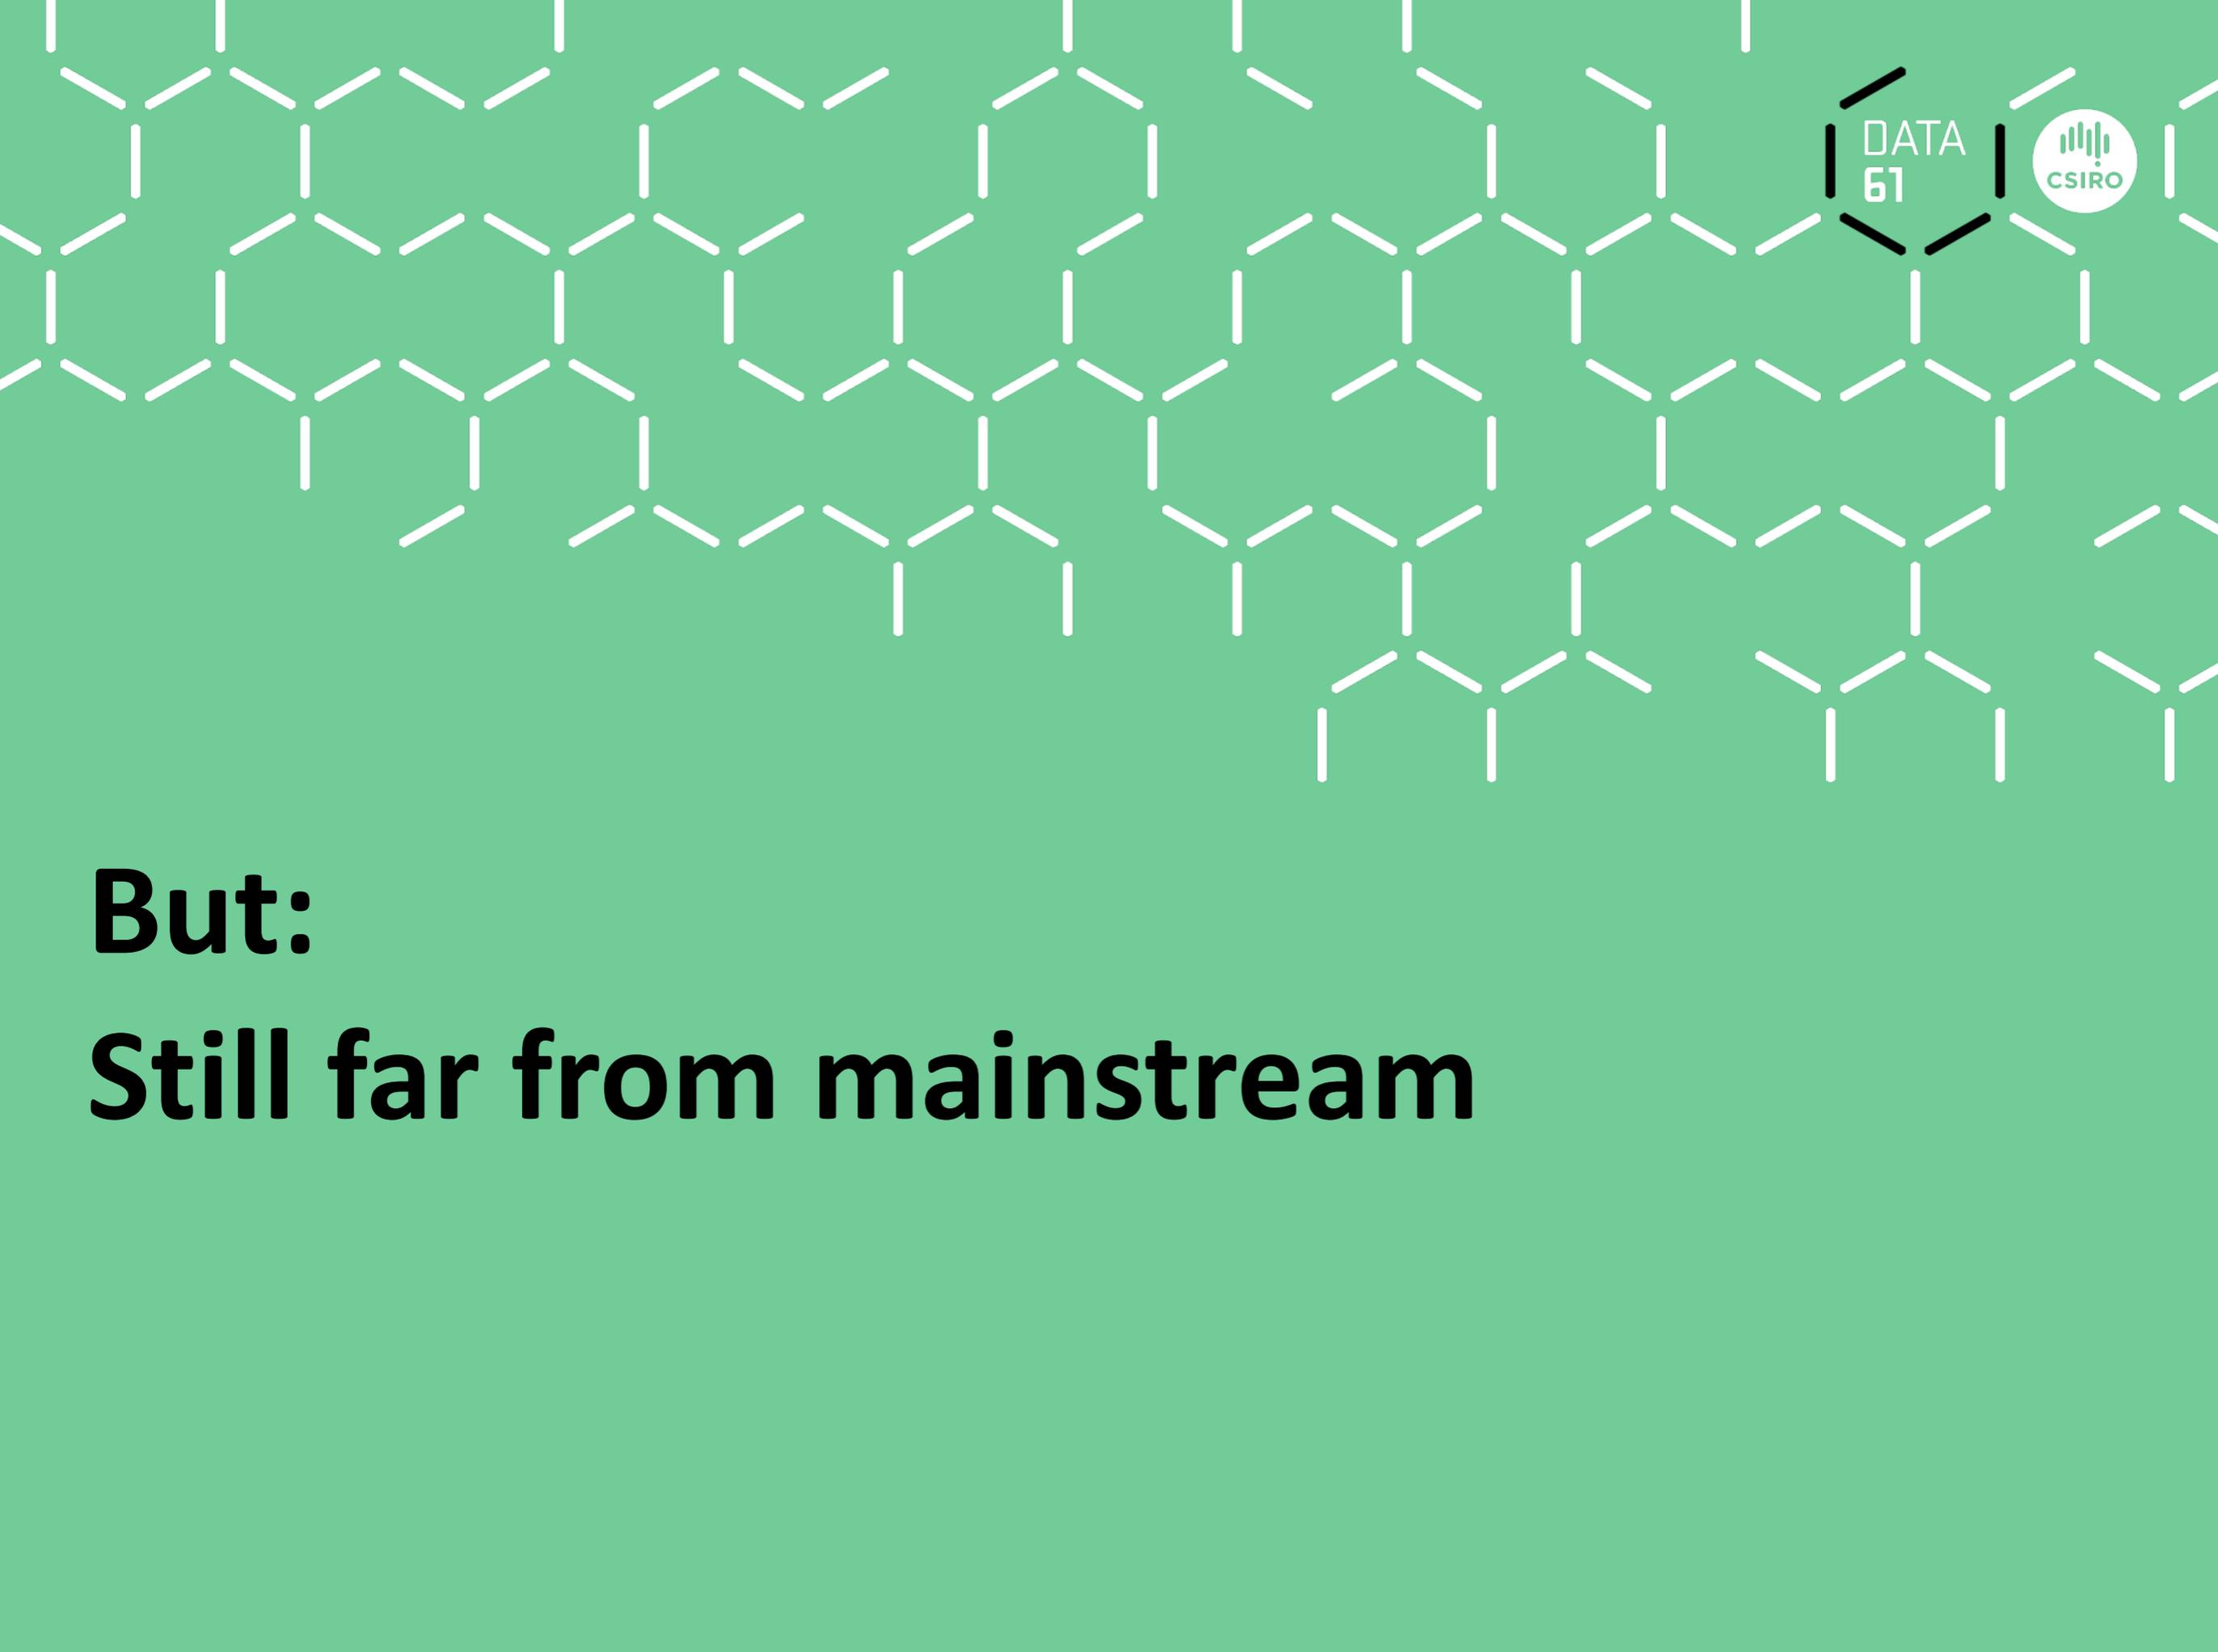
- verified crash resistance file system



- **OpenSSL HMAC**

- verified crypto implementation





DATA
61



But:

Still far from mainstream

Too Expensive



▶ Such projects are still big research results

- Often break new ground
- Multiple person years or person decades
- Real, binary-level results still rare
- Hard to maintain over long periods

Too Expensive



▶ Such projects are still big research results

- Often break new ground
- Multiple person years or person decades
- Real, binary-level results still rare
- Hard to maintain over long periods

▶ Still too expensive

- But not that far off:
 - cheaper than traditional high-assurance dev
 - factor 2-3 over high-quality traditional embedded systems dev

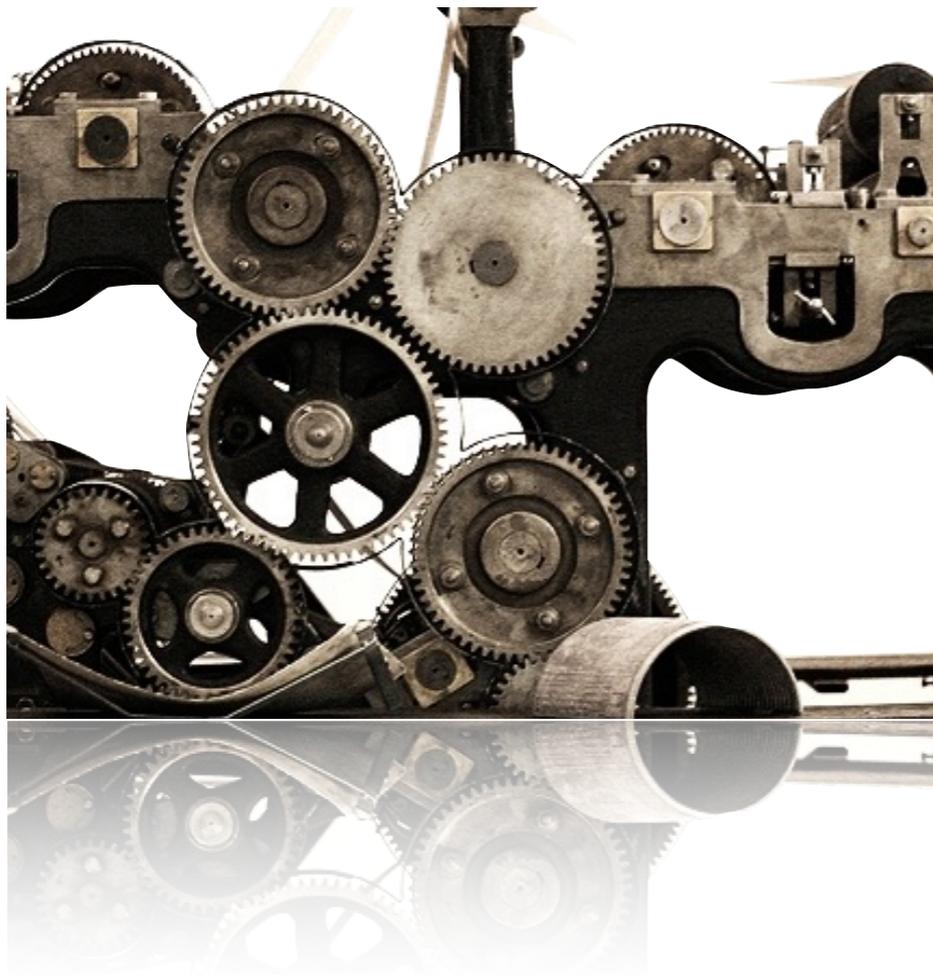
DATA
61



What can be done?

Better, cheaper, faster.

- ▶ Just needs to be **cheaper**:
 - economic pressure wins over time
 - everything else follows



Better, cheaper, faster.

▶ Just needs to be **cheaper**:

- economic pressure wins over time
- everything else follows



▶ **Proof Productivity**:

- Tools
 - more automation, deeper automation, **built for scale**
- Proof Engineering
 - predictability, estimation, **scale**
- Languages
 - design for verification, **increase verification productivity**
- ...

The rest of this talk



▶ seL4

▶ Scale

▶ Proof Engineering

▶ Proof Effort

▶ Future

DATA
61



seL4

seL4: Isolation



Trustworthy Computing Base

- message passing
- virtual memory
- interrupt handling
- access control

Applications

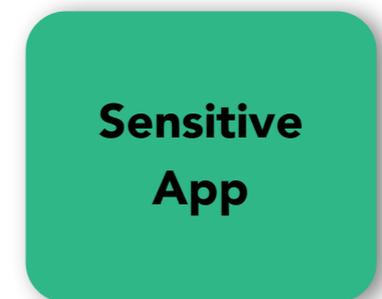
- fault isolation
- fault identification
- IP protection
- modularity

Trusted next to Untrusted

Untrusted



Trusted



seL4: Isolation



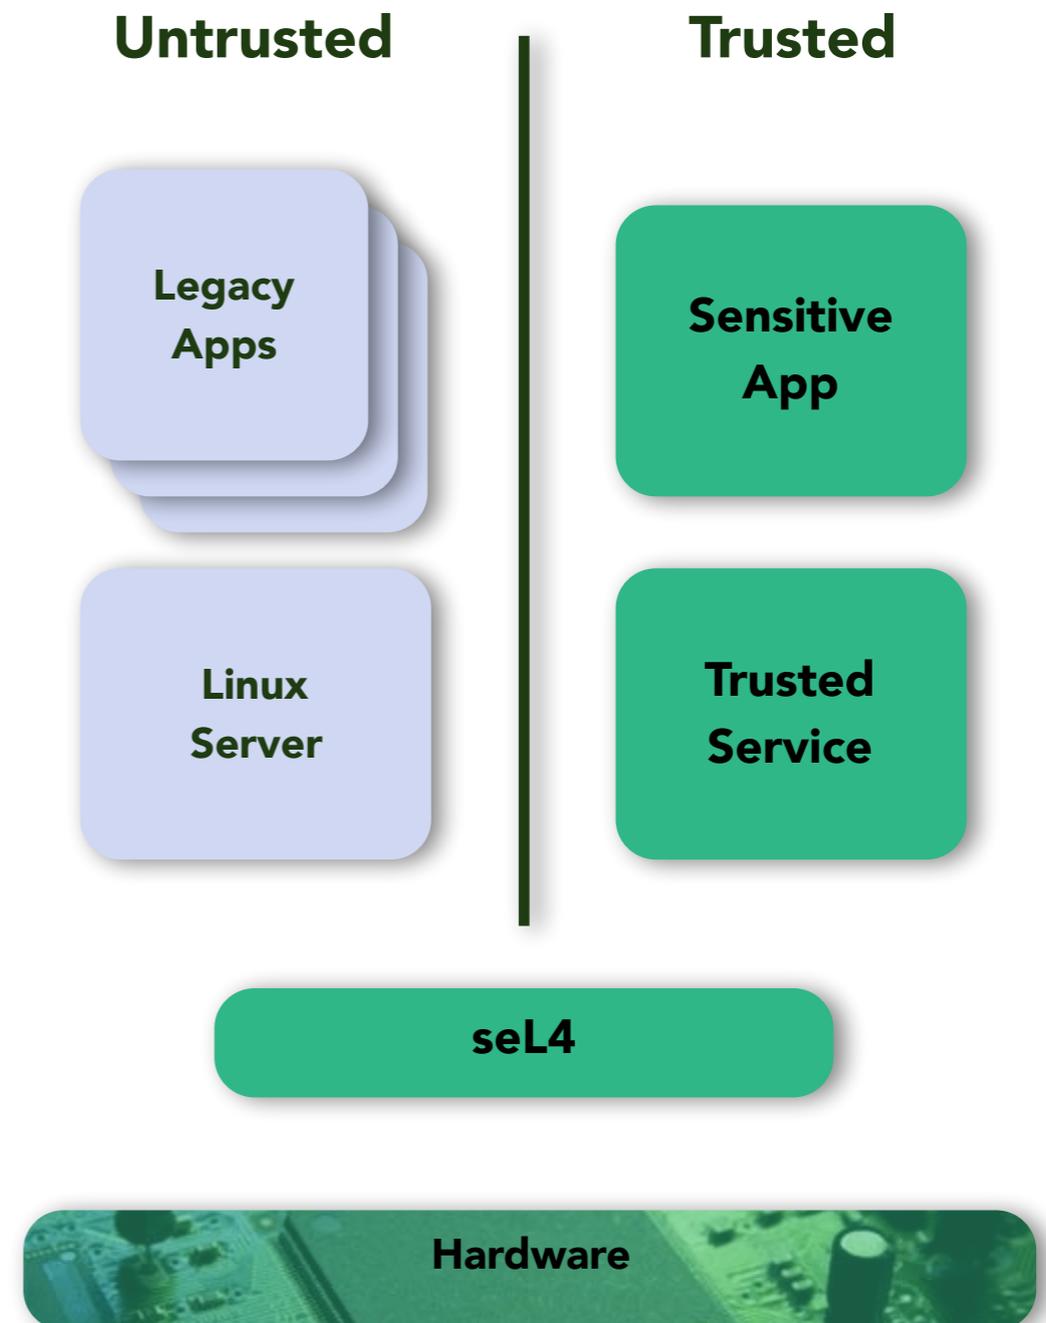
Trustworthy Computing Base

- message passing
- virtual memory
- interrupt handling
- access control

Applications

- fault isolation
- fault identification
- IP protection
- modularity

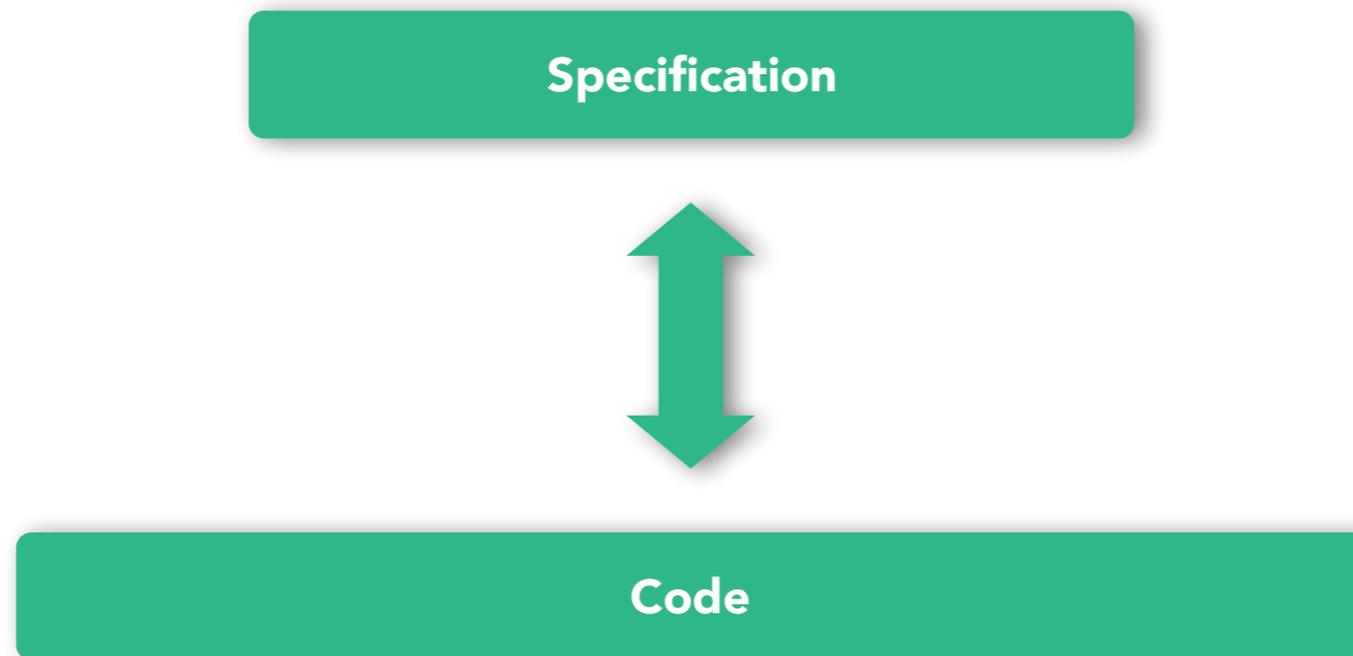
Trusted next to Untrusted



Functional Correctness



Proof



Functional Correctness

What

definition

```
schedule :: unit s_monad where
schedule ≡ do
  threads ← allActiveTCBs;
  thread ← select threads;
  switch_to_thread thread
od
OR switch_to_idle_thread
```

Specification

Proof



Code

Functional Correctness



What

Specification

definition

```
schedule :: unit s_monad where
schedule ≡ do
  threads ← allActiveTCBs;
  thread ← select threads;
  switch_to_thread thread
od
OR switch_to_idle_thread
```

Proof

How

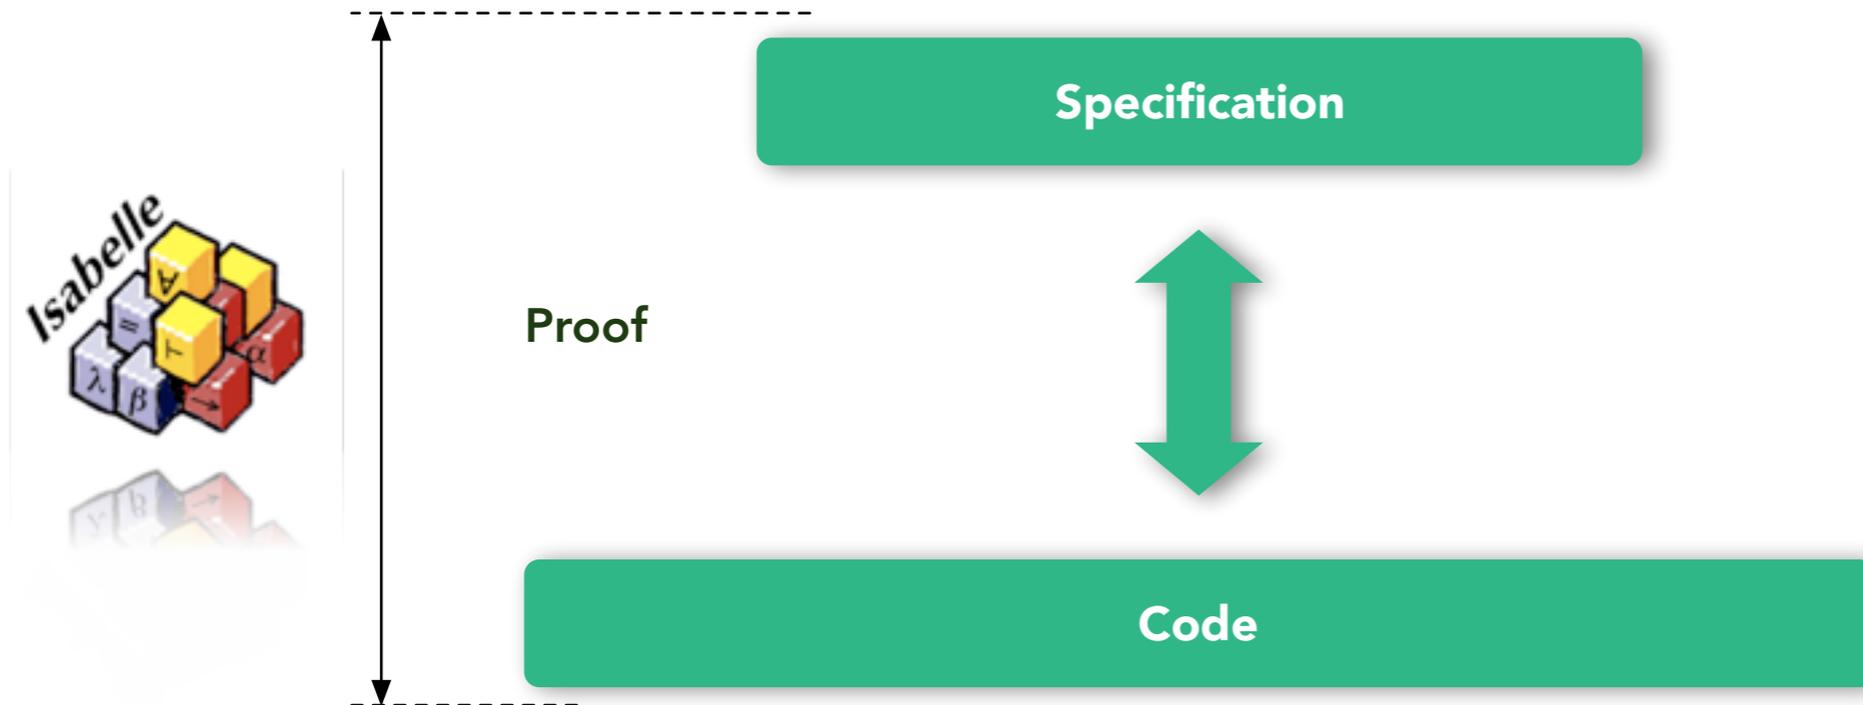
```
void
schedule(void) {
  switch ((word_t)ksSchedulerAction) {
    case (word_t)SchedulerAction_ResumeCurrentThread:
      break;

    case (word_t)SchedulerAction_ChooseNewThread:
      chooseThread();
      ksSchedulerAction = SchedulerAction_ResumeCurrentThread;
      break;

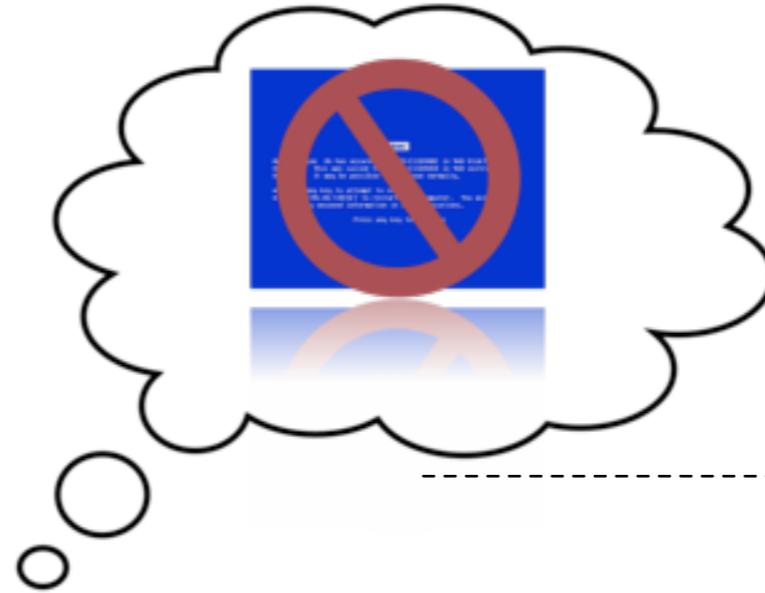
    default: /* SwitchToThread */
      switchToThread(ksSchedulerAction);
      ksSchedulerAction = SchedulerAction_ResumeCurrentThread;
      break;
  }
}

void
chooseThread(void) {
  prio_t prio;
  tcb_t *thread, *next;
```

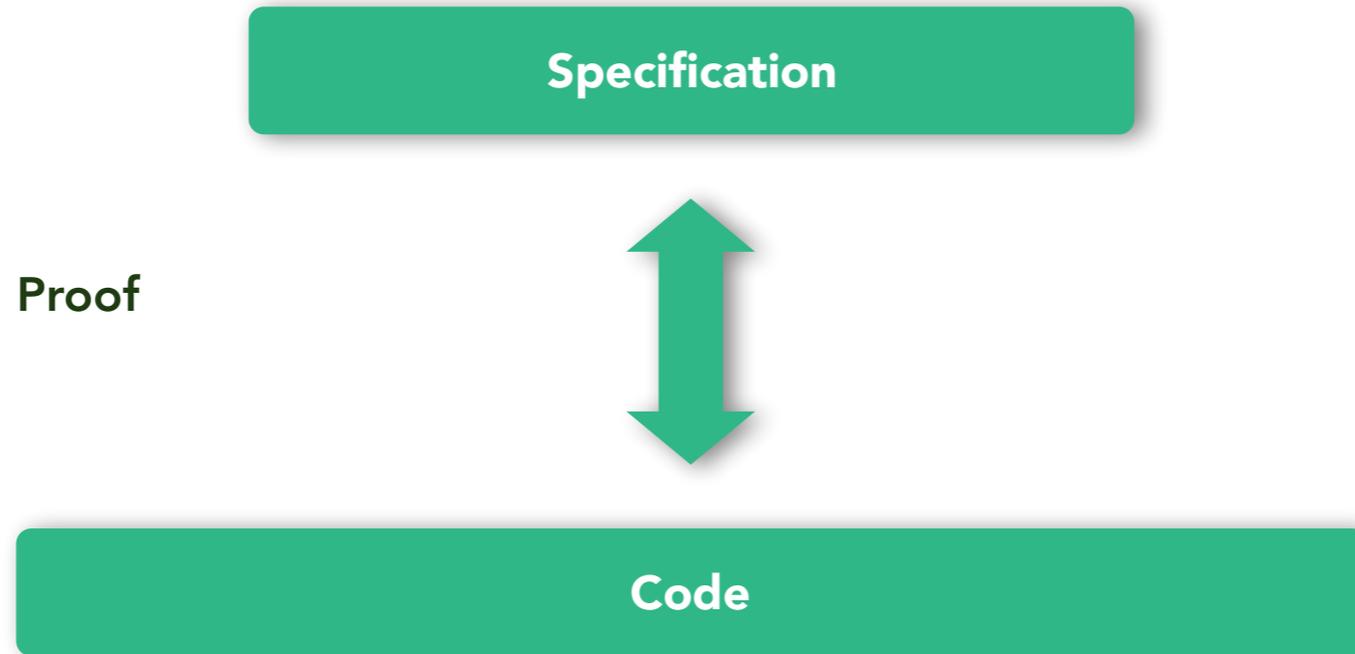
*conditions apply



***conditions apply**



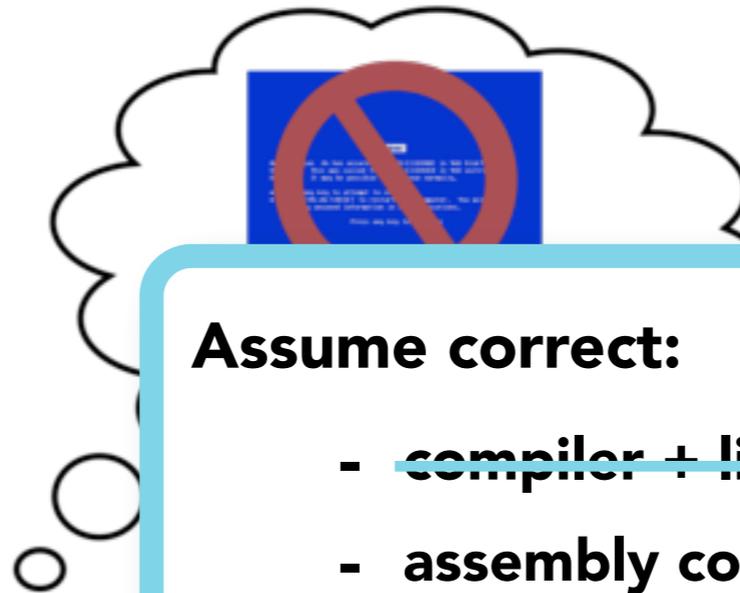
Expectation



Assumptions

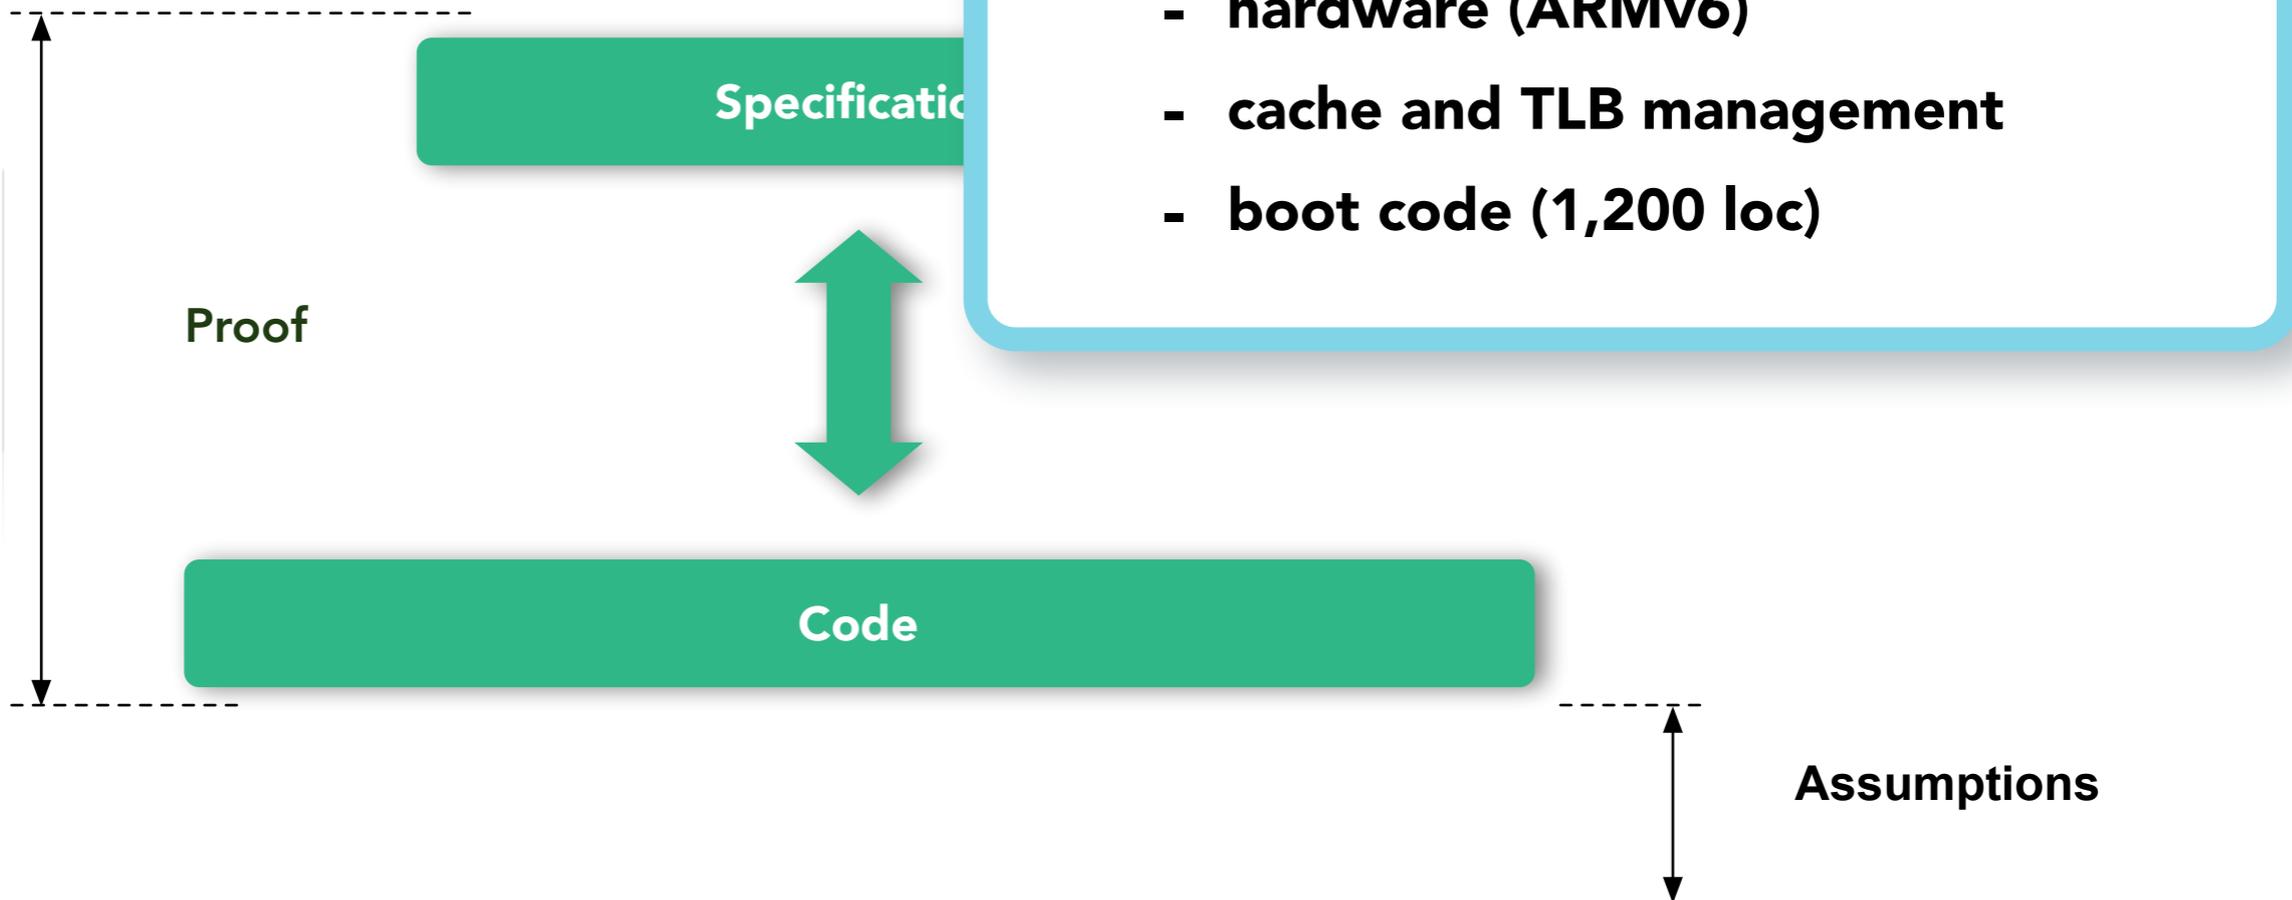


***conditions apply**

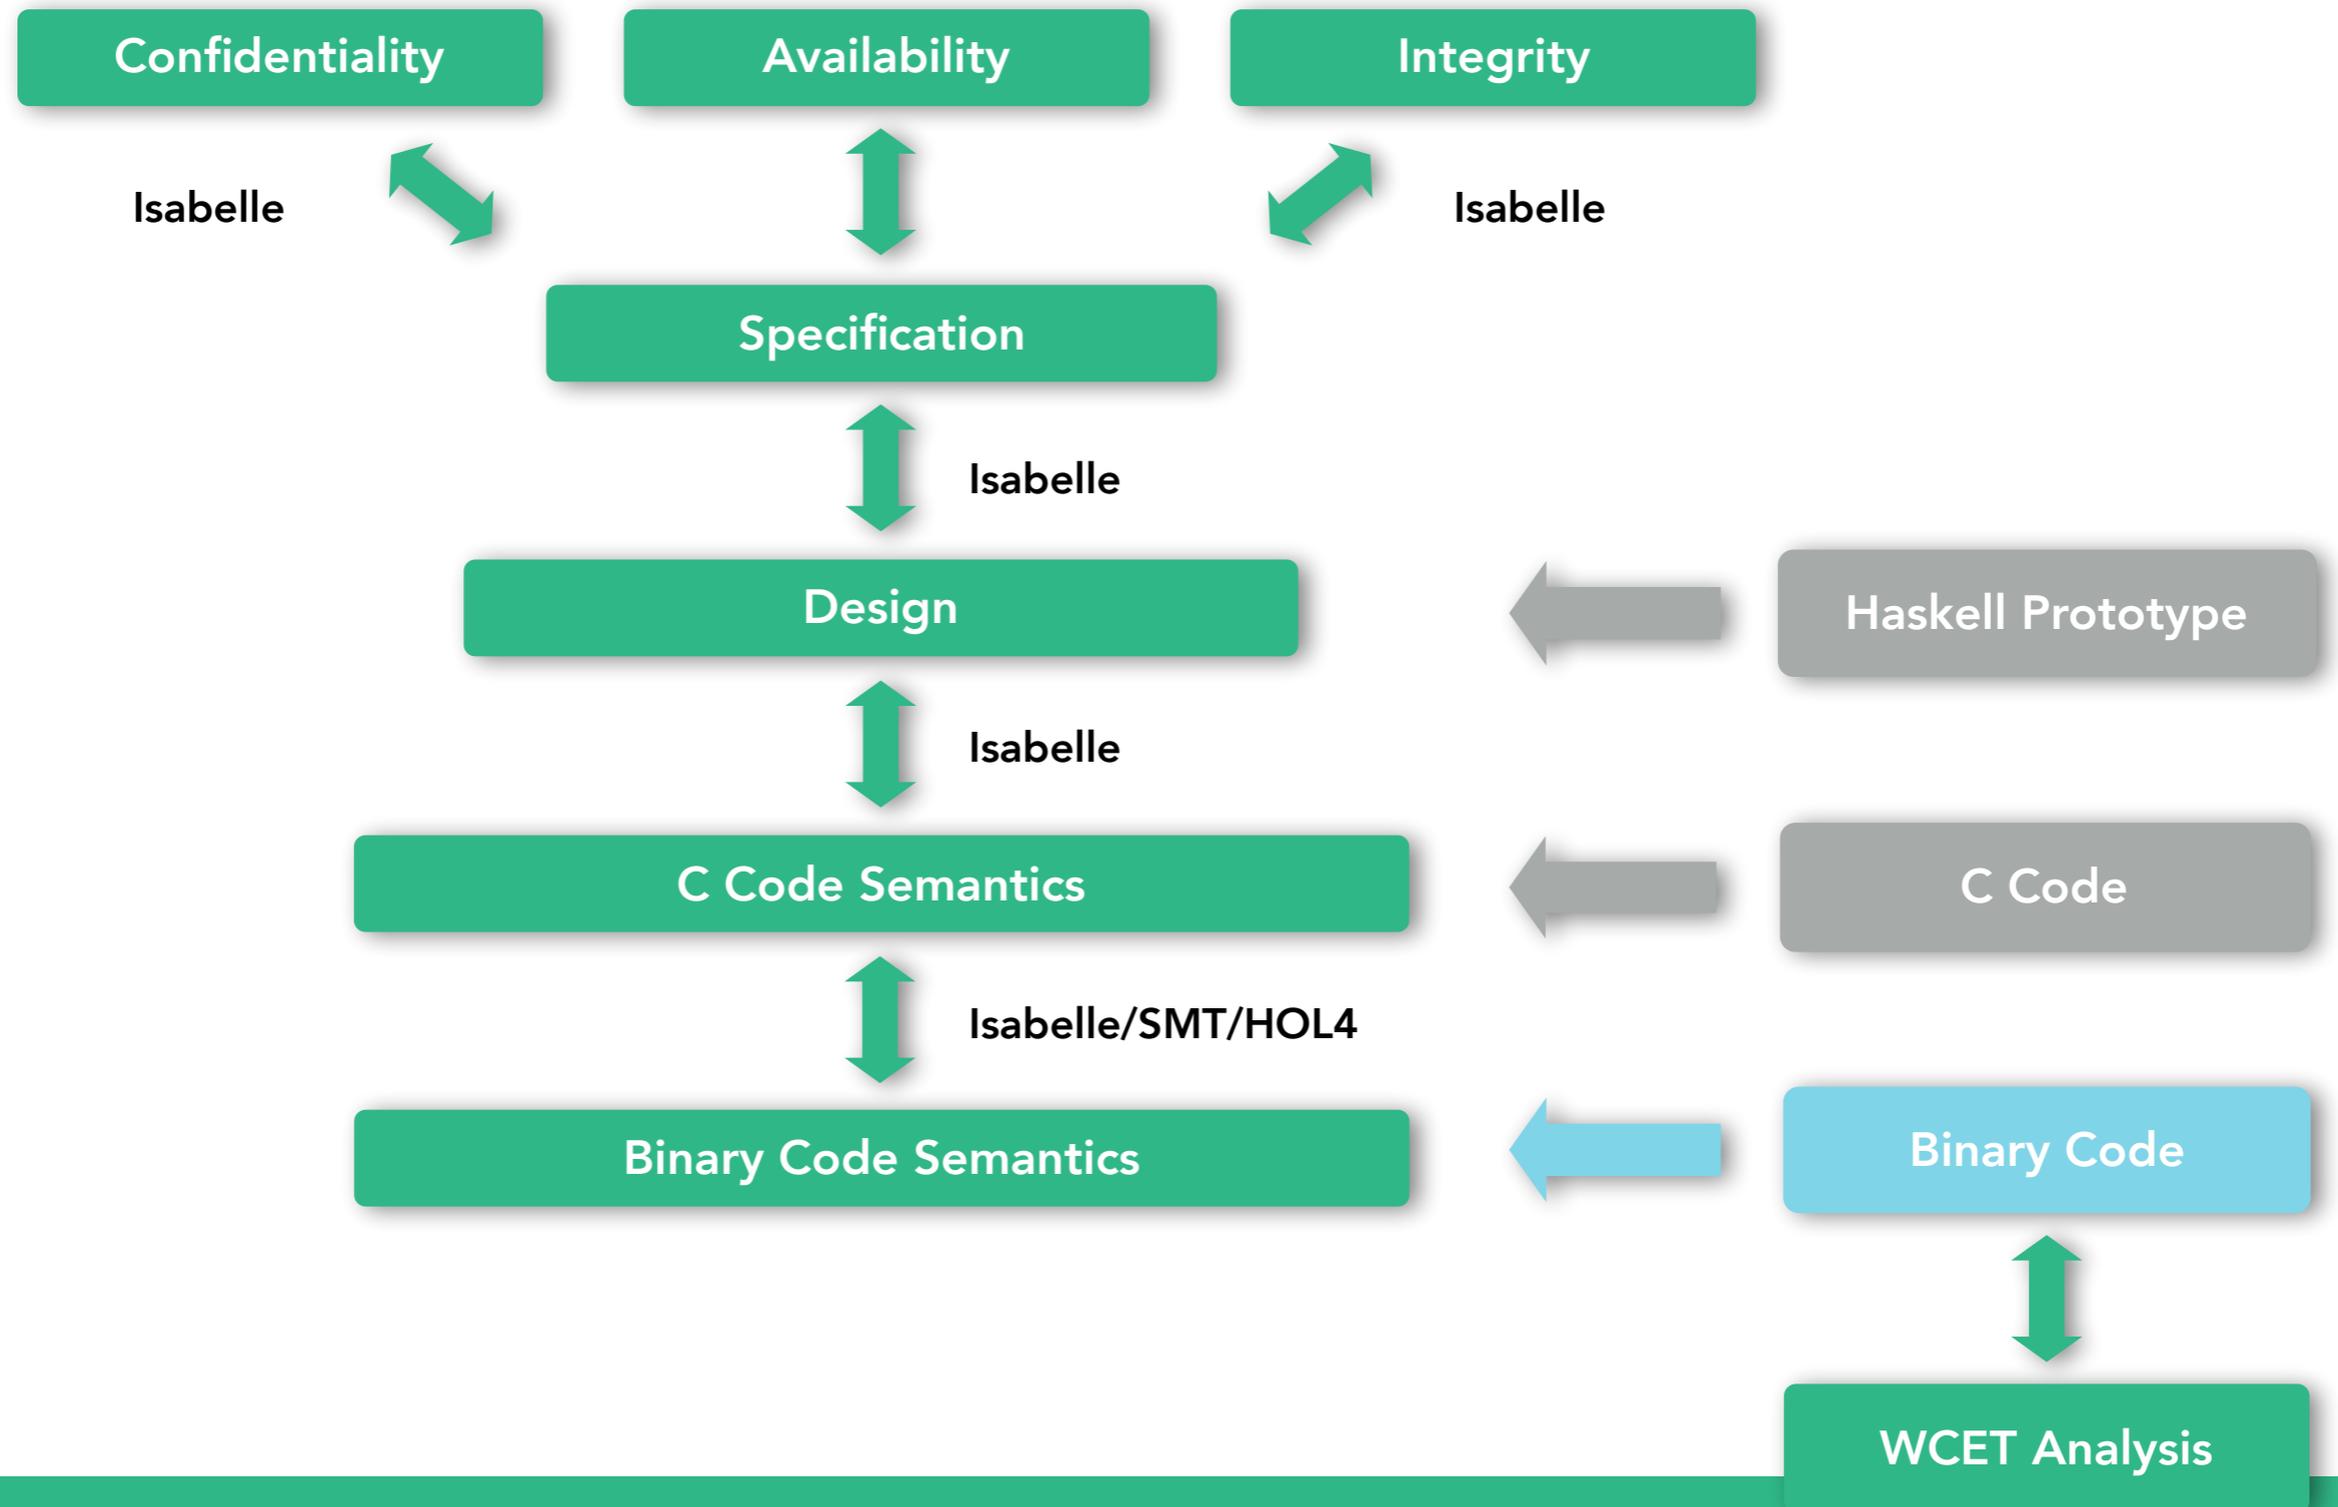


Assume correct:

- ~~compiler + linker (wrt. C op sem)~~
- assembly code (600 loc)
- hardware (ARMv6)
- cache and TLB management
- boot code (1,200 loc)



Proof Architecture Now



Proof Architecture Now

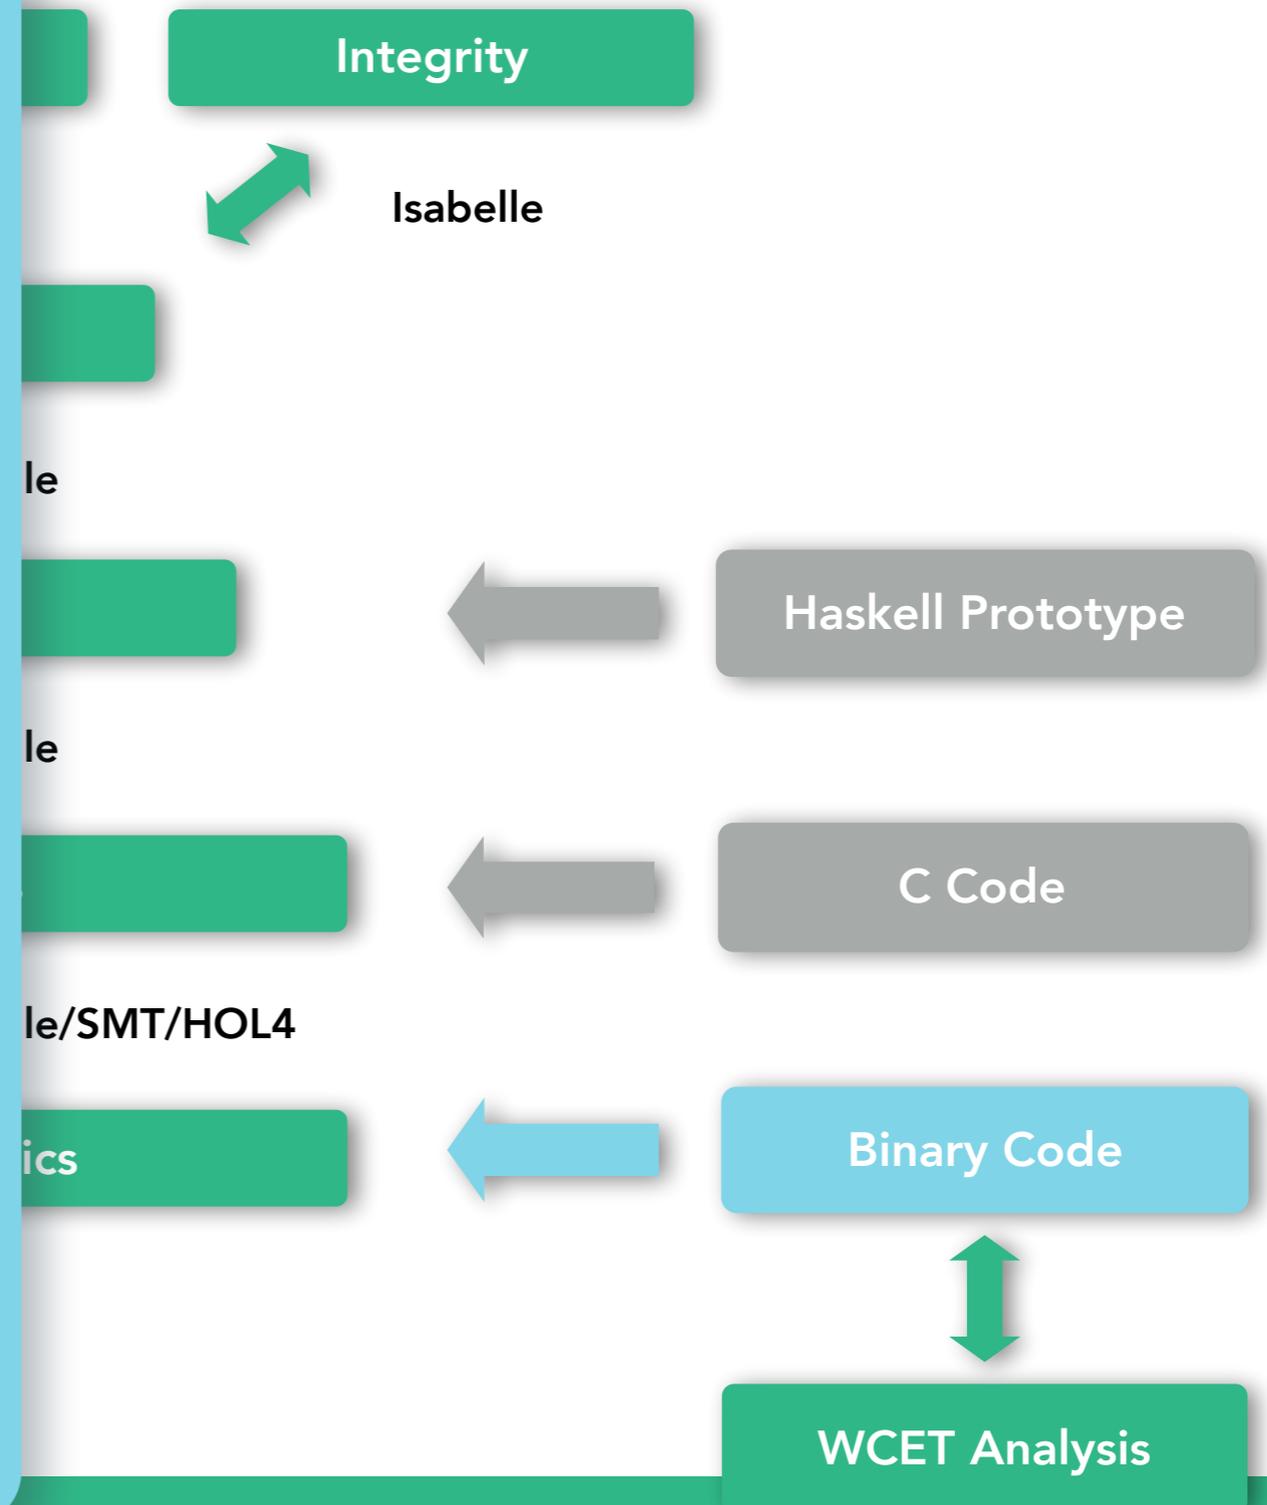


High-level properties:

- functional correctness
- integrity
- authority confinement
- non-interference
- termination
- user-level system initialisation
- verified component platform
- worst-case execution time (by static analysis)

Roadmap:

- verified x64 version
- virtualisation extensions
- mixed-criticality real-time
- timing side-channel elimination



Proof Architecture Now



High-level properties:

- functional correctness
- integrity
- authority confinement
- non-interference
- termination
- user-level system initialisation
- verified component platform
- worst-case execution time (by static analysis)

Roadmap:

- verified x64 version
- virtualisation extensions
- mixed-criticality real-time
- timing side-channel elimination

Integrity

Open Source

<http://seL4.systems>
<https://github.com/seL4/>

le/SMT/HOL4

ics

Binary Code

WCET Analysis

As Real as it Gets

► Autonomous in



As Real as it Gets

- ▶ Autonomous in 3, 2, 1..

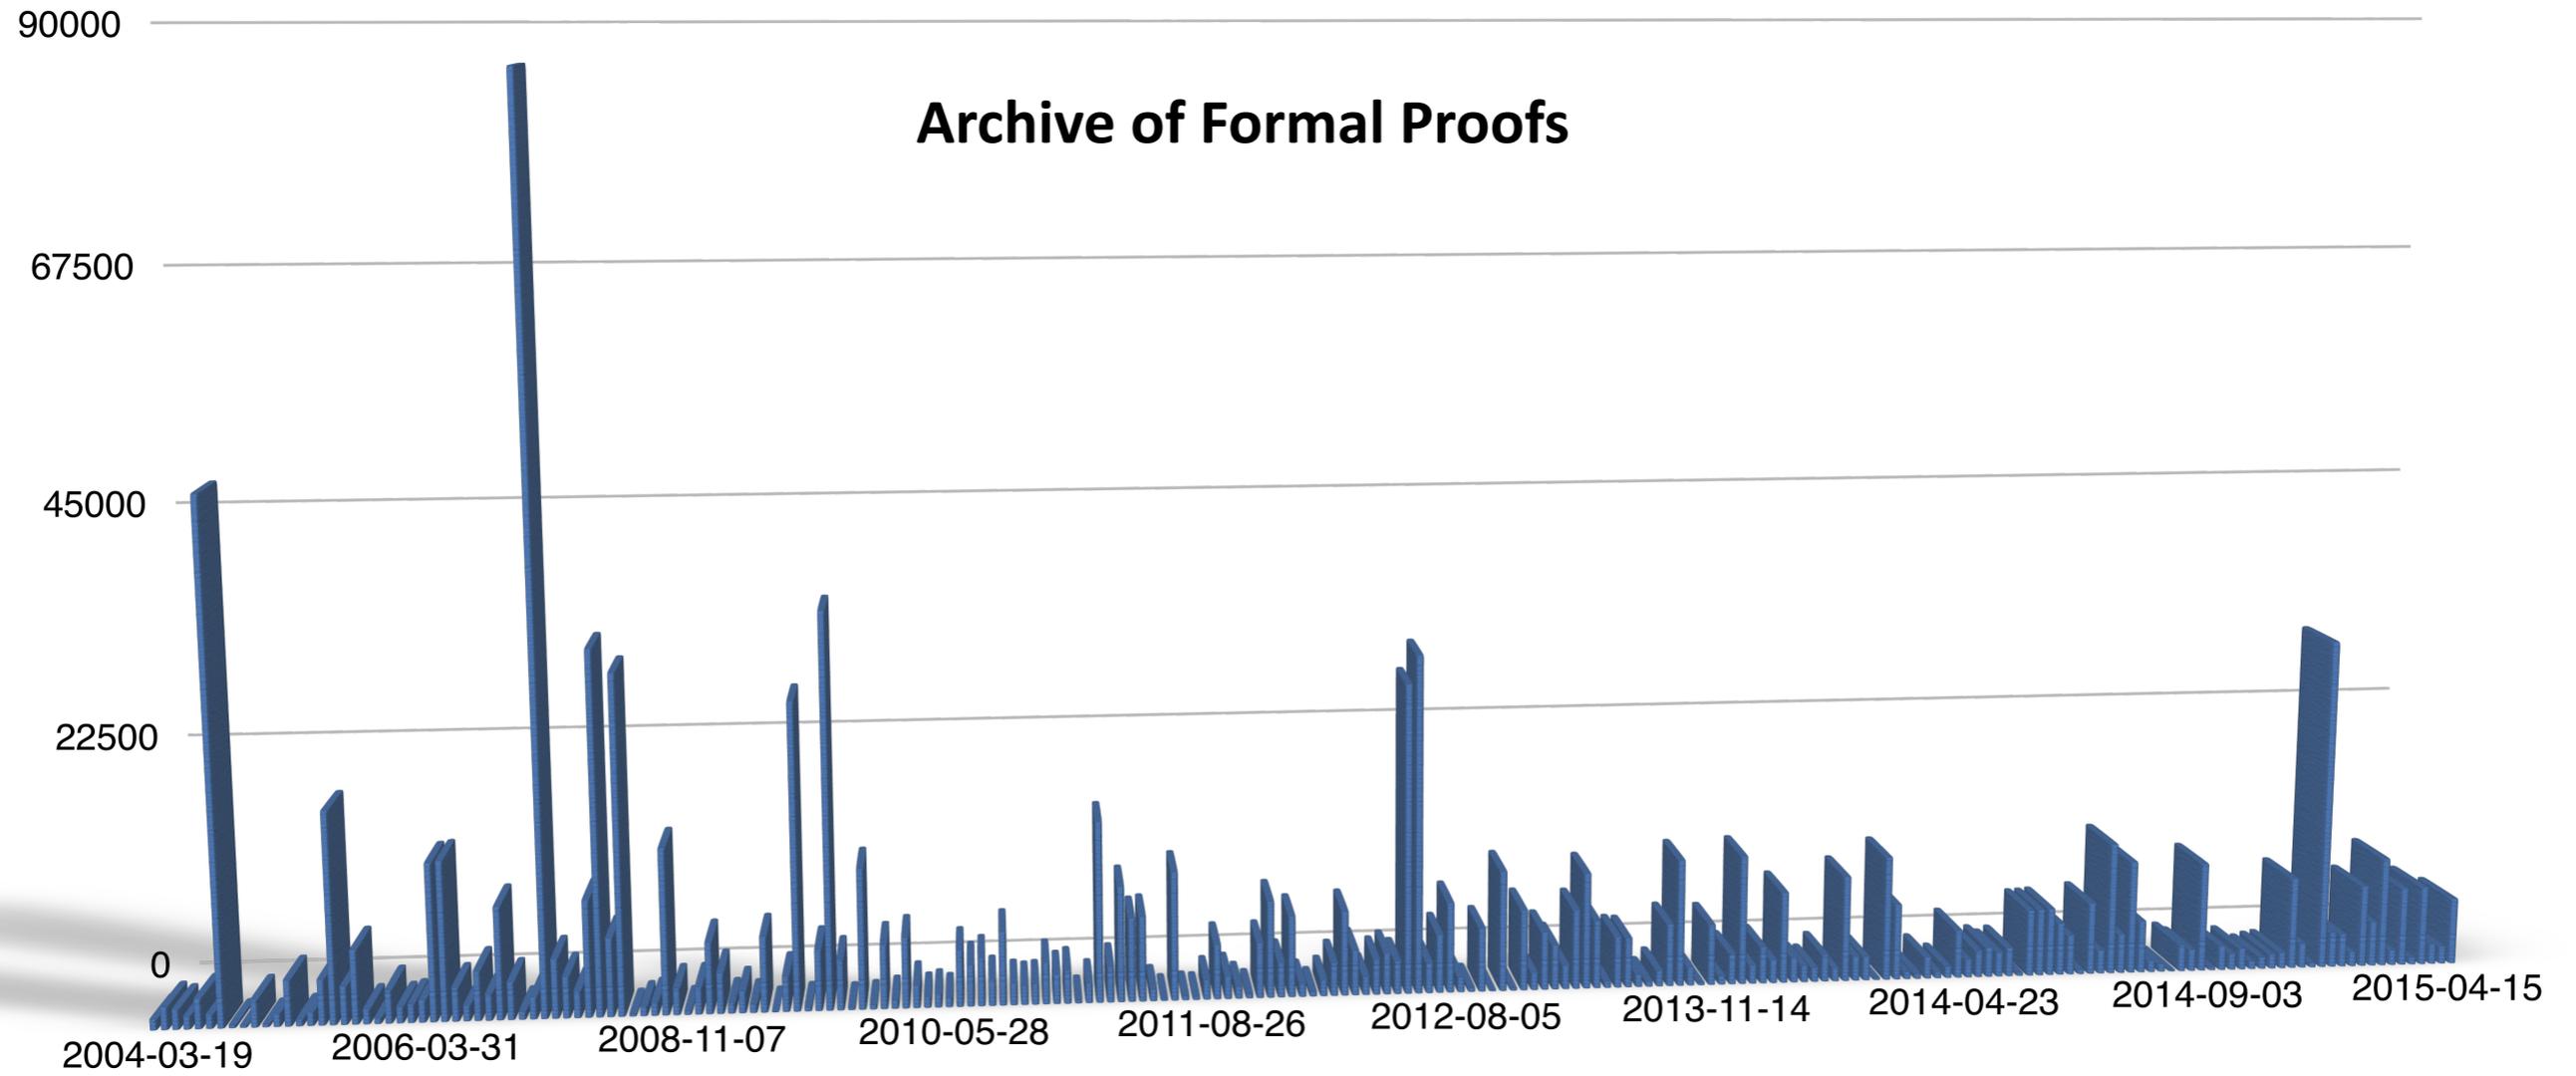


DATA
61



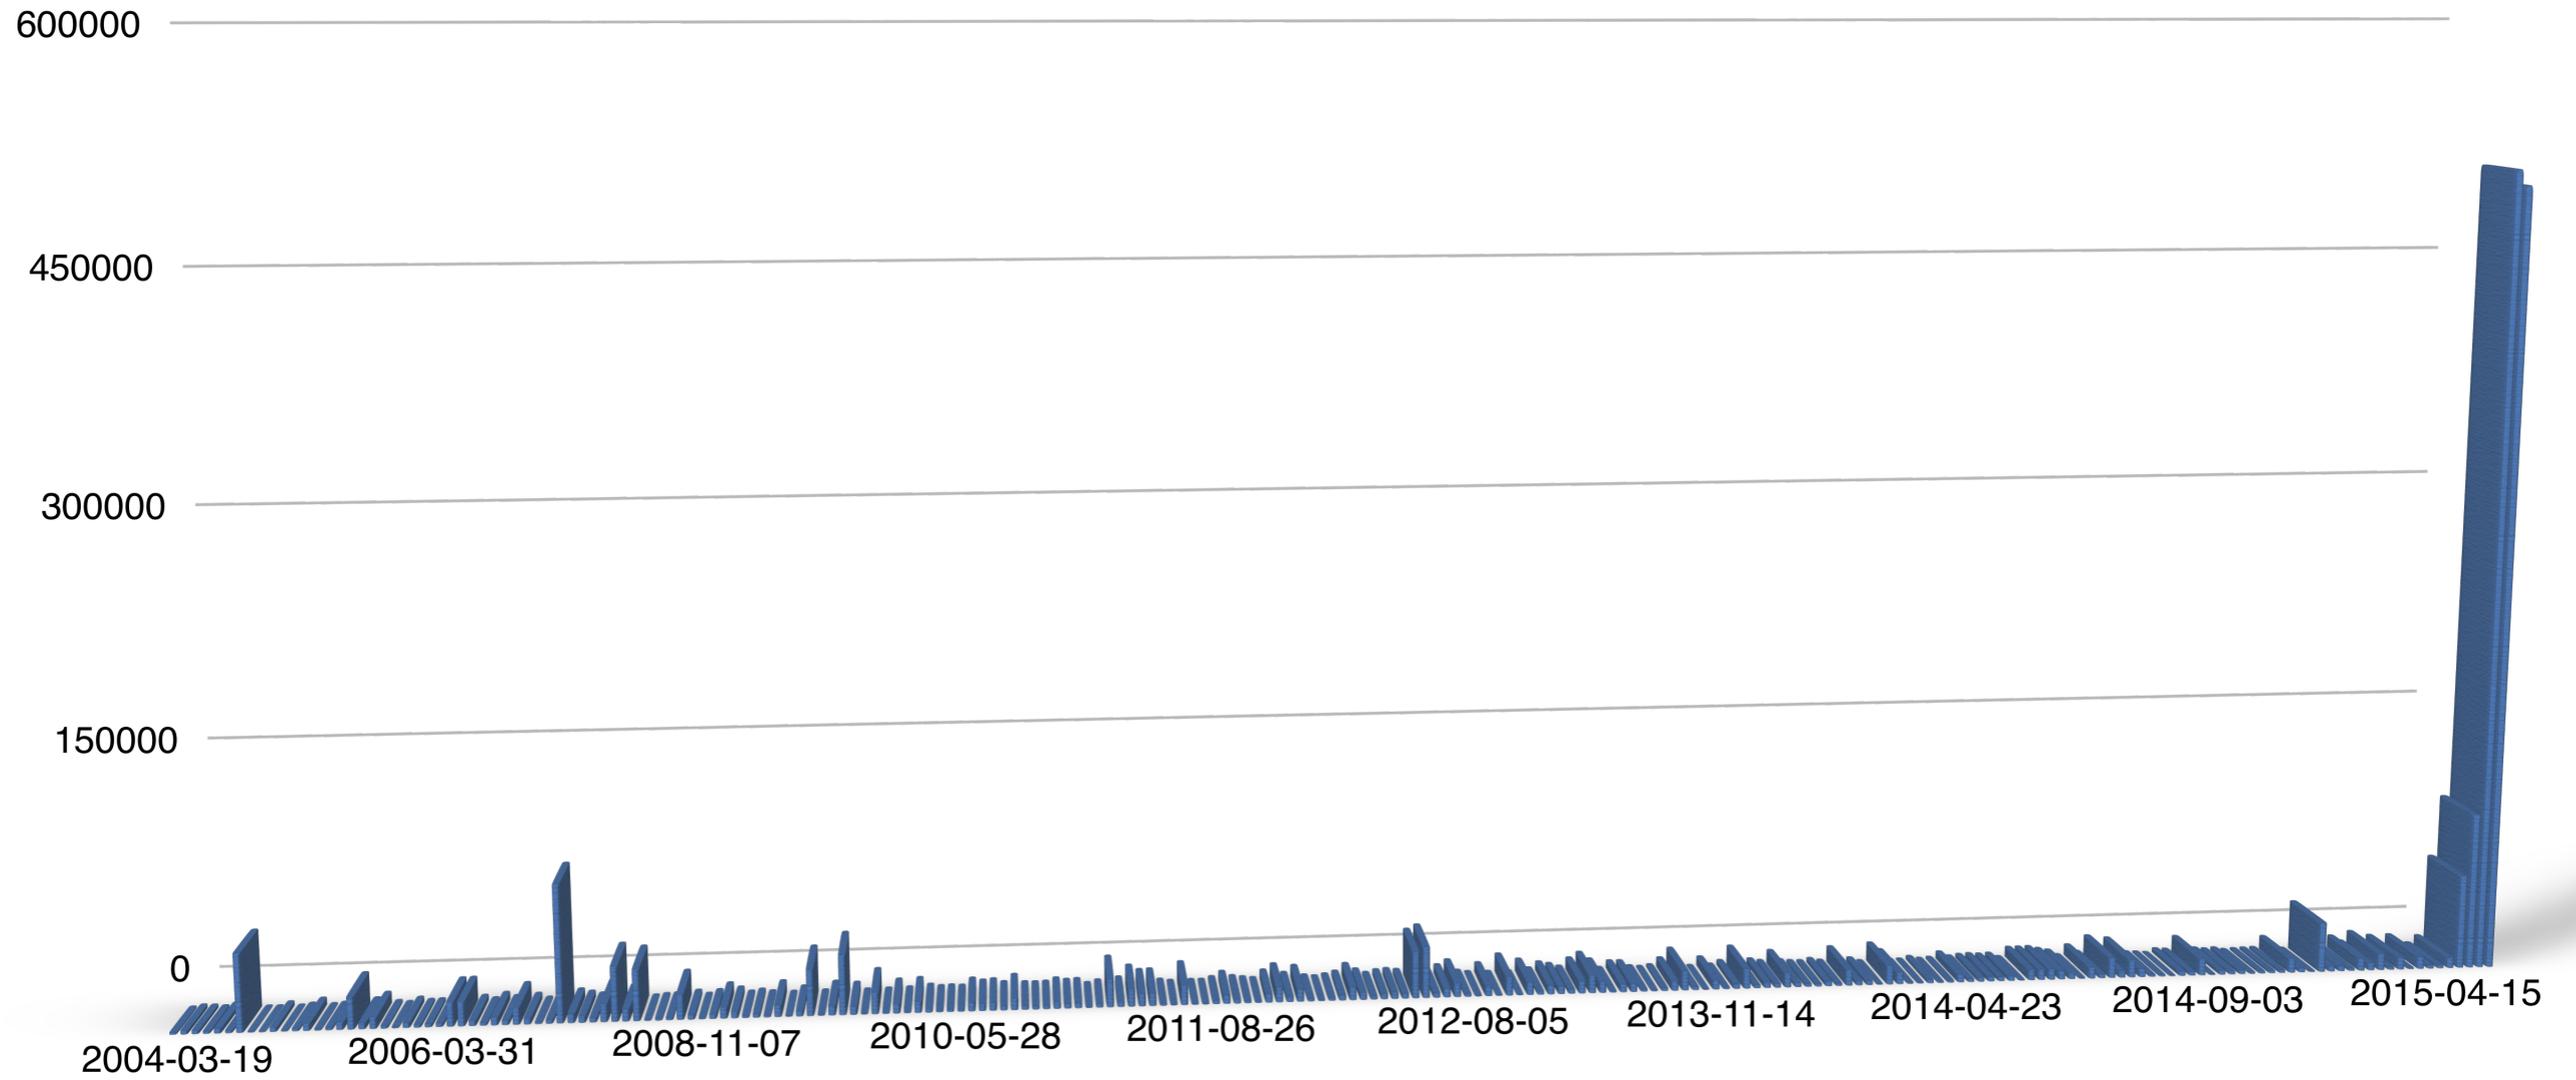
Scale

Scale



size of AFP entries by submission date

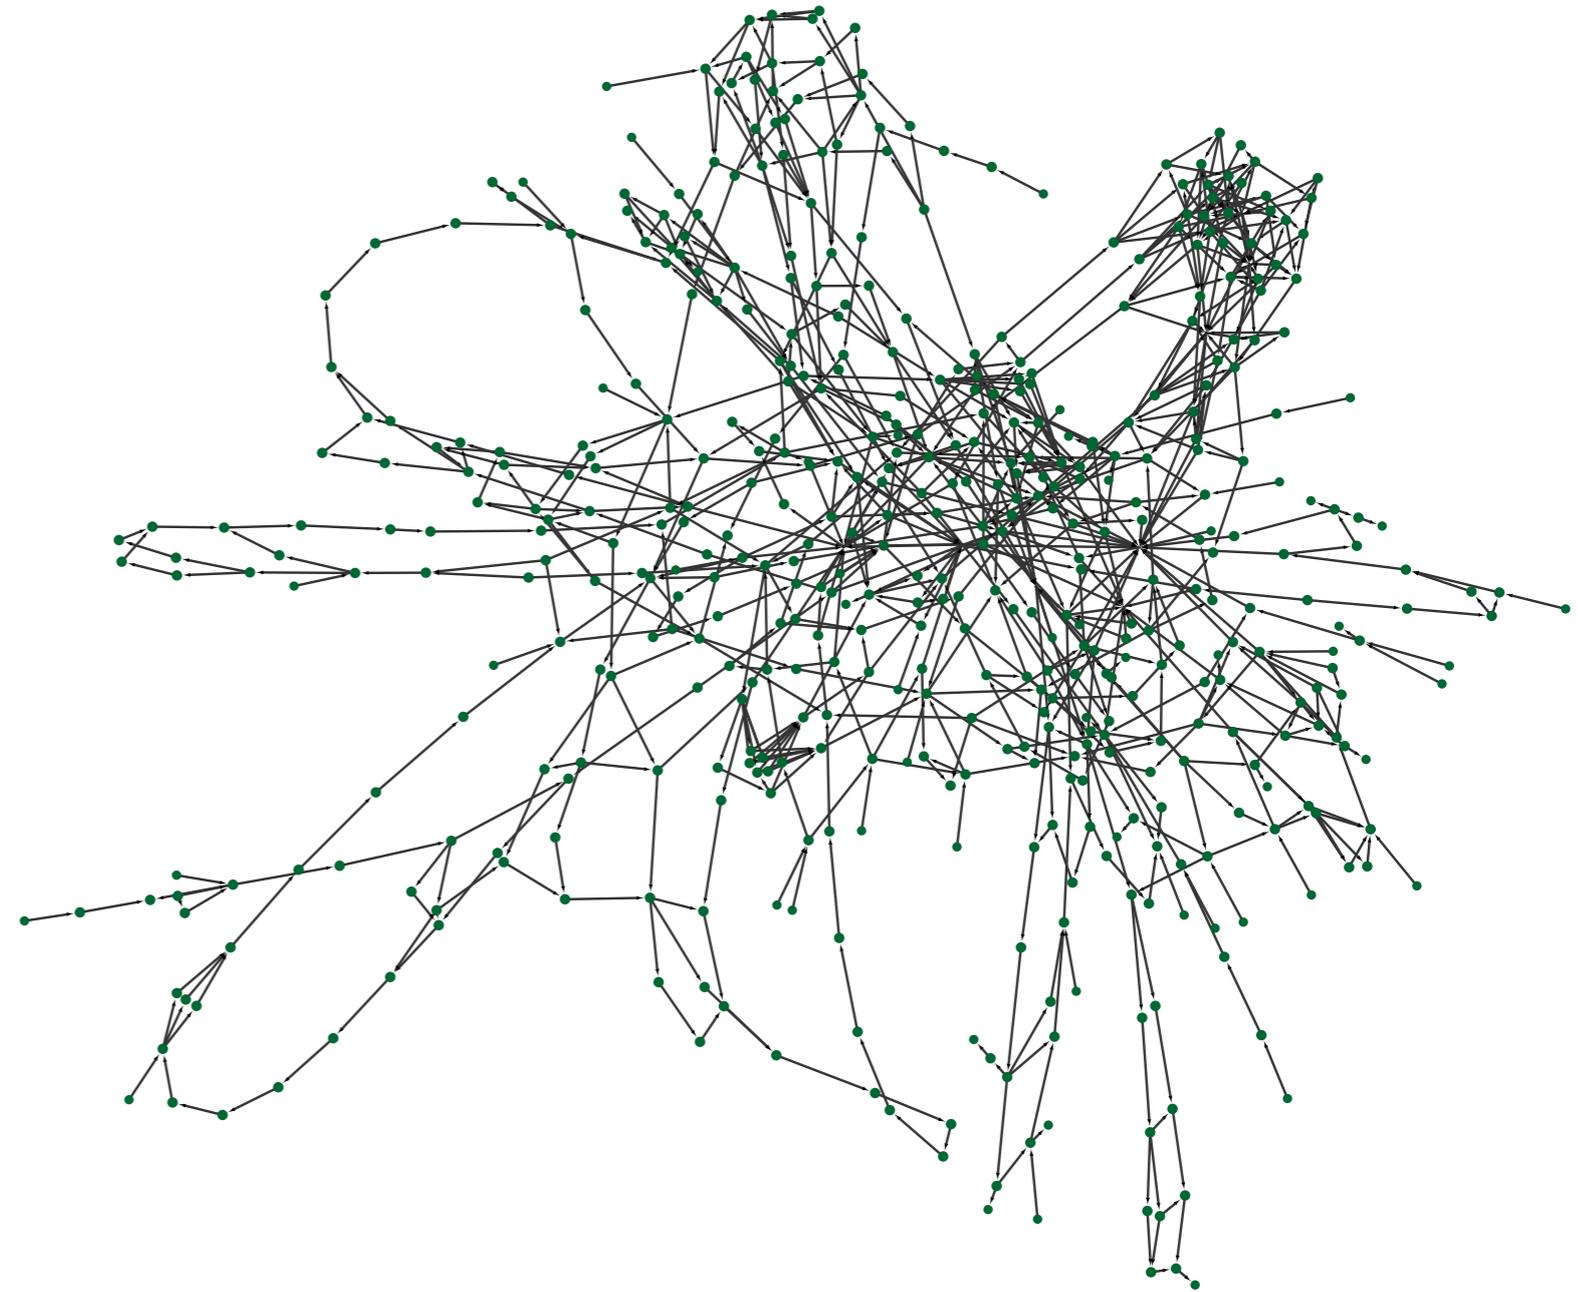
Scale



**size of AFP entries by submission date
with Four-Colour theorem, Odd-Order theorem, Verisoft, seL4**

Proof Introspection

- ▶ 500 files
- ▶ 22,000 lemmas stated
- ▶ 95,000 lemmas proved



Proof Introspection

- ▶ 500 files
- ▶ 22,000 lemmas stated
- ▶ 95,000 lemmas proved



Raf's Observation

The introspection of proof and theories is an essential part of working on a large-scale verification development.

- Learning Isabelle? **Easy.**
- Learning microkernels? **Not too bad.**
- Finding your way in the 500kloc proof jungle? **Hard!**

DATA
61



Proof Engineering

Software vs Proof Engineering



▶ Is Proof Engineering a thing?

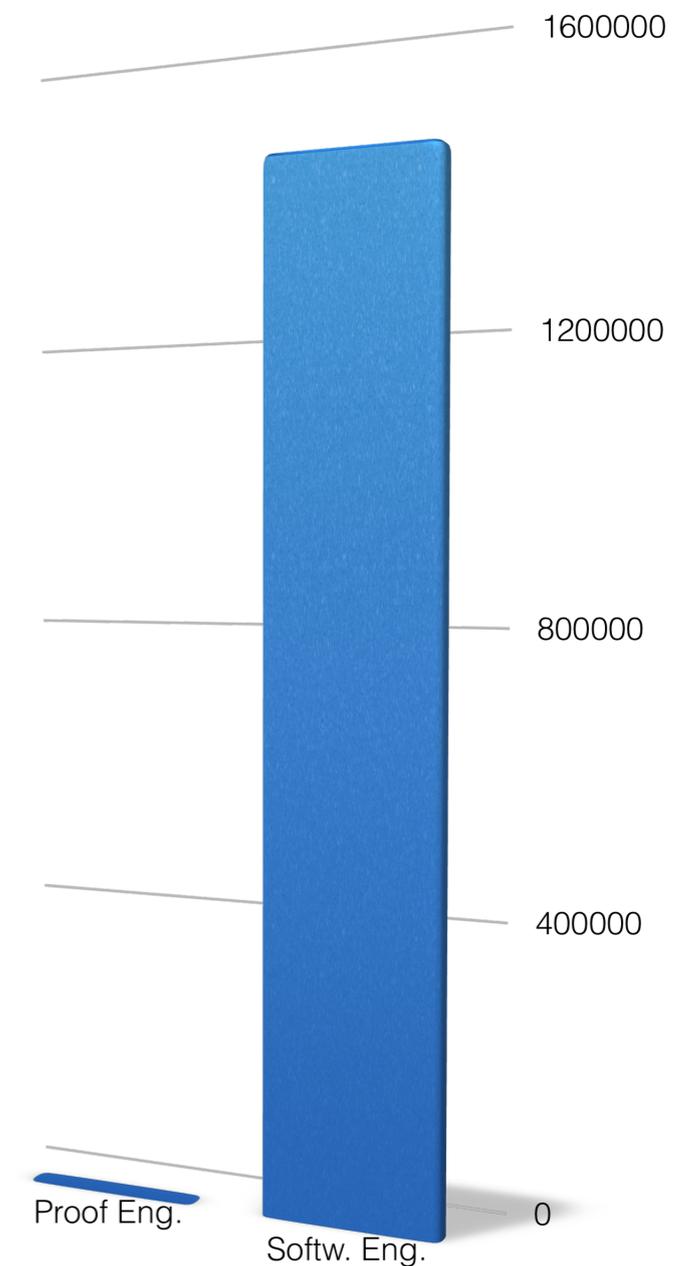
- Google Scholar:
 - “software engineering” 1,430,000 results

Software vs Proof Engineering



► Is Proof Engineering a thing?

- Google Scholar:
 - “software engineering” 1,430,000 results
 - “proof engineering” 564 results



Software vs Proof Engineering

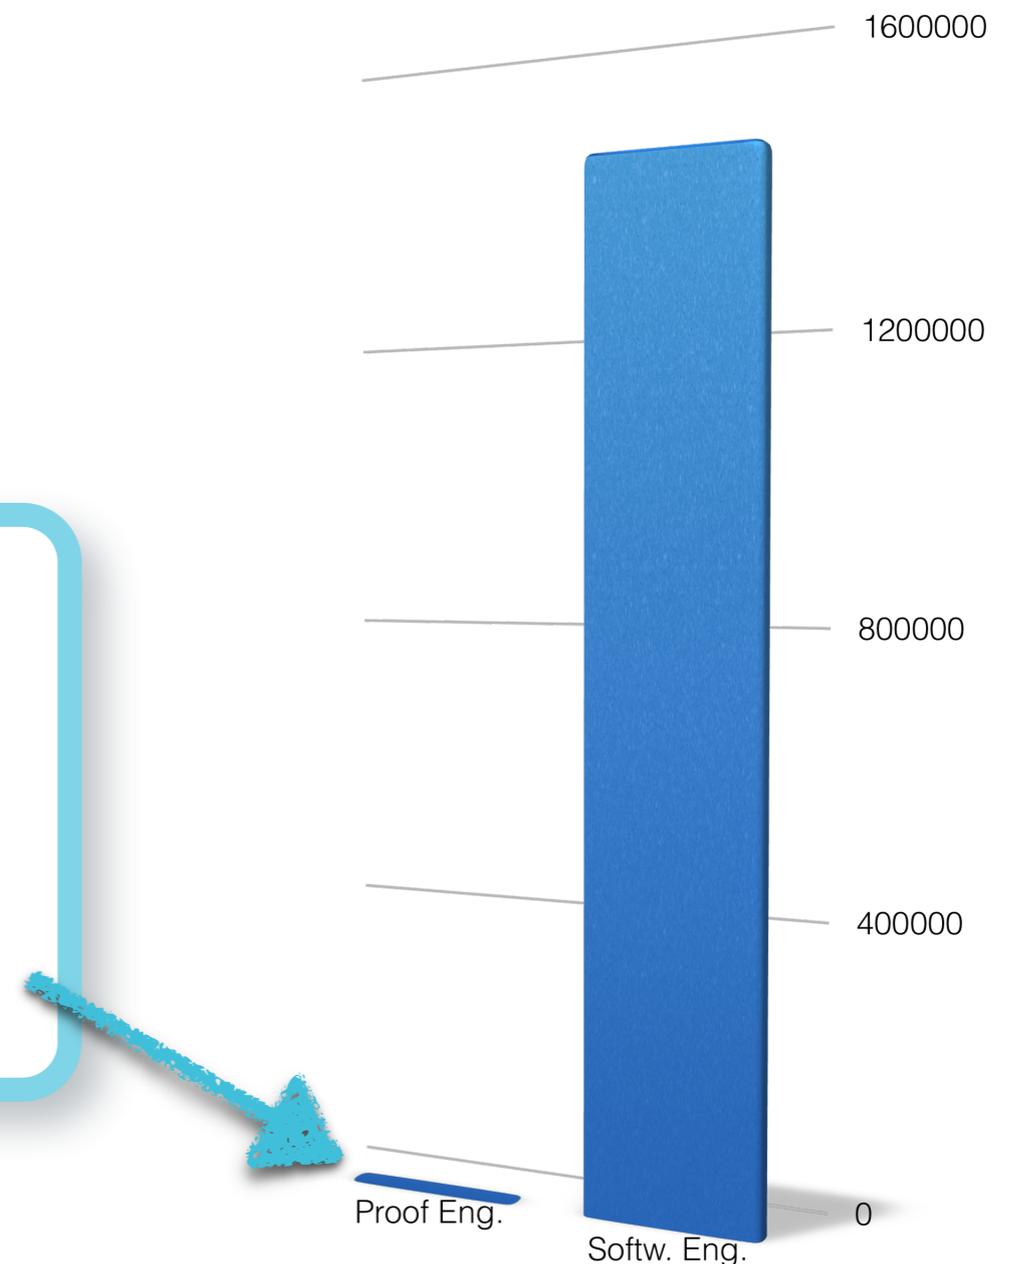


► Is Proof Engineering a thing?

- Google Scholar:
 - “software engineering” 1,430,000 results
 - “proof engineering” 564 results

Includes

“The Fireproof Building” and
“Influence of water permeation and analysis
of treatment for the Longmen Grottoes”

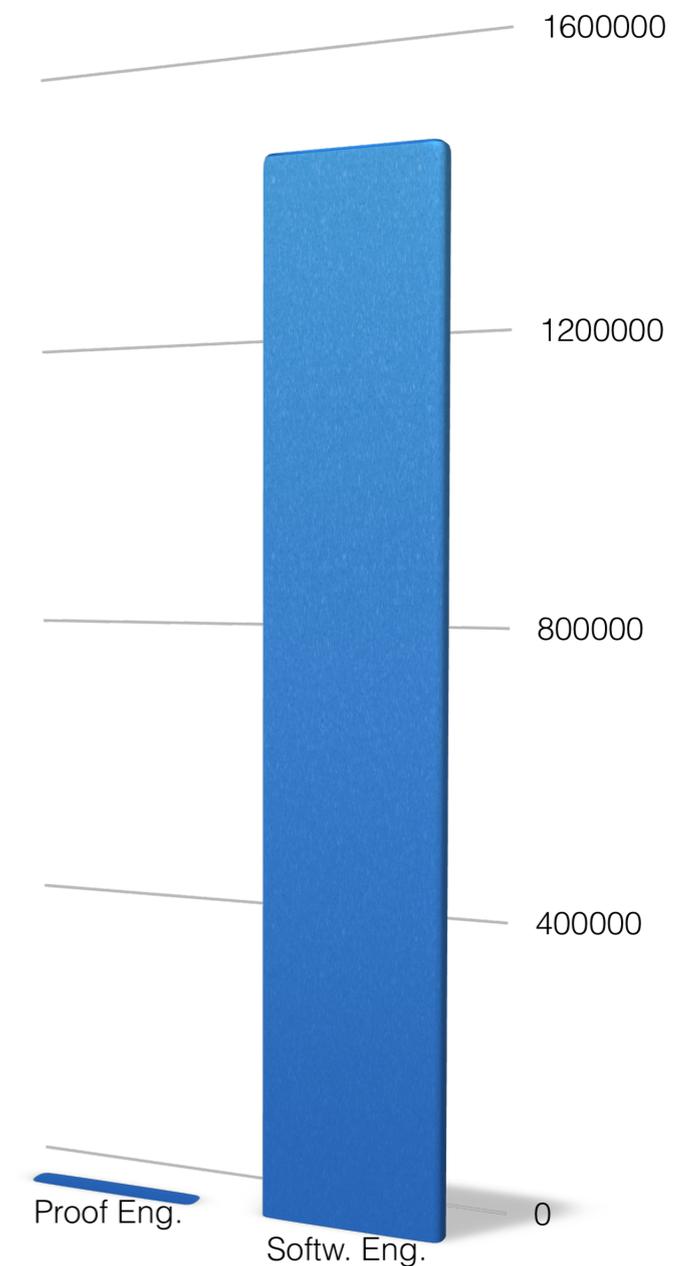


Proof Engineering is The Same



▶ Same kind of artefacts:

- lemmas are functions, modules are modules
- code gets big too
- version control, regressions, refactoring and IDEs apply



Proof Engineering is The Same

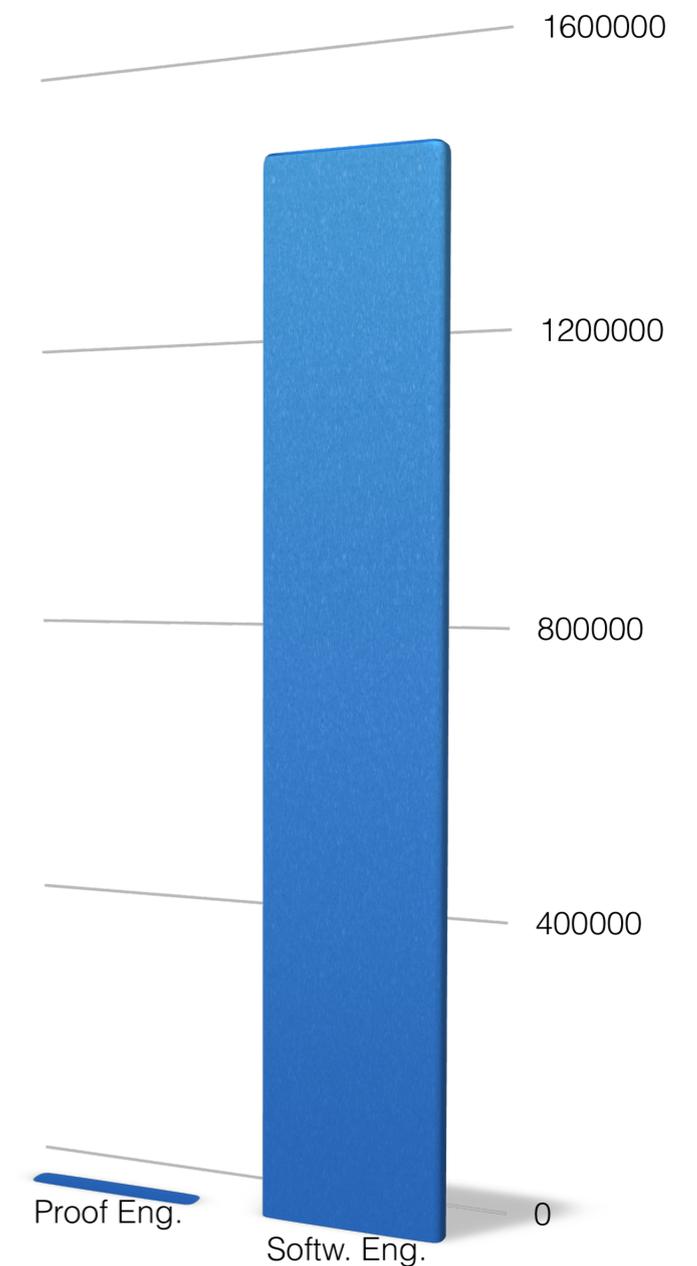


▶ Same kind of artefacts:

- lemmas are functions, modules are modules
- code gets big too
- version control, regressions, refactoring and IDEs apply

▶ Same kind of problems

- managing a large proof base over time
- deliver a proof on time within budget
- dependencies, interfaces, abstraction, etc

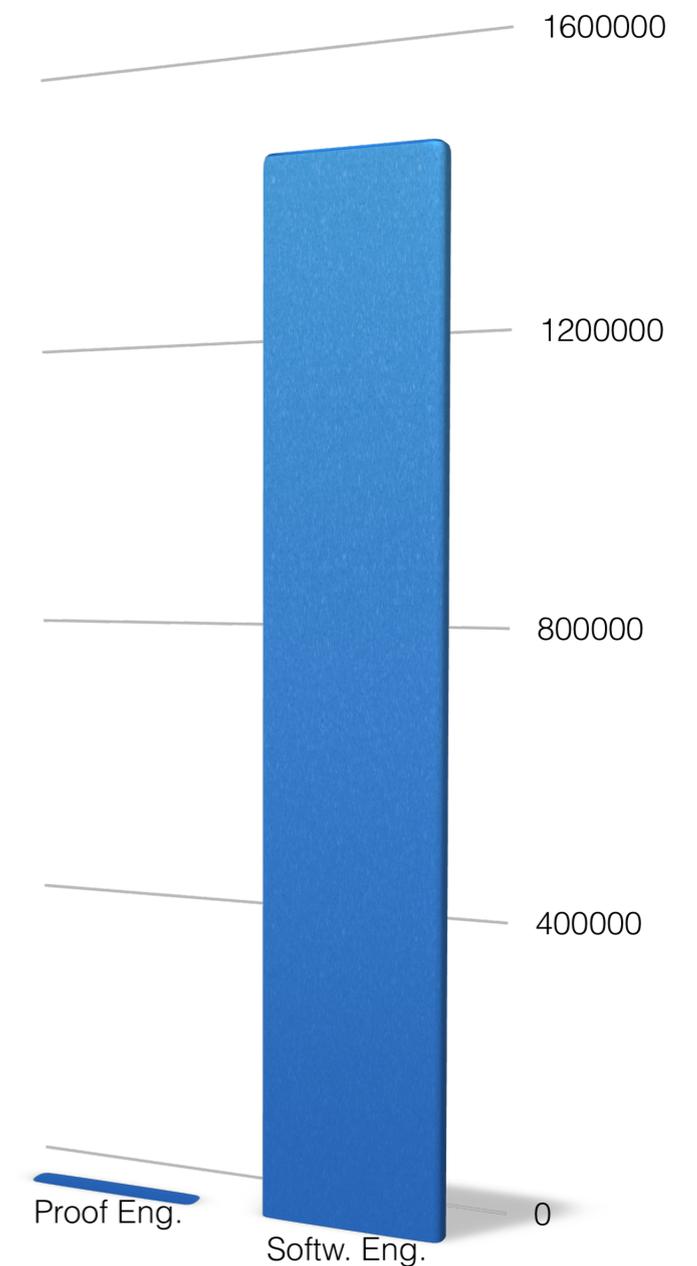


Proof Engineering is Different



► But: New Properties and Problems

- Results are checkable
 - You know when you are done!
 - No testing
 - 95% proof: no such thing
- More dead ends and iteration
- 2nd order artefact
 - Performance less critical
 - Quality less critical
 - Proof Irrelevance
- More semantic context
 - Much more scope for automation

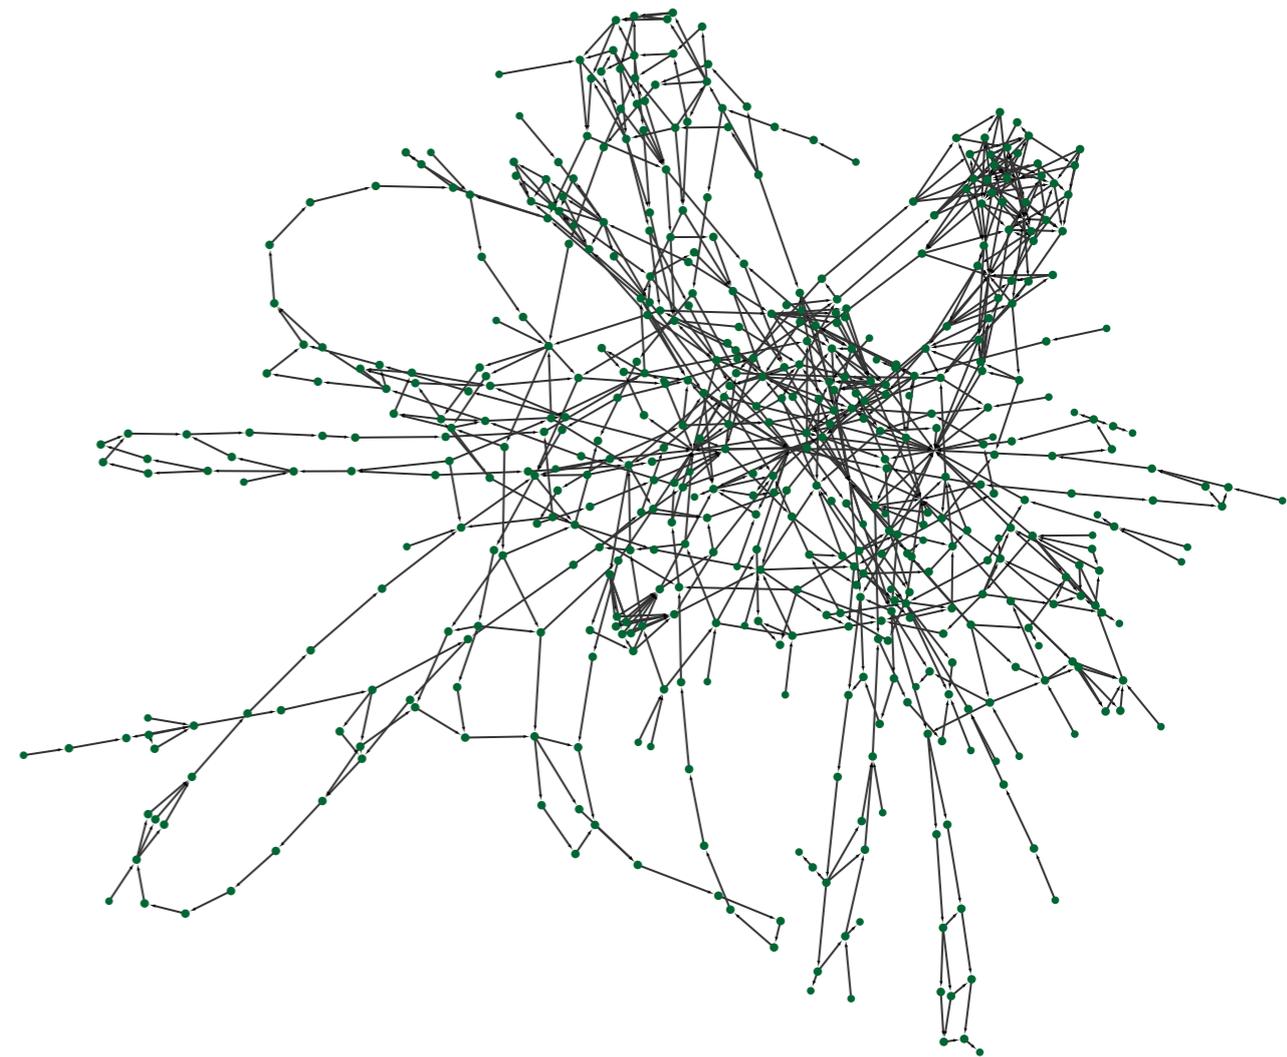


Proof Development



► Proof development

- decomposition of proofs over people,
- custom proof calculus,
- automating mechanical tasks, custom tactics
- proof craft



Proof Development

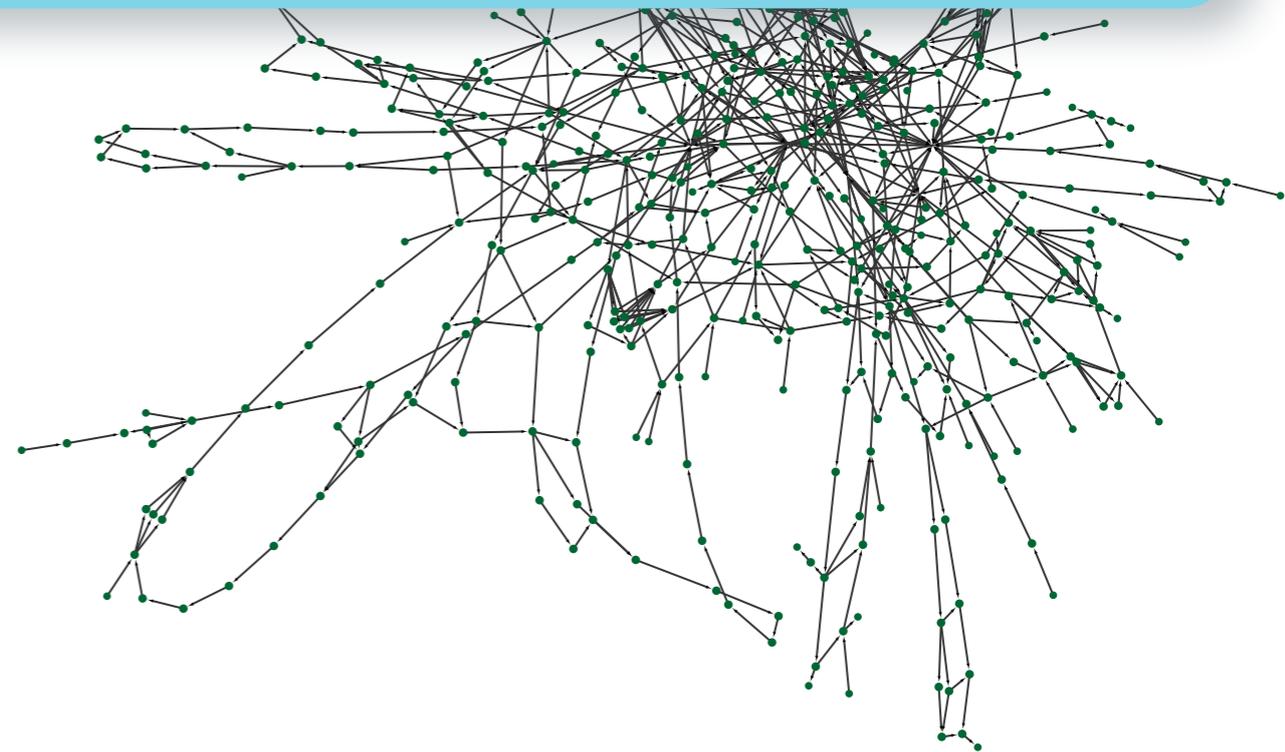


► Proof development

- decomposition of proofs over people
- custom proof calculus,
- automating mechanical tasks, custom
- proof craft

Tim's Statement

Automating “donkey work” allows attention and effort to be focussed where most needed – but it must be done judiciously.



Proof Development



► Proof development

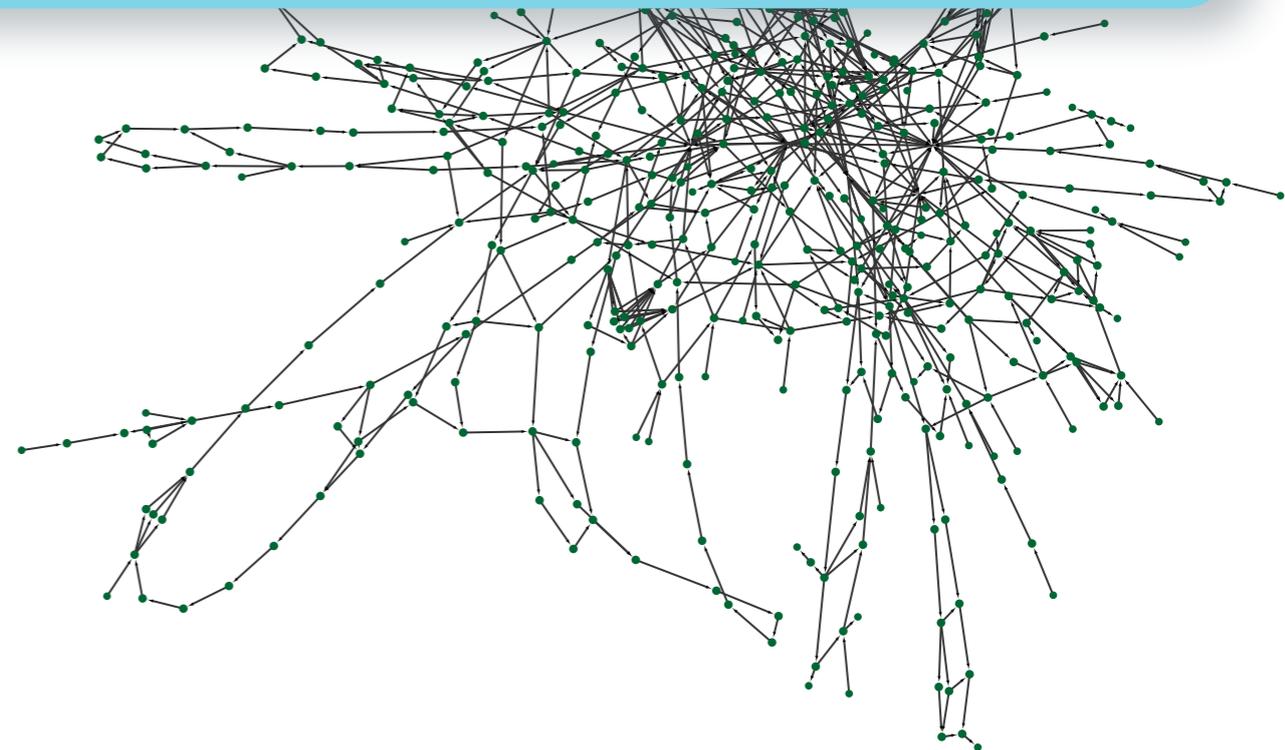
- decomposition of proofs over people
- custom proof calculus,
- automating mechanical tasks, custom
- proof craft

► Challenges

- non-local change,
- speculative change,
- distributed development

Tim's Statement

Automating “donkey work” allows attention and effort to be focussed where most needed – but it must be done judiciously.



Proof Development



▶ Proof development

- decomposition of proofs over people
- custom proof calculus,
- automating mechanical tasks, custom
- proof craft

▶ Challenges

- non-local change,
- speculative change,
- distributed development

Tim's Statement

Automating “donkey work” allows attention and effort to be focussed where most needed – but it must be done judiciously.



Matthias' Conjecture

Over the years, I must have waited weeks for Isabelle. Productivity hinges on a short edit-check cycle; for that, I am even willing to (temporarily) sacrifice soundness.

Problems of Scale

- ▶ **Proof maintenance**
 - changes, updates, new proofs, new features
 - automated regression, keep code in sync
 - refactoring
 - simplification
- ▶ **Original proof: 2005-2009**
- ▶ **Maintenance: 2009-2016 and counting**



Problems of Scale

- ▶ **Proof maintenance**
 - changes, updates, new proofs, new features
 - automated regression, keep code in sync
 - refactoring
 - simplification
- ▶ **Original proof: 2005-2009**
- ▶ **Maintenance: 2009-2016 and counting**

Dan's Conclusion

Verification is fast, maintenance is forever.



Proof Engineering Tools



► User Interface

- could proof IDEs be more powerful than code IDEs?
- more semantic information
- proof completion and suggestion?

The screenshot shows the Isabelle proof IDE interface. The main window displays a theory named 'Example' with the following code:

```
theory Example
imports Base
begin

inductive path for R :: "'a ⇒ 'a ⇒ bool" where
  base: "path R x x"
| step: "R x y ⇒ path R y z ⇒ path R x z"

theorem example:
  fixes x z :: 'a assumes "path R x z" shows "P x z"
  using assms
proof induct
  case (base x)
  show "P x x" by auto
next
  case (step x y z)
  note `R x y` and `path R y z`
  moreover note `P y z`
  ultimately show "P x z" by auto
qed
end
```

The right-hand side of the IDE shows a project browser with a tree view of the theory structure. The status bar at the bottom indicates the version '5,1 (35/405)' and the current session information '(isabelle,sidekick,UTF-8-Isabelle)Nm r o UC467120Mb 3:38 PM'.

Proof Engineering Tools



► User Interface

- could proof IDEs be more powerful than code IDEs?
- more semantic information
- proof completion and suggestion?

► Refactoring

- less constrained,
new kinds of refactoring possible, e.g.
 - move to best position in library
 - generalise lemma
 - recognise proof patterns

A screenshot of the Isabelle proof IDE interface. The main window displays a theory file named 'Example.thy'. The code is as follows:

```
theory Example
imports Base
begin

inductive path for R :: "'a => 'a => bool" where
  base: "path R x x"
| step: "R x y => path R y z => path R x z"

theorem example:
  fixes x z :: 'a assumes "path R x z" shows "P x z"
  using assms
proof induct
  case (base x)
  show "P x x" by auto
next
  case (step x y z)
  note `R x y` and `path R y z`
  moreover note `P y z`
  ultimately show "P x z" by auto
qed
end
```

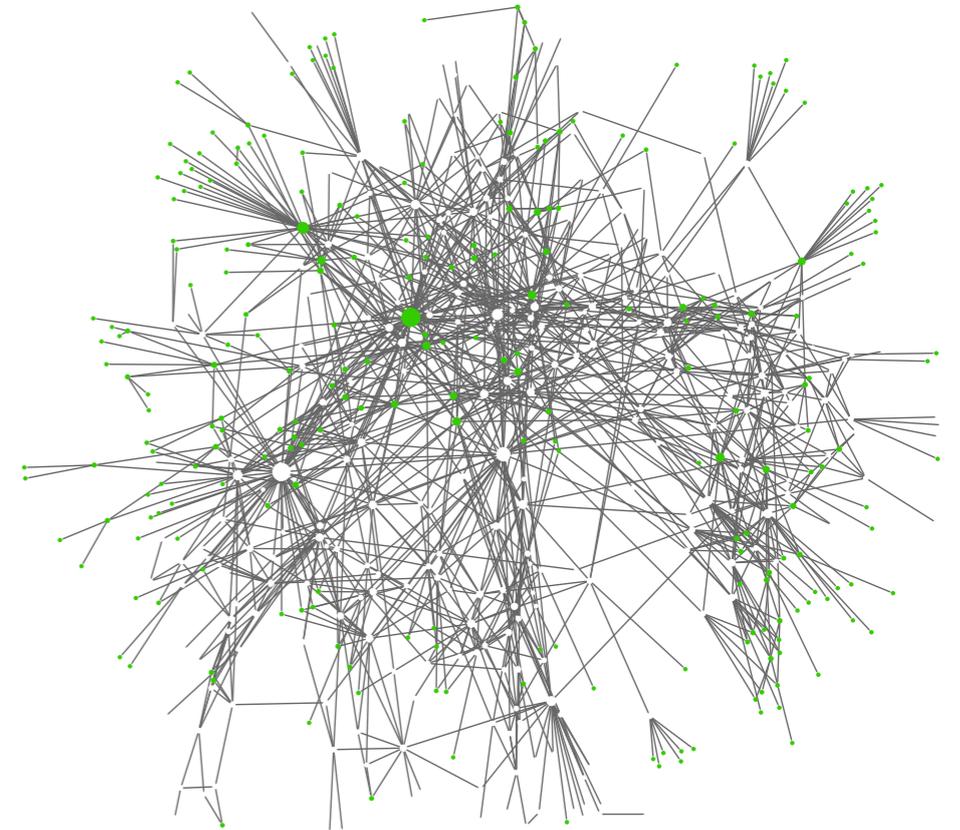
The interface includes a sidebar on the right with a search filter and a project tree showing the current theory. At the bottom, there are tabs for 'Output', 'Prover Session', and 'Raw Output', and a status bar showing '5,1 (35/405)' and '(isabelle,sidekick,UTF-8-Isabelle)Nm r o UG467.120Mb 3:38 PM'.

Proof Patterns



► Large-scale Libraries

- architecture:
 - layers, modules, components, abstractions, genericity
- proof interfaces
- proof patterns



Proof Patterns

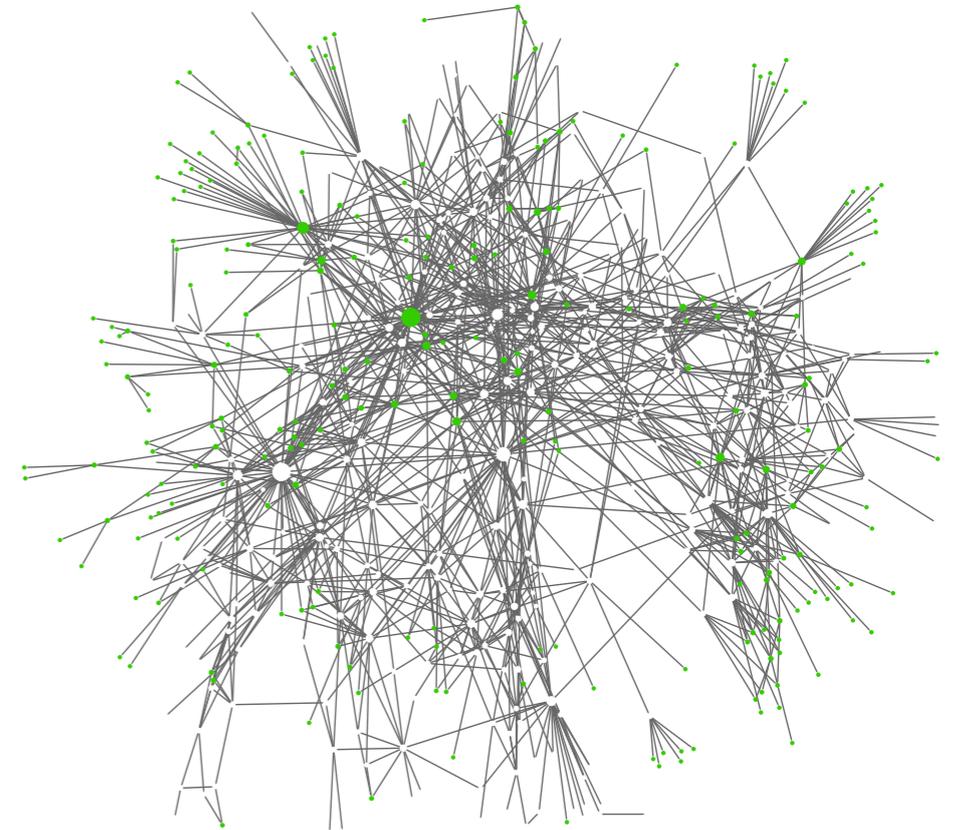


▶ Large-scale Libraries

- architecture:
 - layers, modules, components, abstractions, genericity
- proof interfaces
- proof patterns

▶ Technical Debt

- what does a clean, maintainable proof look like?
- which techniques will make future change easier?
- readability important? is documentation?



Proof Effort

Predictions

Can we predict for proofs:

- **how large will it be?**
- **how long will it take?**
- **how much will it cost?**



Predictions

Can we predict for proofs:

- **how large will it be?**
- **how long will it take?**

Of course not.

Many hard problems look deceptively easy.



Predictions

Can we predict for proofs:

- **how large will it be?**
- **how long will it take?**

Of course not.

Many hard problems look deceptively easy.

But maybe for program verification?

At least statistically, some of the time?



Predictions

Can we predict for proofs:

- **how large will it be?**
- **how long will it take?**

Of course

Many have

But many

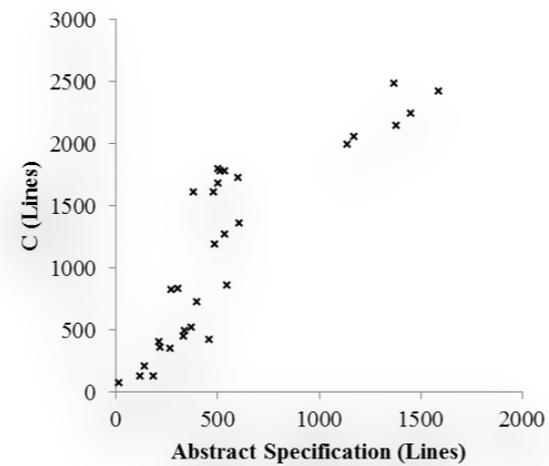
At least

**We have large proofs.
Let's crunch some data!**



Some Hope

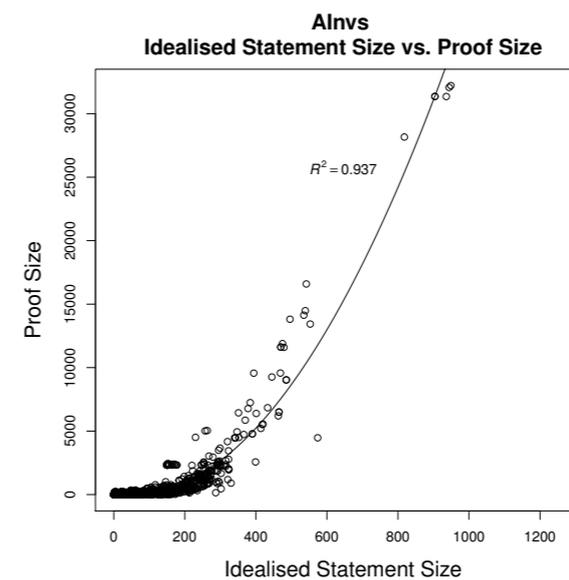
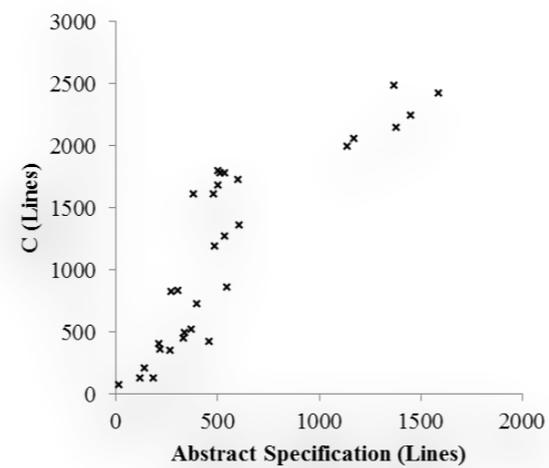
Code Size is correlated with Spec Size



Some Hope

Code Size is correlated with Spec Size

Spec Size is correlated with Proof Size

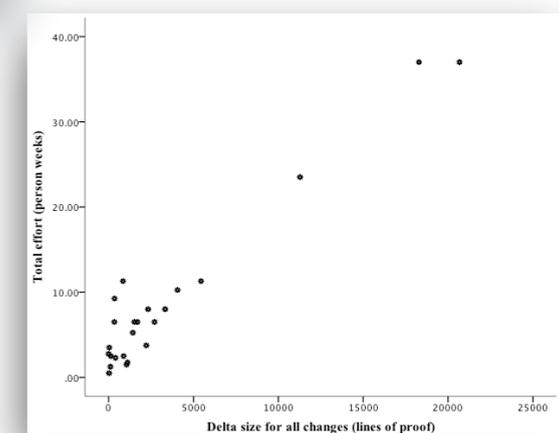
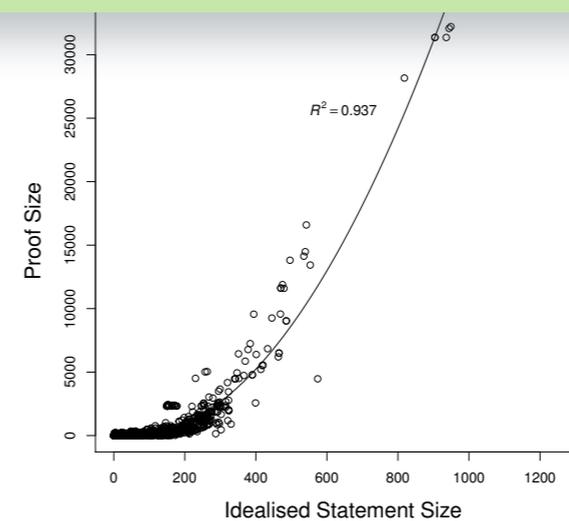
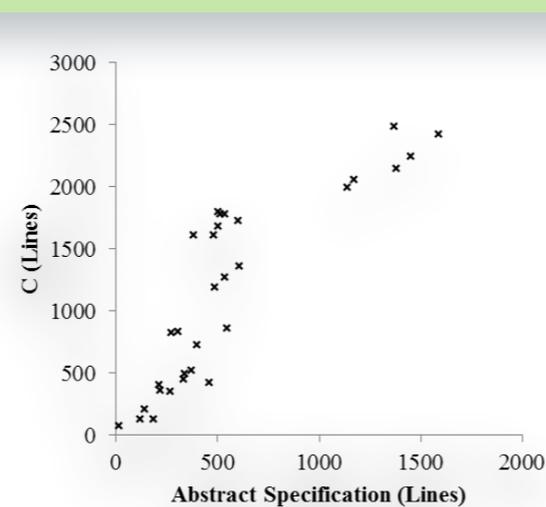


Some Hope

Code Size is correlated with Spec Size

Spec Size is correlated with Proof Size

Proof Size is correlated with Effort



Some Hope

Code Size is correlated with Spec Size

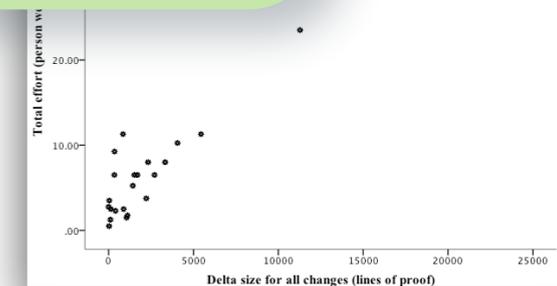
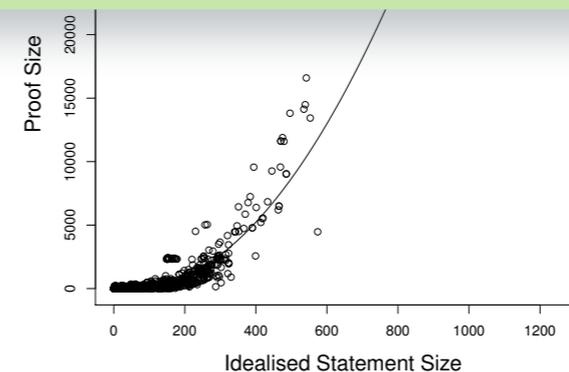
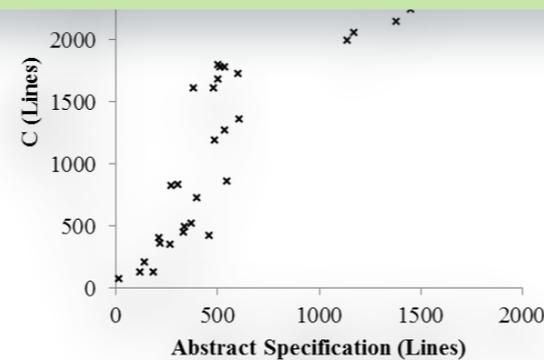
Spec

There may be hope for a prediction model.

Probably applies to verification of non-modular code.

Pro

Unlikely to work for other kinds of proofs, but likely to transfer to other interactive provers.



DATA
61



The Future

The Future: Integration



- ▶ **No method fits all**
- ▶ **Use seL4 isolation!**
 - don't verify all components
 - mix verification approaches

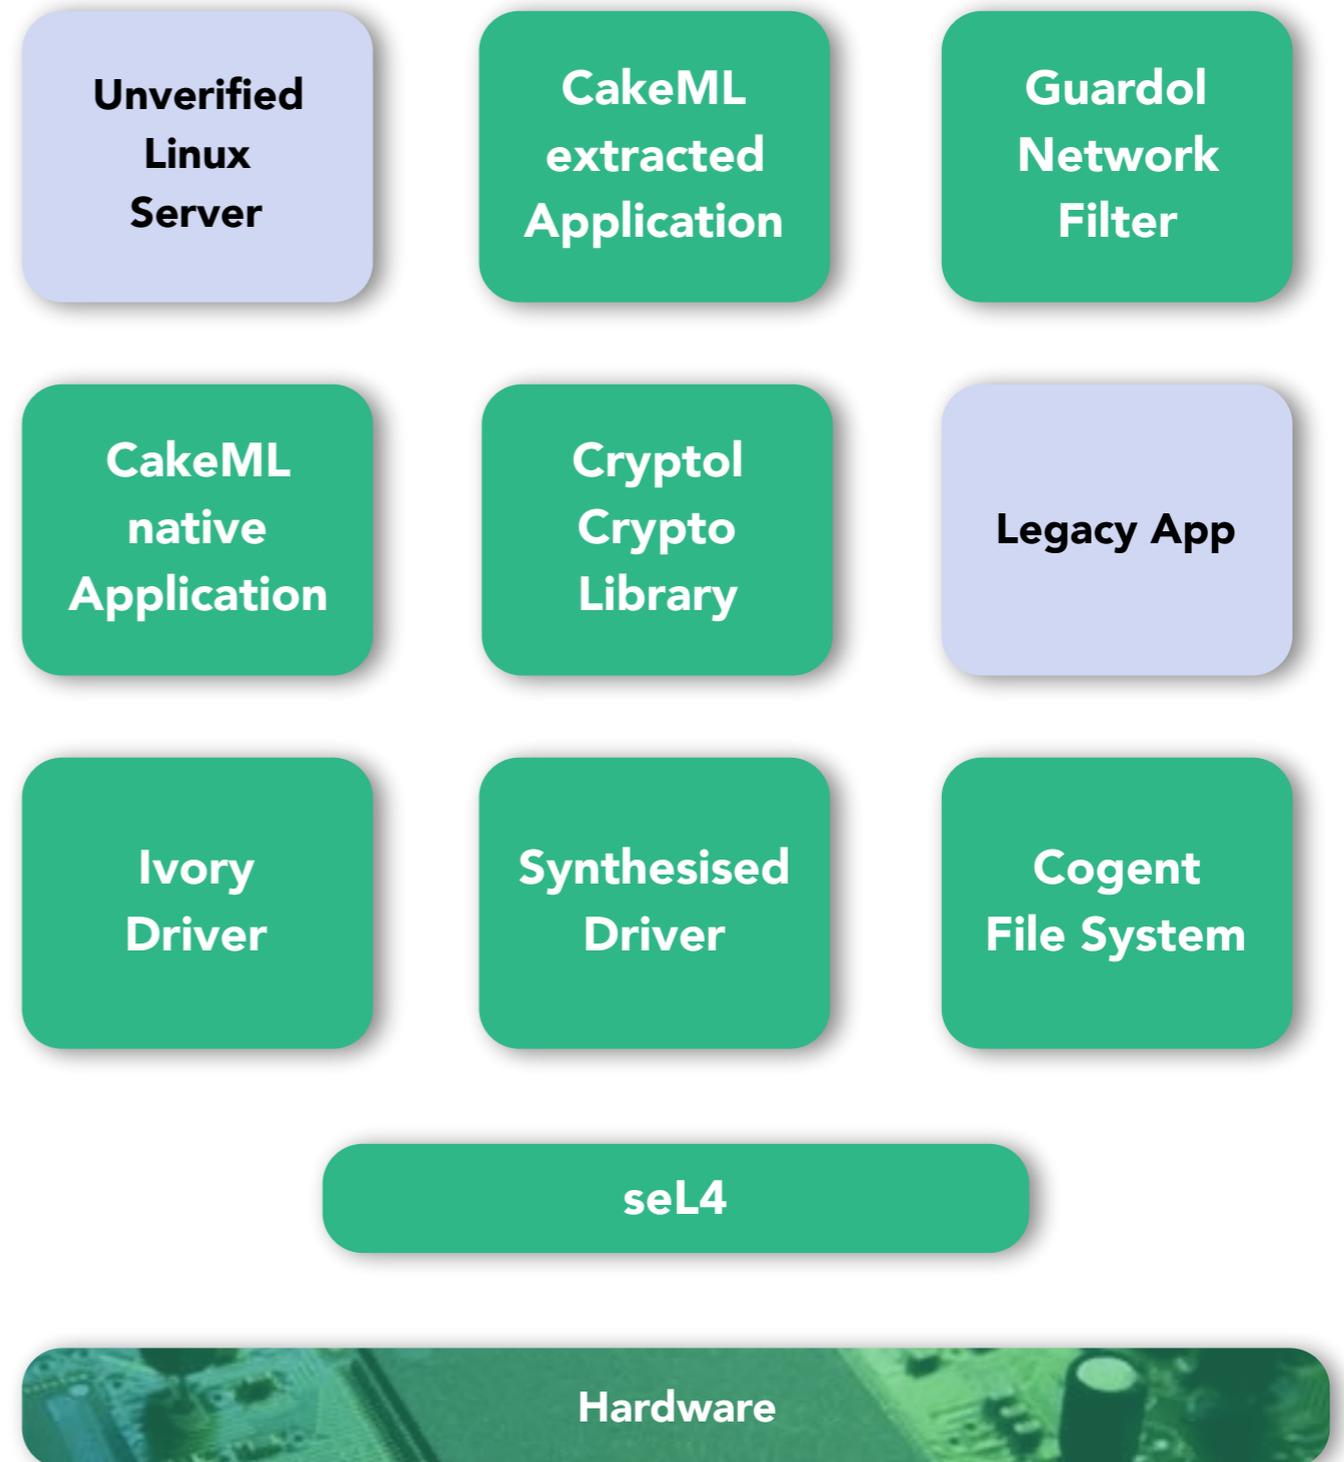
A horizontal banner image showing a close-up of a green printed circuit board (PCB) with various electronic components. The word 'Hardware' is overlaid in white text on the right side of the image.

Hardware

The Future: Integration



- ▶ **No method fits all**
- ▶ **Use seL4 isolation!**
 - don't verify all components
 - mix verification approaches



The Future: Integration



- ▶ **No method fits all**
- ▶ **Use seL4 isolation!**
 - don't verify all components
 - mix verification approaches

Unverified
Linux
Server

CakeML
extracted
Application

Guardol
Network
Filter

CakeML
native
Application

Cryptol
Crypto
Library

Legacy App

Ivory

Synthesised

Cogent
File System

Will need formal interfaces

seL4

Hardware

Summary



- **Verification of real systems is happening**

Summary



- **Verification of real systems is happening**

- **It's still too expensive**

Summary



- **Verification of real systems is happening**

- **It's still too expensive**

- **There is hope**

Summary



- **Verification of real systems is happening**

- **It's still too expensive**

- **There is hope**

- **Ongoing work on**
 - **Proof Engineering**
 - **Languages for verification productivity**
 - **Increased Automation**

Summary



- **Verification of real systems is happening**

- **It's still too expensive**

- **There is hope**

- **Ongoing work on**
 - **Proof Engineering**
 - **Languages for verification productivity**
 - **Increased Automation**

- **Integration will be key**



Thank You

Trustworthy Systems
Gerwin Klein

t +61 2 8306 0578
e gerwin.klein@nicta.com.au
w <http://trustworthy.systems>

data61.csiro.au

