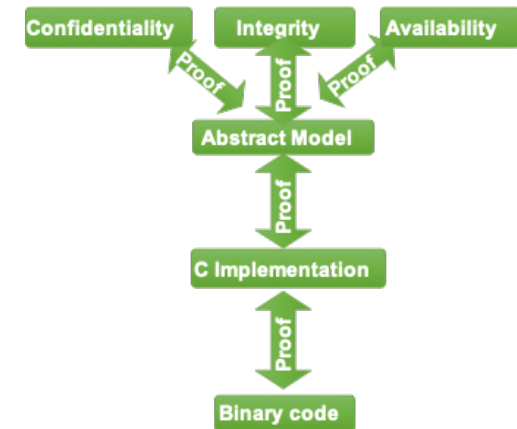


School of Computer Science & Engineering
COMP9242 Advanced Operating Systems

2021 T2 Week 10 Part 1

Formal Verification and seL4

@GernotHeiser



Copyright Notice

These slides are distributed under the Creative Commons Attribution 3.0 License

- You are free:
 - to share—to copy, distribute and transmit the work
 - to remix—to adapt the work
- under the following conditions:
 - **Attribution:** You must attribute the work (but not in any way that suggests that the author endorses you or your use of the work) as follows:

“Courtesy of Gernot Heiser, UNSW Sydney”

The complete license text can be found at
<http://creativecommons.org/licenses/by/3.0/legalcode>

Assurance and Verification

Refresher: Assurance and Formal Verification

- **Assurance:**
 - systematic evaluation and testing
 - essentially an intensive and onerous form of quality assurance
- **Formal verification:**
 - mathematical proof
- **Certification:** independent examination
 - confirming that the assurance or verification was done right

Assurance and formal verification aim to establish correctness of

- mechanism design
- mechanism implementation

Assurance: Substantiating Trust

- Specification
 - Unambiguous description of desired behaviour
- System design
 - Justification that it meets specification
- Implementation
 - Justification that it implements the design
- Maintenance
 - Justifies that system use meets assumptions

Informal (English)
or formal (maths)

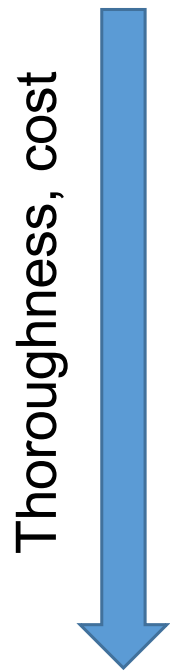
Compelling argument
or formal proof

Code inspection,
rigorous testing,
proof

Common Criteria

- *Common Criteria for IT Security Evaluation* [ISO/IEC 15408, 99]
 - ISO standard, for general use
 - Evaluates QA used to ensure systems meet their requirements
 - Developed out of the famous US DOD “Orange Book”:
Trusted Computer System Evaluation Criteria [1985]
- Terminology:
 - *Target of evaluation* (TOE): Evaluated system
 - *Security target* (ST): Defines requirements
 - *Protection profile* (PP): Standardised ST template
 - *Evaluation assurance level* (EAL): Defines thoroughness of evaluation
 - PPs have maximum EAL they can be used for

CC: Evaluation Assurance Levels



Level	Requirements	Specification	Design	Implementation
EAL1	not evaluated	Informal	not eval	not evaluated
EAL2	not evaluated	Informal	Informal	not evaluated
EAL3	not evaluated	Informal	Informal	not evaluated
EAL4	not evaluated	Informal	Informal	not evaluated
EAL5	not evaluated	Semi-Formal	Semi-Formal	Informal
EAL6	Formal	Semi-Formal	Semi-Formal	Informal
EAL7	Formal	Formal	Formal	Informal

Common Criteria: Protection Profiles (PPs)

- *Controlled Access PP* (CAPP)
 - standard OS security, up to EAL3
- *Single Level Operating System PP*
 - superset of CAPP, up to EAL4+
- *Labelled Security PP* (LSPP)
 - MAC for COTS OSes
- *Multi-Level Operating System PP*
 - superset of CAPP, LSPP, up to EAL4+
- *Separation Kernel Protection Profile* (SKPP)
 - strict partitioning, for EAL6-7

COTS OS Certifications

- EAL3:
 - 2010 Mac OS X (10.6)
- EAL4:
 - 2003: Windows 2000
 - 2005: SuSE Enterprise Linux
 - 2006: Solaris 10 (EAL4+)
 - against CAPP (an EAL3 PP!)
 - 2007: Red Hat Linux (EAL4+)
- EAL6:
 - 2008: Green Hills INTEGRITY-178B (EAL6+)
 - against SKPP, relatively simple PPC-based hardware platform in TOE
- EAL7:
 - 2019: Prove & Run PROVENCORE



Get regularly
hacked!

SKPP on Commodity Hardware

- SKPP: OS provides only separation
- One Box One Wire (OB1) Project
 - Use INTEGRITY-178B to isolate VMs on commodity desktop hardware
 - Leverage existing INTEGRITY certification
 - by “porting” it to commodity platform


NSA subsequently dis-endorsed SKPP,
discontinued certifying \geq EAL5

Conclusion [NSA, March 2010]:

- SKPP validation for commodity hardware platforms infeasible due to their complexity
- SKPP has limited relevance for these platforms

Common Criteria Limitations

- Very expensive
 - rule of thumb: EAL6+ costs \$1K/LOC design-implementation-evaluation-certification
- Too much focus on development process
 - rather than the product that was delivered
- Lower EALs of little practical use for OSEs
 - c.f. COTS OS EAL4 certifications
- Commercial Licensed Evaluation Facilities licenses rarely revoked
 - Leads to potential “race to the bottom” [Anderson & Fuloria, 2009]



Effectively dead
in 5-Eyes defence

Formal Verification

- Prove properties about a mathematical model of a system

Model checking / abstract interpretation:

- ☐ Cannot generally prove code correct
 - Proves specific properties
- ☐ Generally have to
 - over-approximate (false positives), or
 - under-approximate (false negatives, unsound)
- ☐ Suffers state-space explosion
- ✓ Automatic
- ✓ May scale to large code bases

Theorem proving:

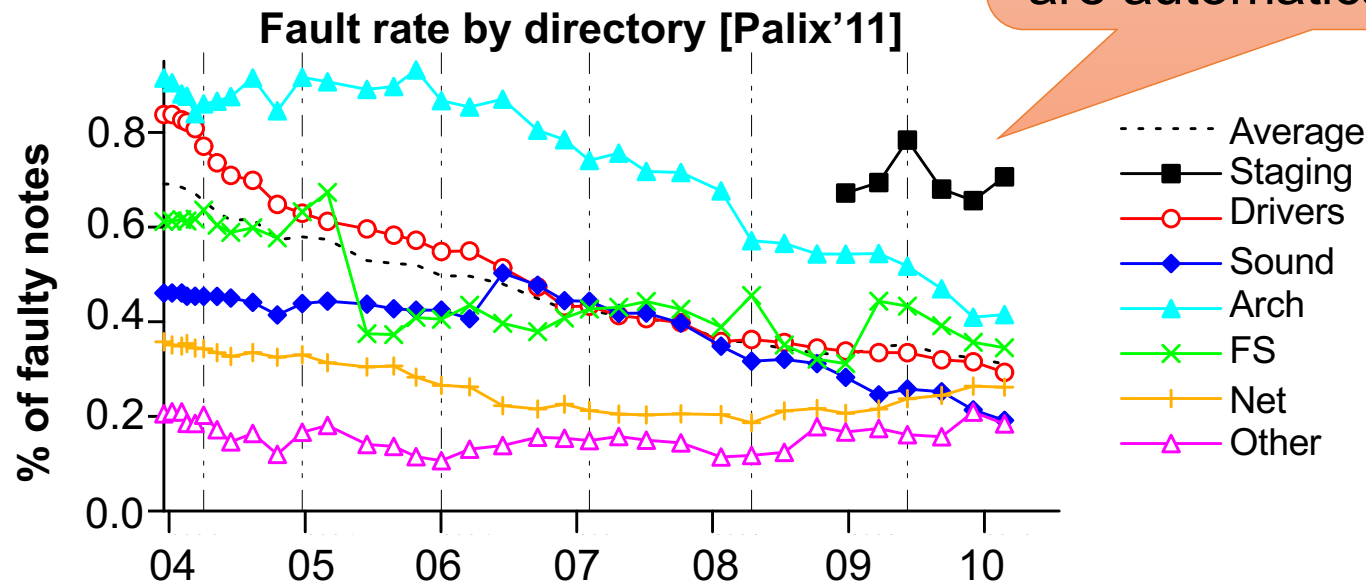
- ✓ Can deal with large (even infinite) state spaces
- ✓ Can prove functional correctness against a spec
- ☐ Very labour-intensive

Recent work automatically proved functional correctness of simple systems using SMT solvers [Hyperkernel, SOSP'17]

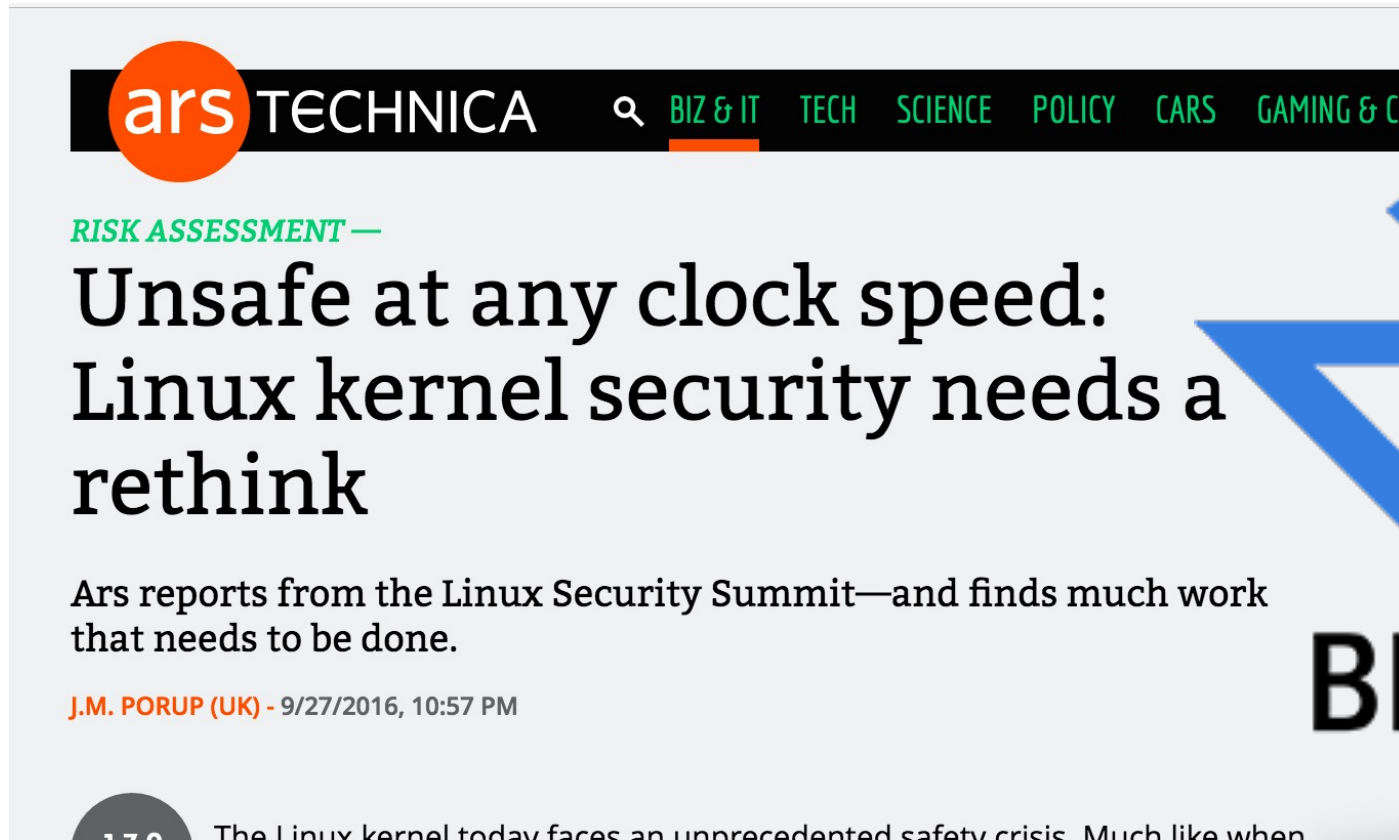
Model Checking and Linux: A Sad Story

- Static analysis of Linux source [Chou & al, 2001]
 - Found high density of bugs, especially in device drivers
- Re-analysis 10 years later [Palix & al, 2011]

Disappointing rate of improvement for bugs that are automatically detectable!



And the Result?



The screenshot shows the top of an Ars Technica article. The header includes the 'ars TECHNICA' logo and a navigation bar with links for 'BIZ & IT', 'TECH', 'SCIENCE', 'POLICY', 'CARS', and 'GAMING & CULTURE'. The article is categorized under 'RISK ASSESSMENT'. The title is 'Unsafe at any clock speed: Linux kernel security needs a rethink'. The sub-headline reads: 'Ars reports from the Linux Security Summit—and finds much work that needs to be done.' The author is 'J.M. PORUP (UK)' and the date is '9/27/2016, 10:57 PM'. The beginning of the article text is visible: 'The Linux kernel today faces an unprecedented safety crisis. Much like when'.

ars TECHNICA

RISK ASSESSMENT —

Unsafe at any clock speed: Linux kernel security needs a rethink

Ars reports from the Linux Security Summit—and finds much work that needs to be done.

J.M. PORUP (UK) - 9/27/2016, 10:57 PM

The Linux kernel today faces an unprecedented safety crisis. Much like when





August 2009

A NICTA bejelentette a világ első, formális módszerekkel igazolt,



► [Stories](#) [Recent](#) [Popular](#) [Search](#)

Slashdot is powered by [your submissions](#)

+ - **Technology: World's First**

Posted by [Soulskill](#) on Thursday Aug 6, 2009
from the wait-for-it dept.

An anonymous reader writes

"Operating systems usually have bugs and so forth are known by almost everyone. To [prove that a particular OS kernel](#) is formally verified, and as such it is the first. Researchers used an executable model of the Isabelle theorem prover to generate a proof that matches the executable and the kernel.

Does it run Linux? "We're pleased to say that it does."



New Scientist

Saturday 29/8/2009

Page: 21

Section: General News

Region: National

Type: Magazines Science / Technology

Size: 196.31 sq.cms.

Published: -----S-

The ultimate way to keep your computer safe from harm

FLAWS in the code, or "kernel", that sits at the heart of modern computers leave them prone to occasional malfunction and vulnerable to attack by worms and viruses. So the development of a secure general-purpose microkernel could pave the

way to a more secure system, says Klein.

His team formulated a model with more than 200,000 logical steps which allowed them to prove that the program would always behave as its

eredményeiképpen pedig egy olyan megbízhatóságot kapnak a szoftvertől, amely e



10 BREAKTHROUGH TECHNOLOGIES

[Share](#)

2011

Crash-Proof Code

Making critical software safer

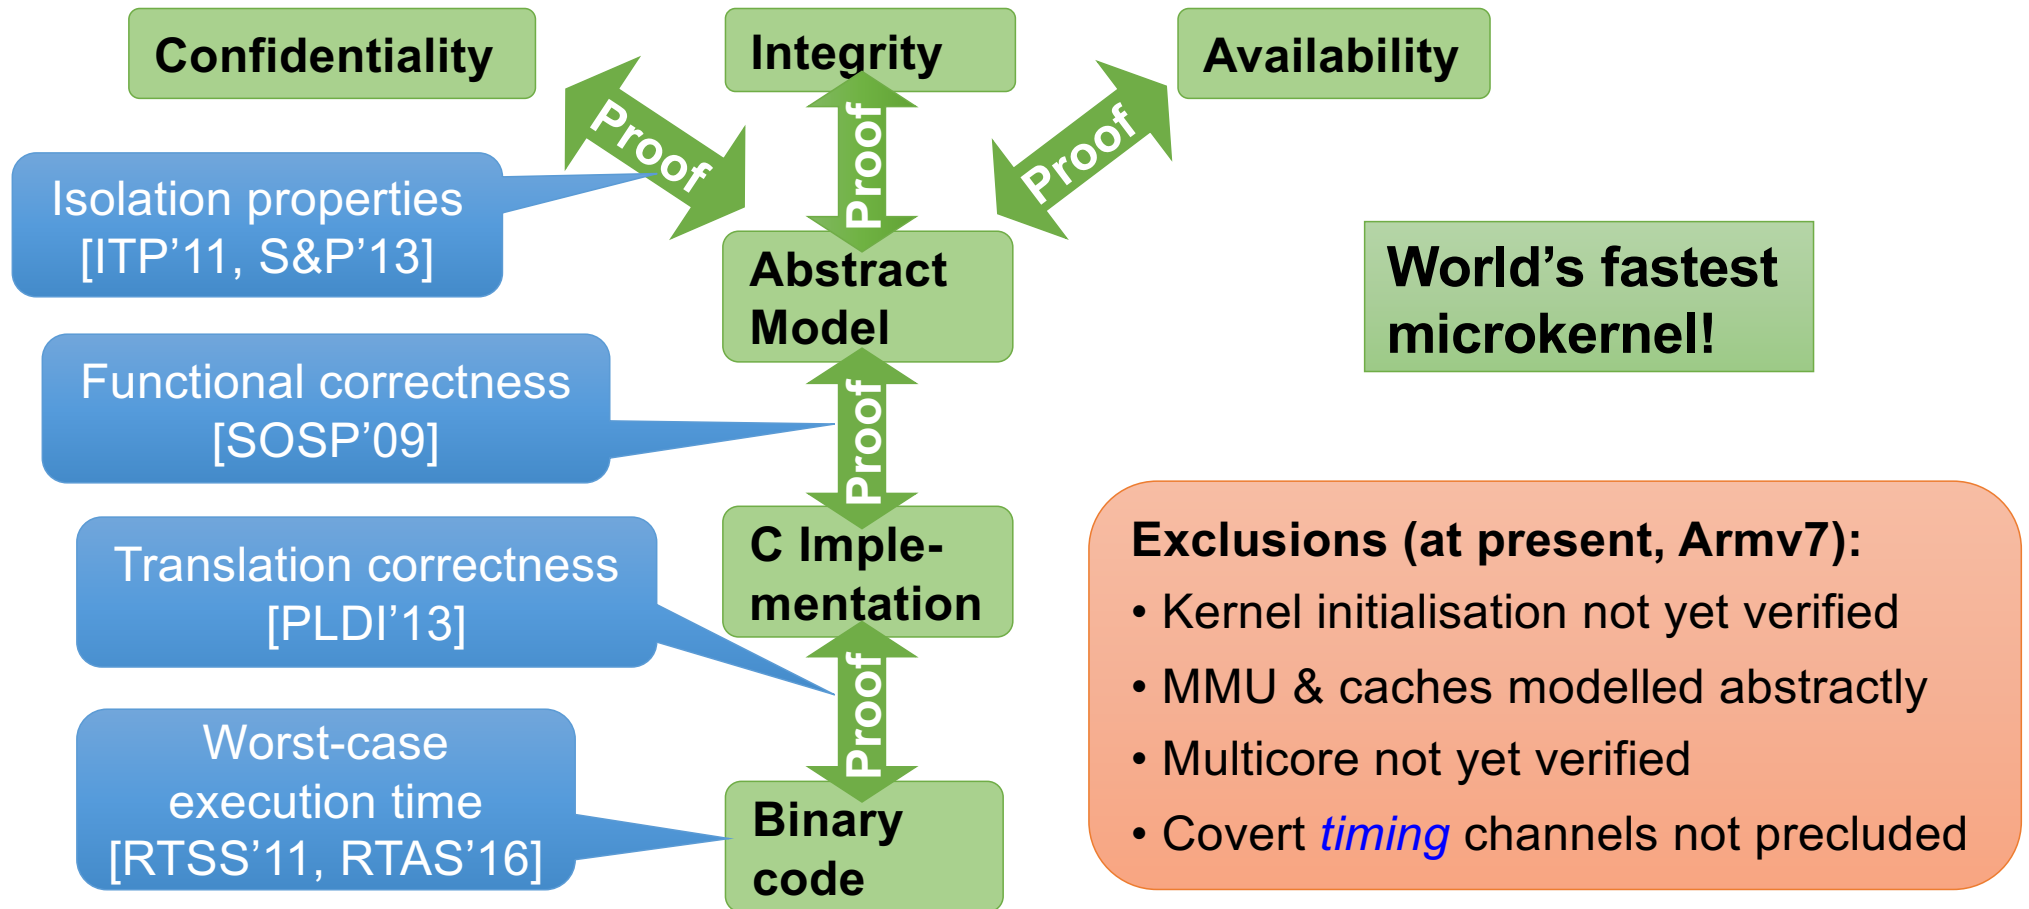
7 comments

WILLIAM BULKELEY

May/June 2011

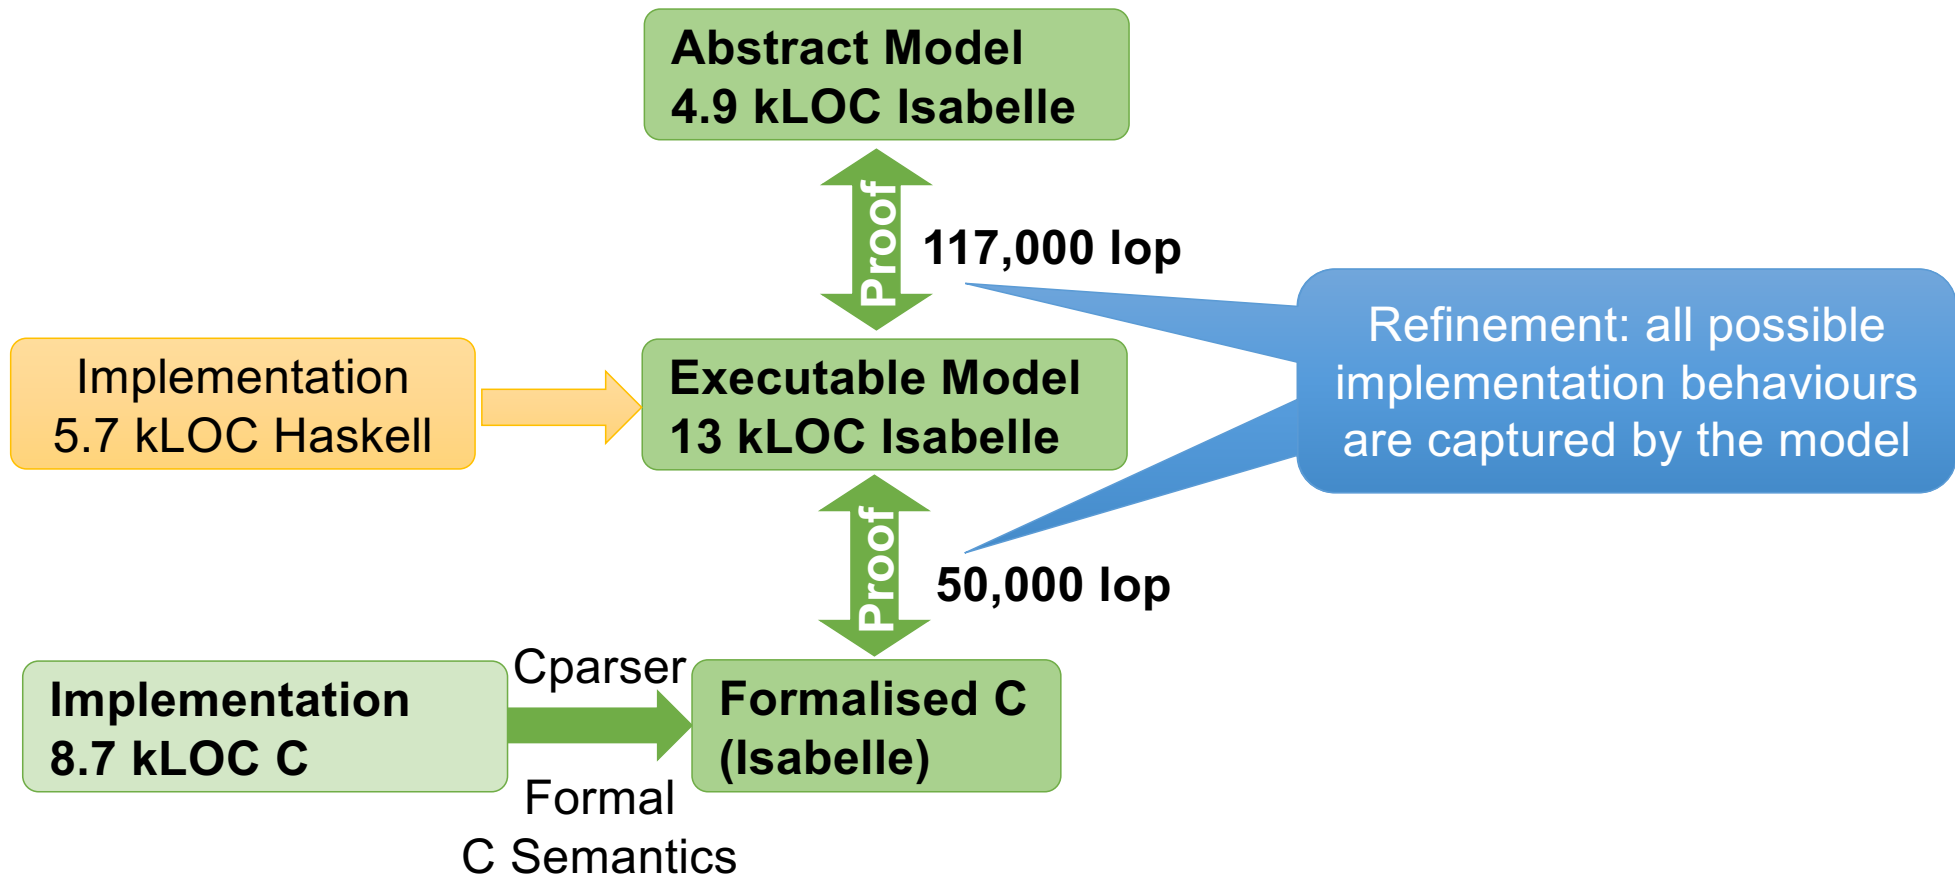


seL4 Proving Security and Safety (Armv6/7)



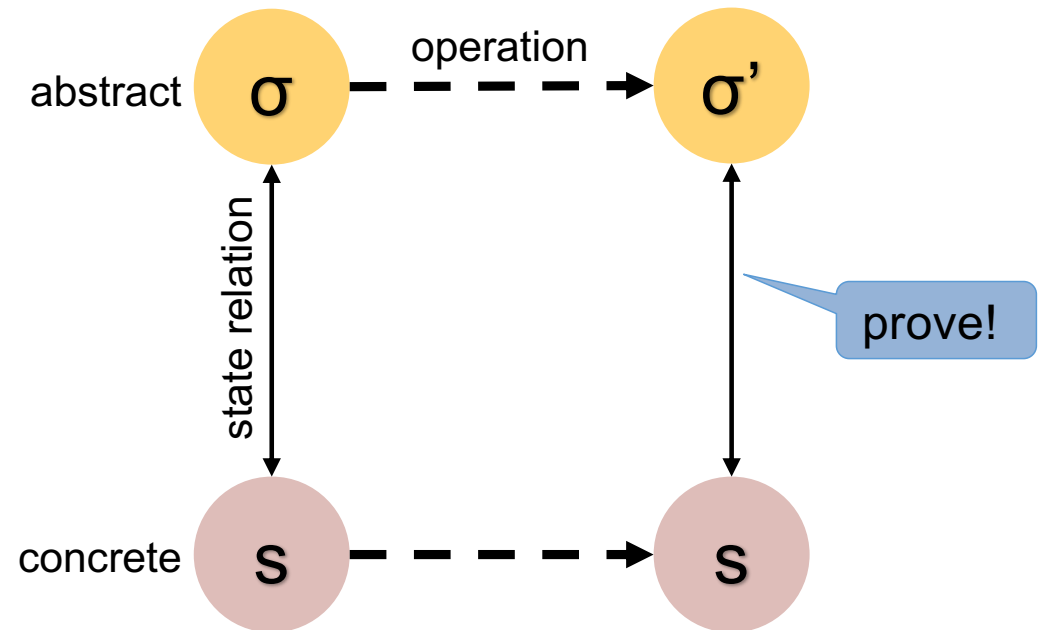
Functional Correctness

seL4 Proving Functional Correctness



seL4 How do these proofs work?

“Forward simulation”:
Prove state correspondence
of abstract and concrete levels



Functional Correctness Summary

Kinds of properties proved

- Behaviour of C code is fully captured by abstract model
- Behaviour of C code is fully captured by executable model
- Kernel never fails, behaviour is always well-defined
 - assertions never fail
 - will never de-reference null pointer
 - will never access array out of bounds
 - cannot be subverted by malformed input
- All syscalls terminate, reclaiming memory is safe, ...
- Well typed references, aligned objects, kernel always mapped...
- Access control is decidable

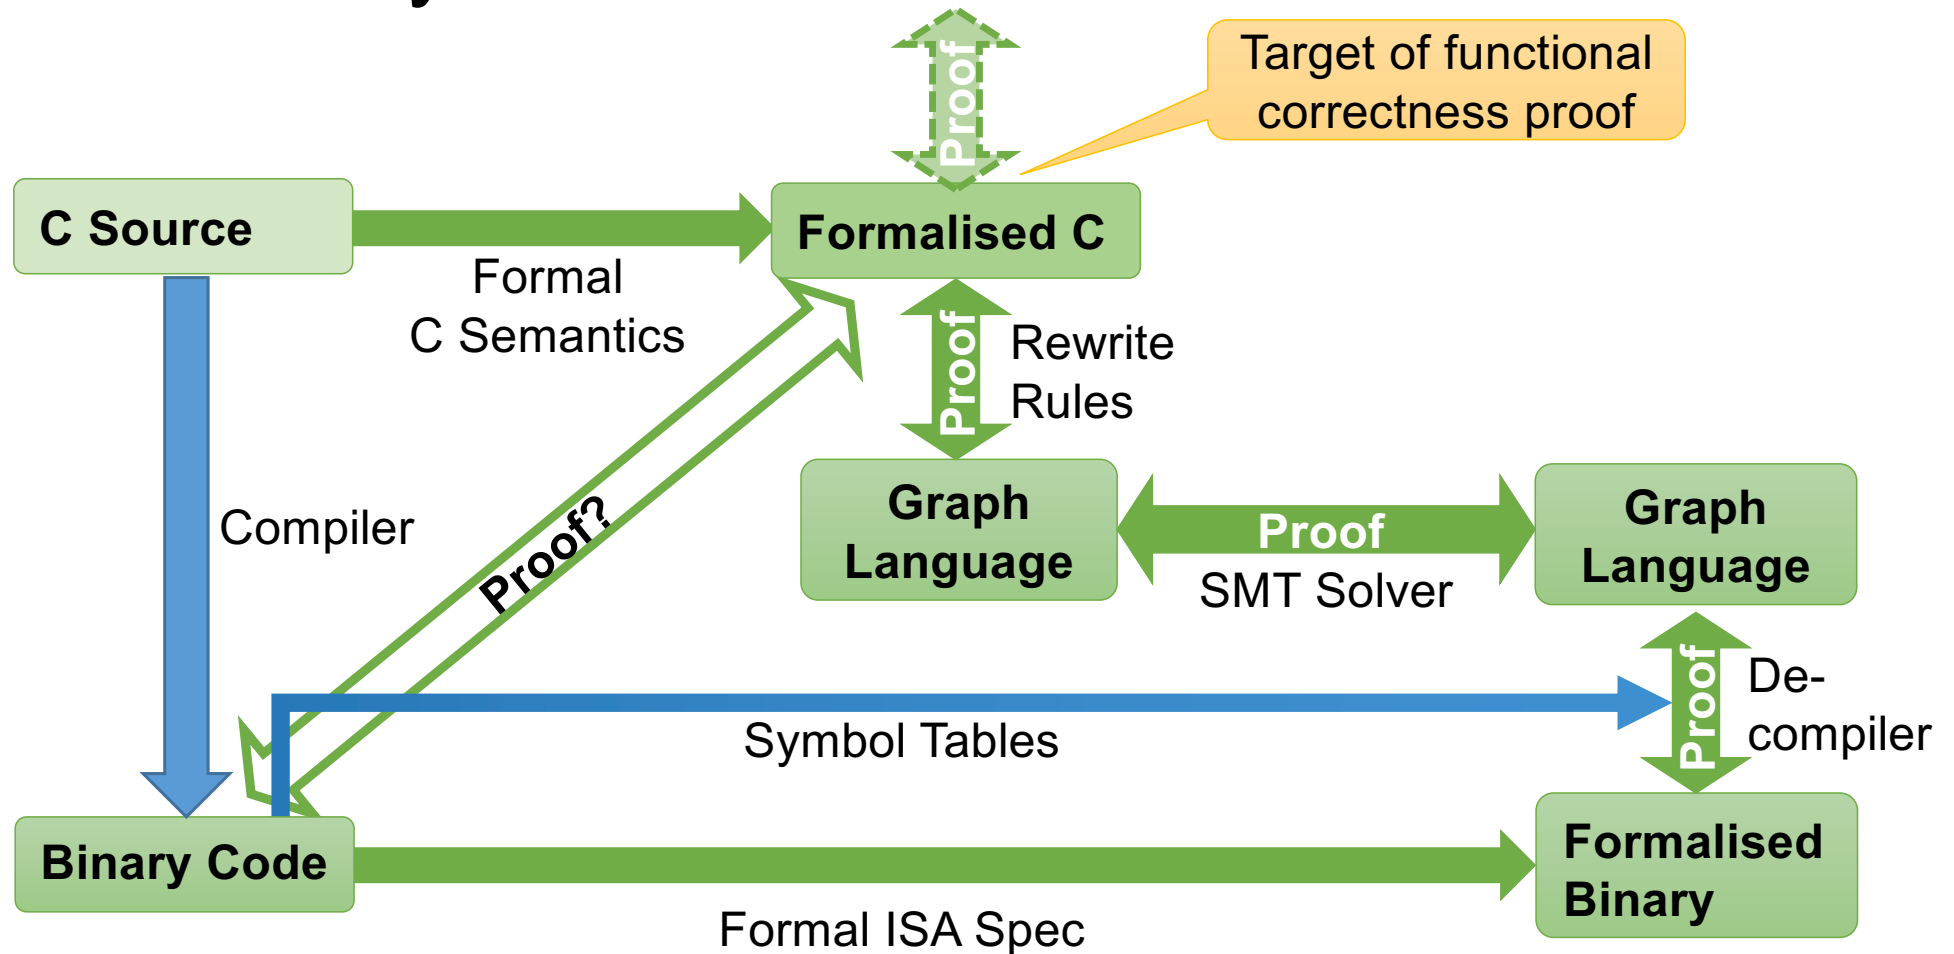
Can prove further properties on abstract level!

Bugs found:

- 16 in (shallow) testing
- 460 in verification
 - 160 in C,
 - 150 in design,
 - 150 in spec

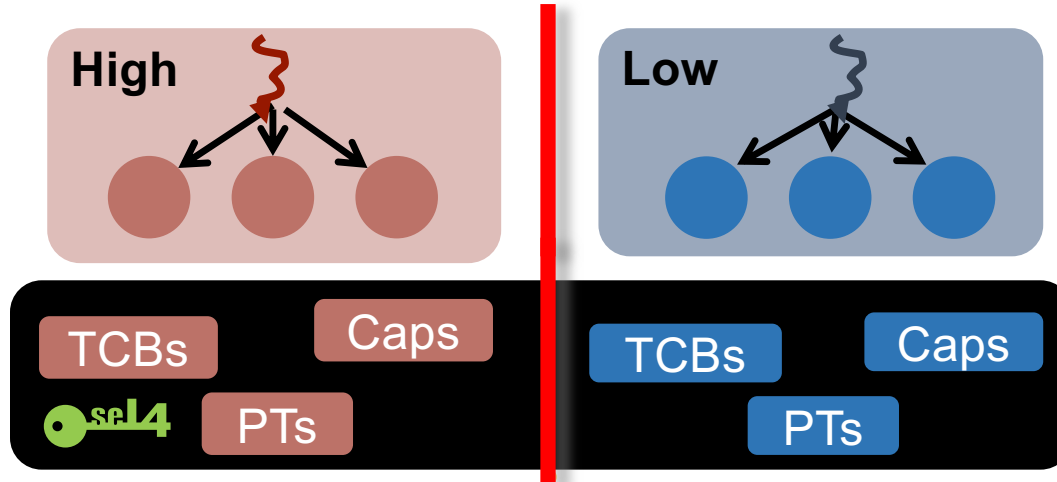
Binary Correctness

seL4 Binary Verification: Translation Validation



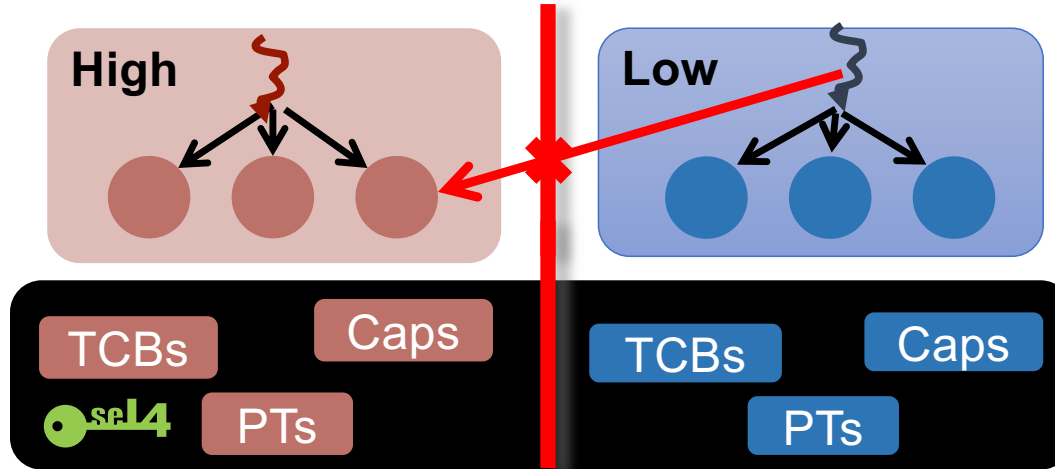
Security Enforcement

seL4 Isolation Goes Deep



Kernel data
partitioned
like user data

seL4 Integrity: Control Write Access



To prove:

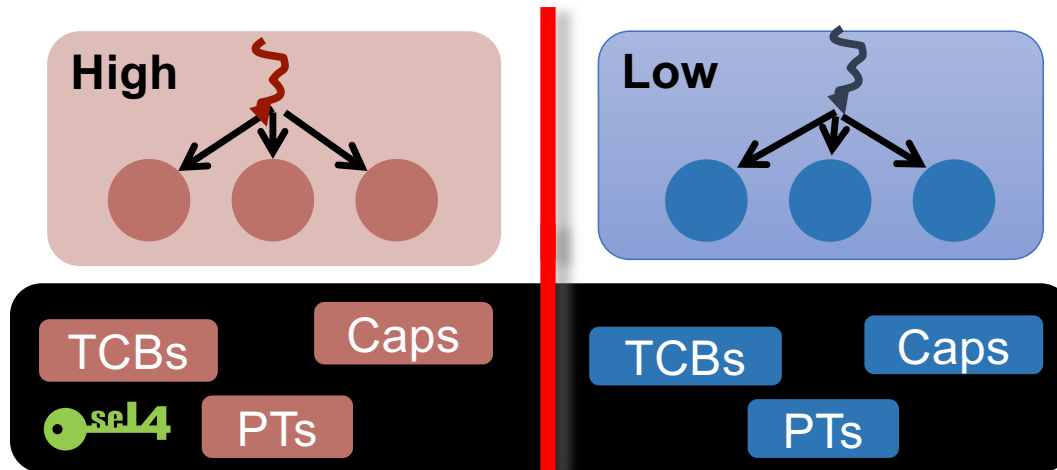
Low has no *write* capabilities to High objects
⇒ no action of Low will modify High state

Specifically, *kernel does not modify on Low's behalf!*

Event-based kernel
always operates on be-
half of well-defined user:

- Prove kernel only
modifies data if
presented write cap

seL4 Availability: Ensuring Resource Access

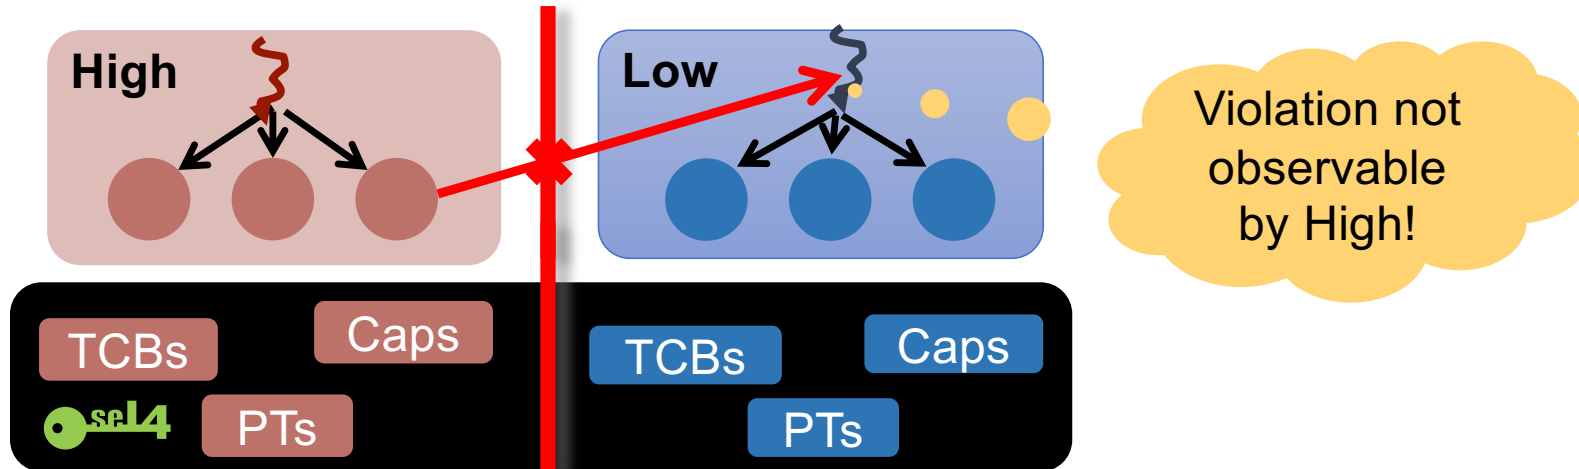


Nothing to do, implied by other properties!

Strict separation of kernel resources
⇒ Low cannot deny High access to resources



Confidentiality: Control Information Flow



Non-interference proof:

- Evolution of Low does not depend on High state
- Also shows absence of covert *storage channels*

To prove:

Low has no *read* capabilities to High objects
⇒ no action will reveal High state to Low



Confidentiality Proof Challenge

Spec

```
bool a();
```

```
bool b() {  
    int secret;  
}
```

Idiotic but valid refinement

Implementation

```
bool a() {  
    return !secret;  
}
```

Non-determinism
breaks confidentiality
under refinement!

Solution:

- Remove non-determinism where it affects confidentiality
- Eg: scheduler strictly round-robin

Infoword is very strong
property, requiring
restrictions rarely met
in real world

Limitations

seL4 Verification Assumptions

1. Hardware behaves as expected

- Formalised hardware-software contract (ISA)
- Hardware implementation free of bugs, Trojans, ...

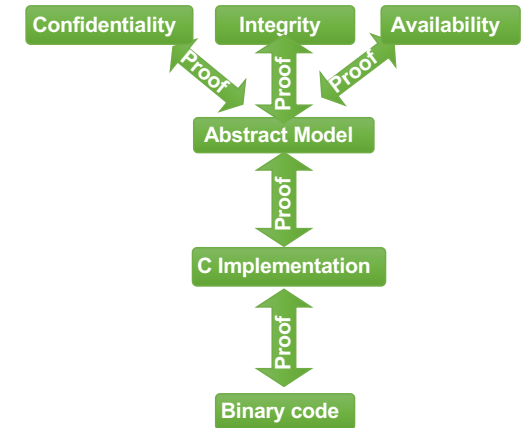
2. Spec matches expectations

- Can only prove “security” if specify what “security” means
- Spec may not be what we think it is

3. Proof checker is correct

- Isabel/HOL checking core that validates proofs against logic

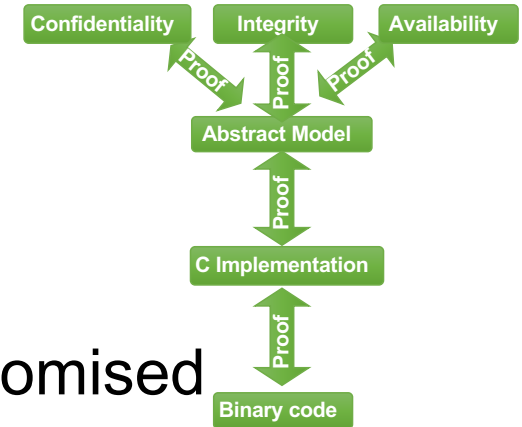
With binary verification do
not need to trust C compiler!



seL4 Present Verification Limitations

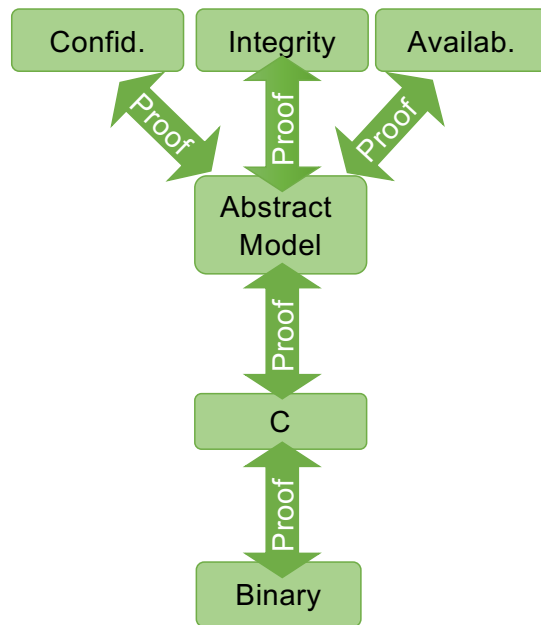
- Not verified boot code
 - **Assume** it leaves kernel in safe state
- Caches/MMU presently modeled at high level / axiomised
 - MMU model finished by recent PhD
- Not proved any temporal properties
 - Presently not proved scheduler observes priorities, properties needed for RT
 - WCET analysis applies only to dated ARM11/A8 cores
 - No proofs about timing channels

Present research!

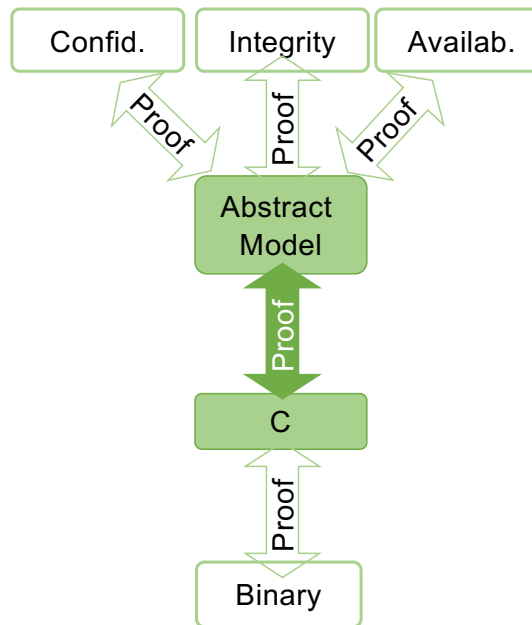


seL4 Present Status

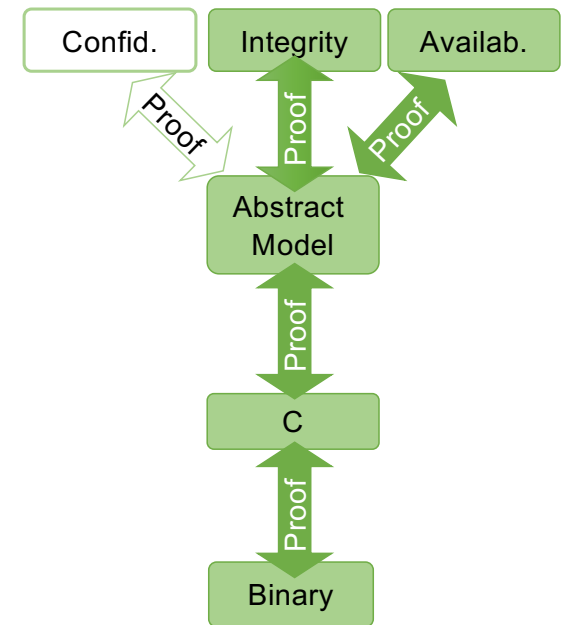
32b Arm




64b x86



64b RISC-V

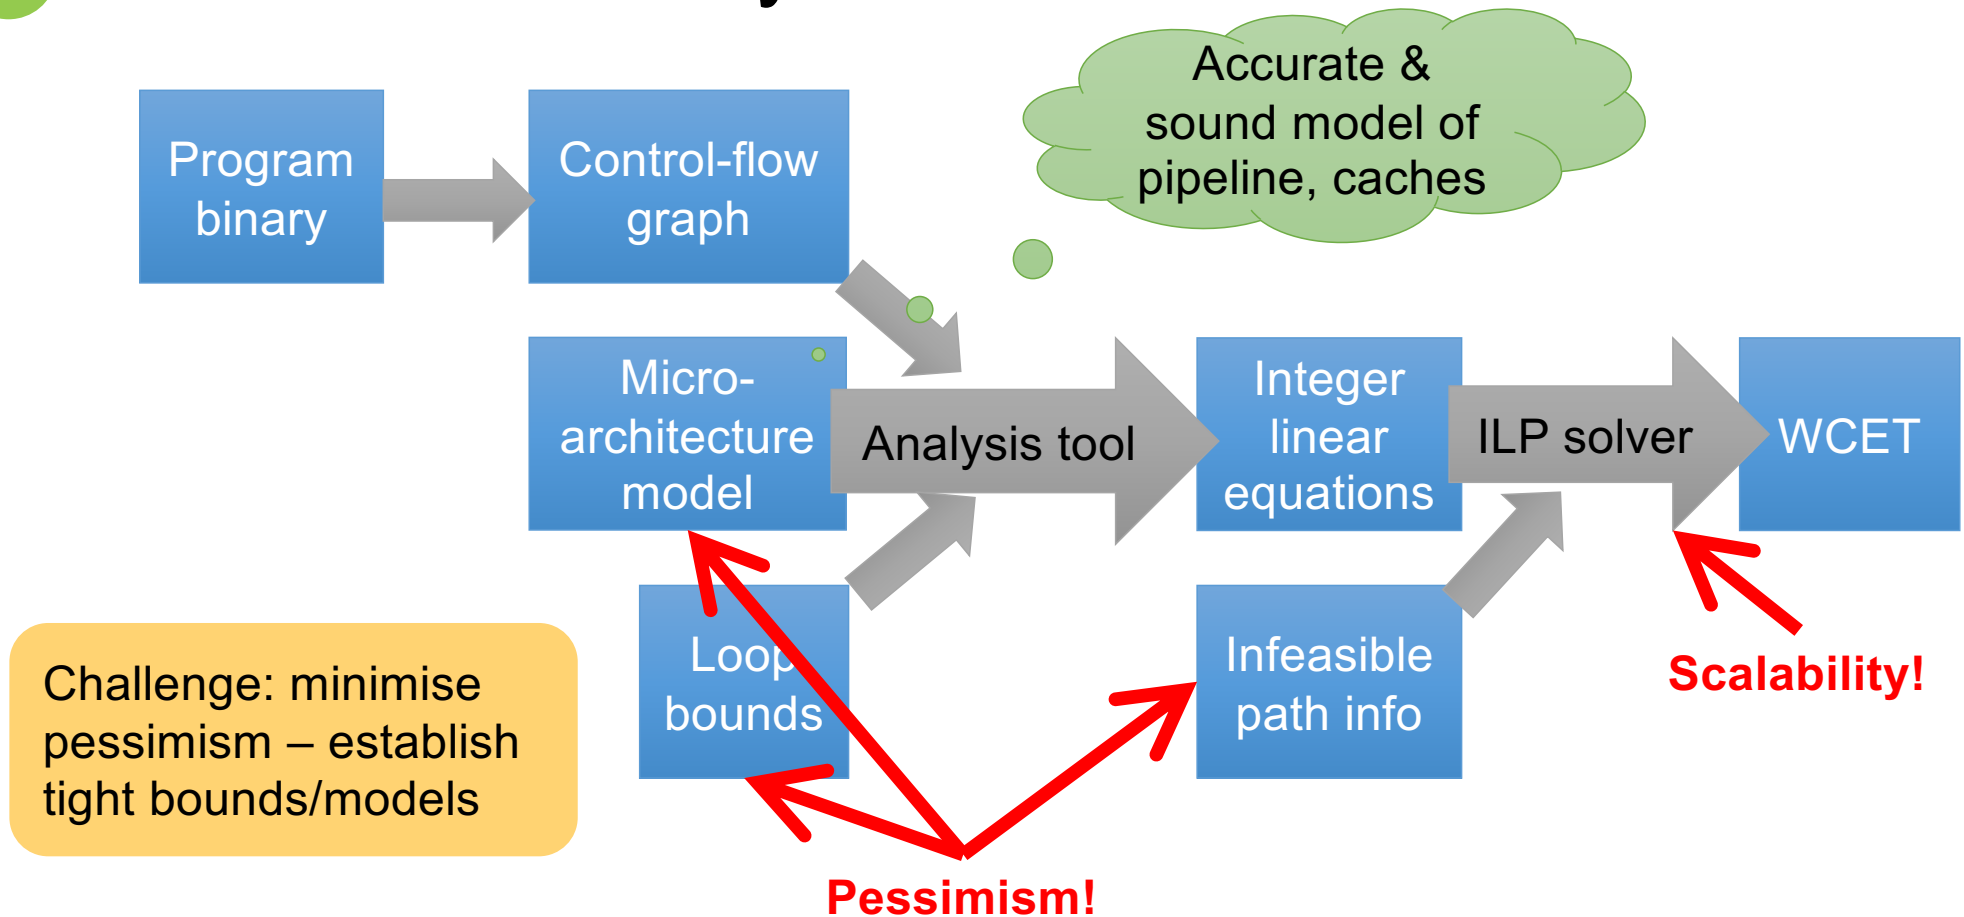


Common Criteria?

Level	Requirements	Specification	Design	Implementation
EAL1	not evaluated	Informal	not eval	not evaluated
EAL2	not evaluated	Informal	Informal	not evaluated
EAL3	not evaluated	Informal	Informal	not evaluated
EAL4	not evaluated	Informal	Informal	not evaluated
EAL5	not evaluated	Semi-Formal	Semi-Formal	Informal
EAL6	Formal	Semi-Formal	Semi-Formal	Informal
EAL7	Formal	Formal	Formal	Informal
	Formal	Formal	Formal	Formal

WCET Analysis

seL4 WCET Analysis

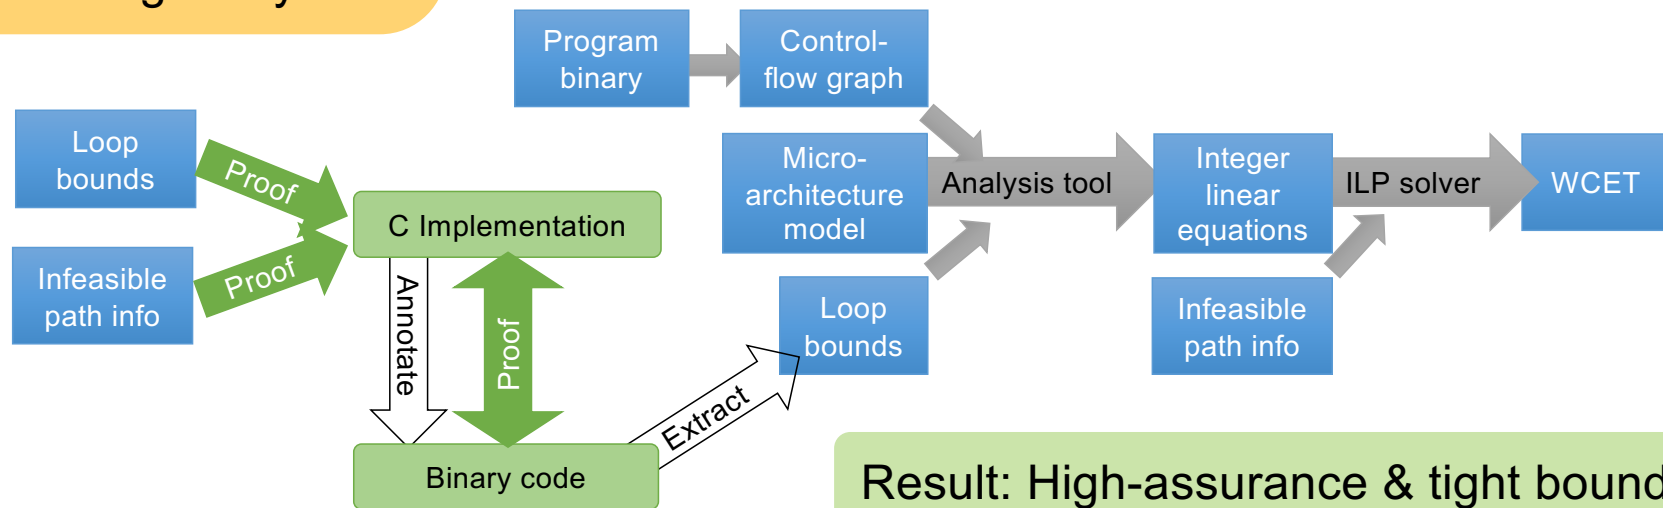


seL4 Loop Bounds & Infeasible Paths

Tight loop bounds and infeasible path refutations infeasible to obtain from binary – lack of semantic information, especially pointer aliasing analysis.

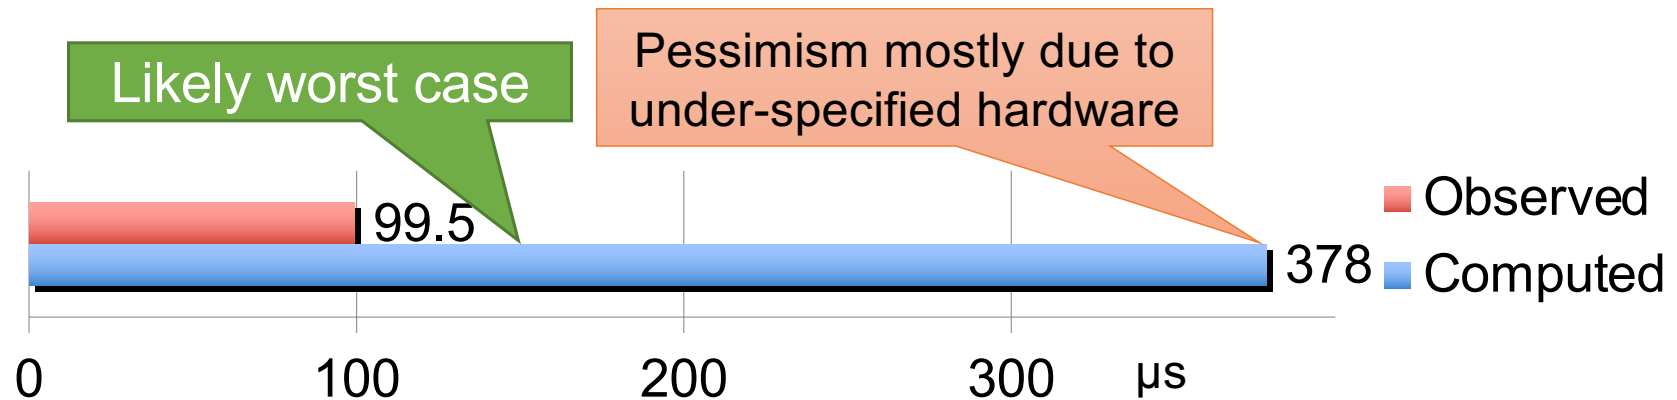
Idea:

- prove on C level
- transfer to binary using translation-validation toolchain



Result: High-assurance & tight bounds!

seL4 WCET Analysis on ARM11



WCET presently limited by verification practicalities

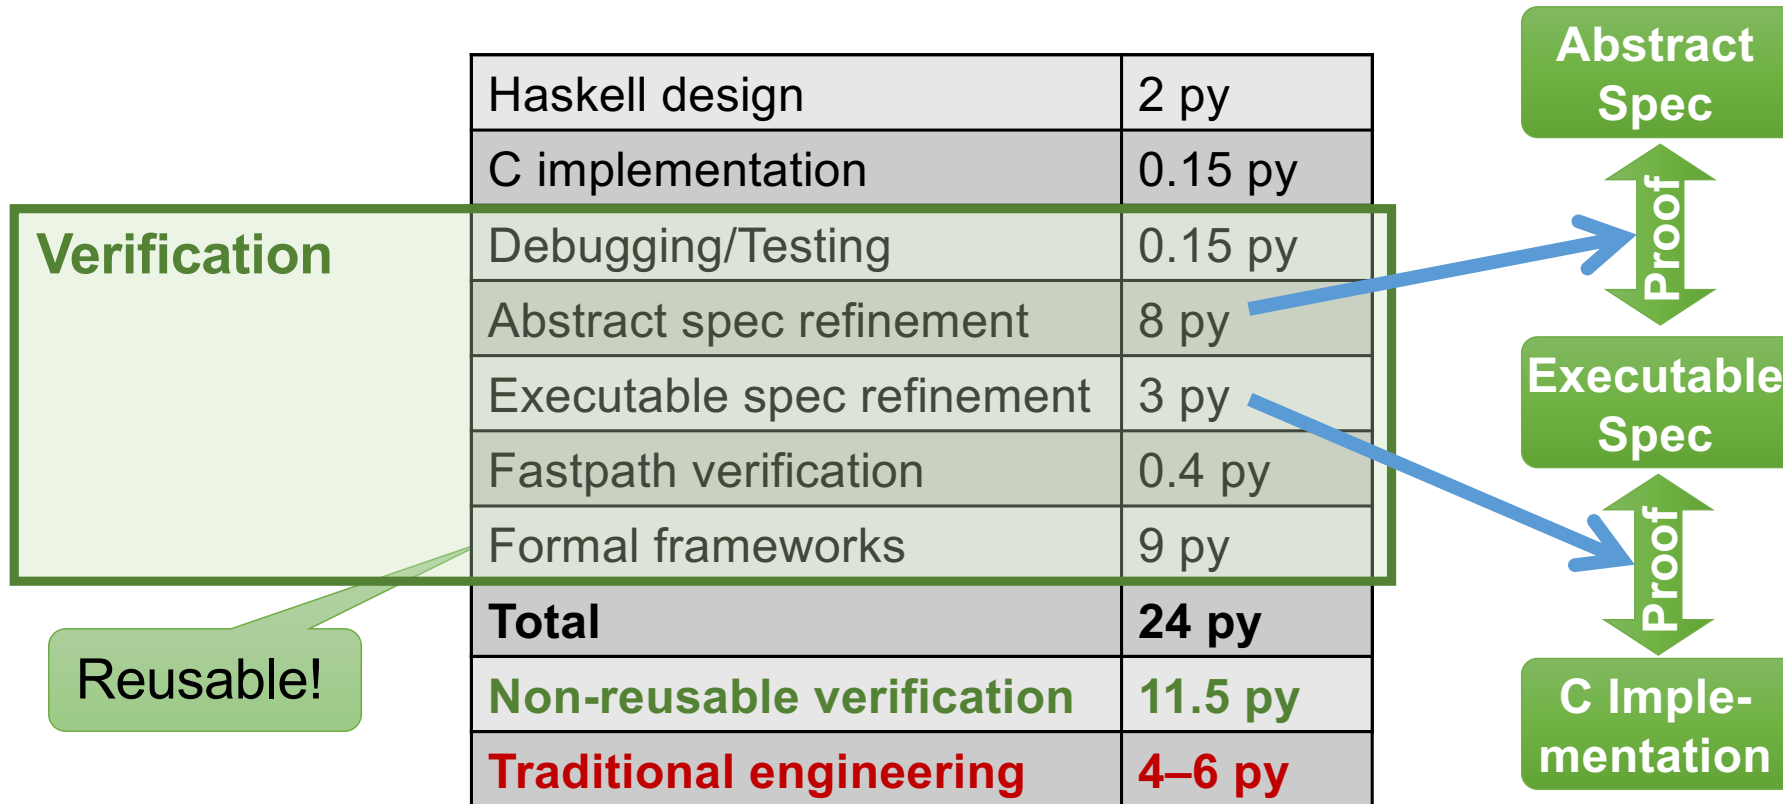
- without regard to verification achieved 50 μs
- 10 μs seem achievable
- BCET $\sim 1\mu\text{s}$
- [Blackham'11, '12] [Sewell'16]

Problem: Latency information no longer published by Arm!



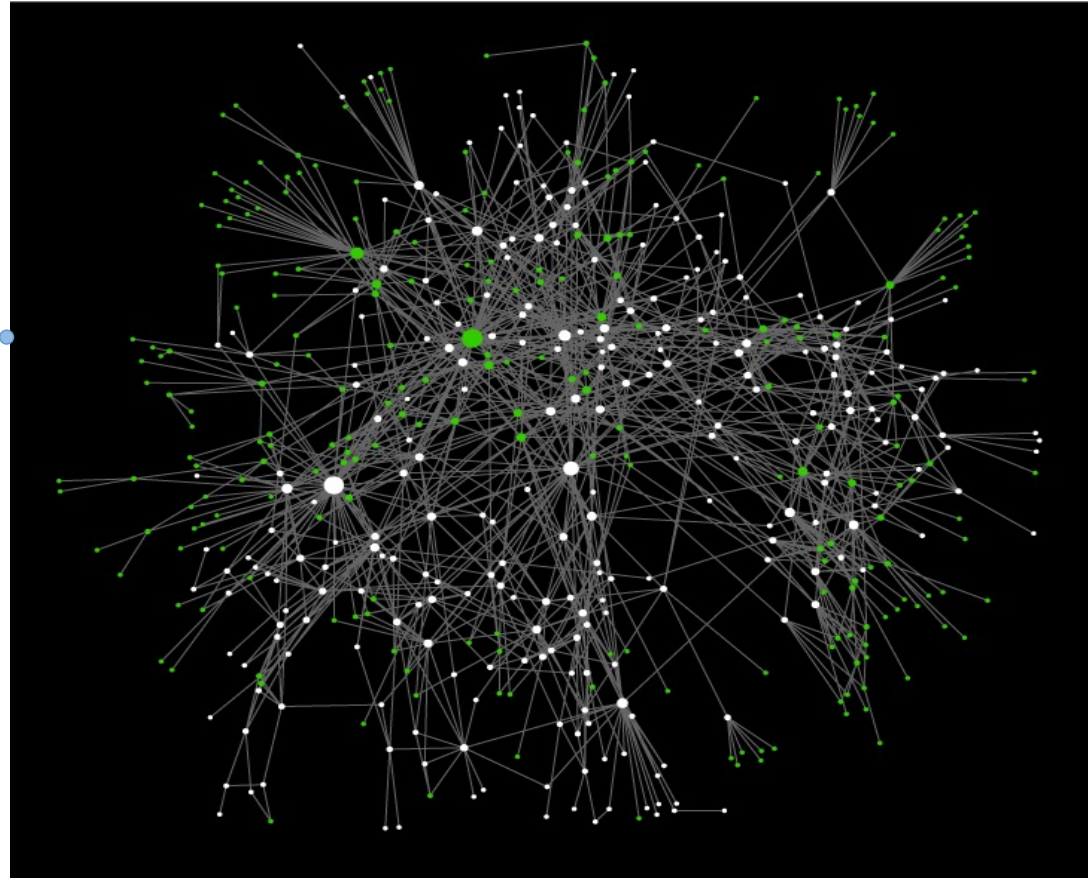
Cost of Verification

seL4 Verification Cost Breakdown

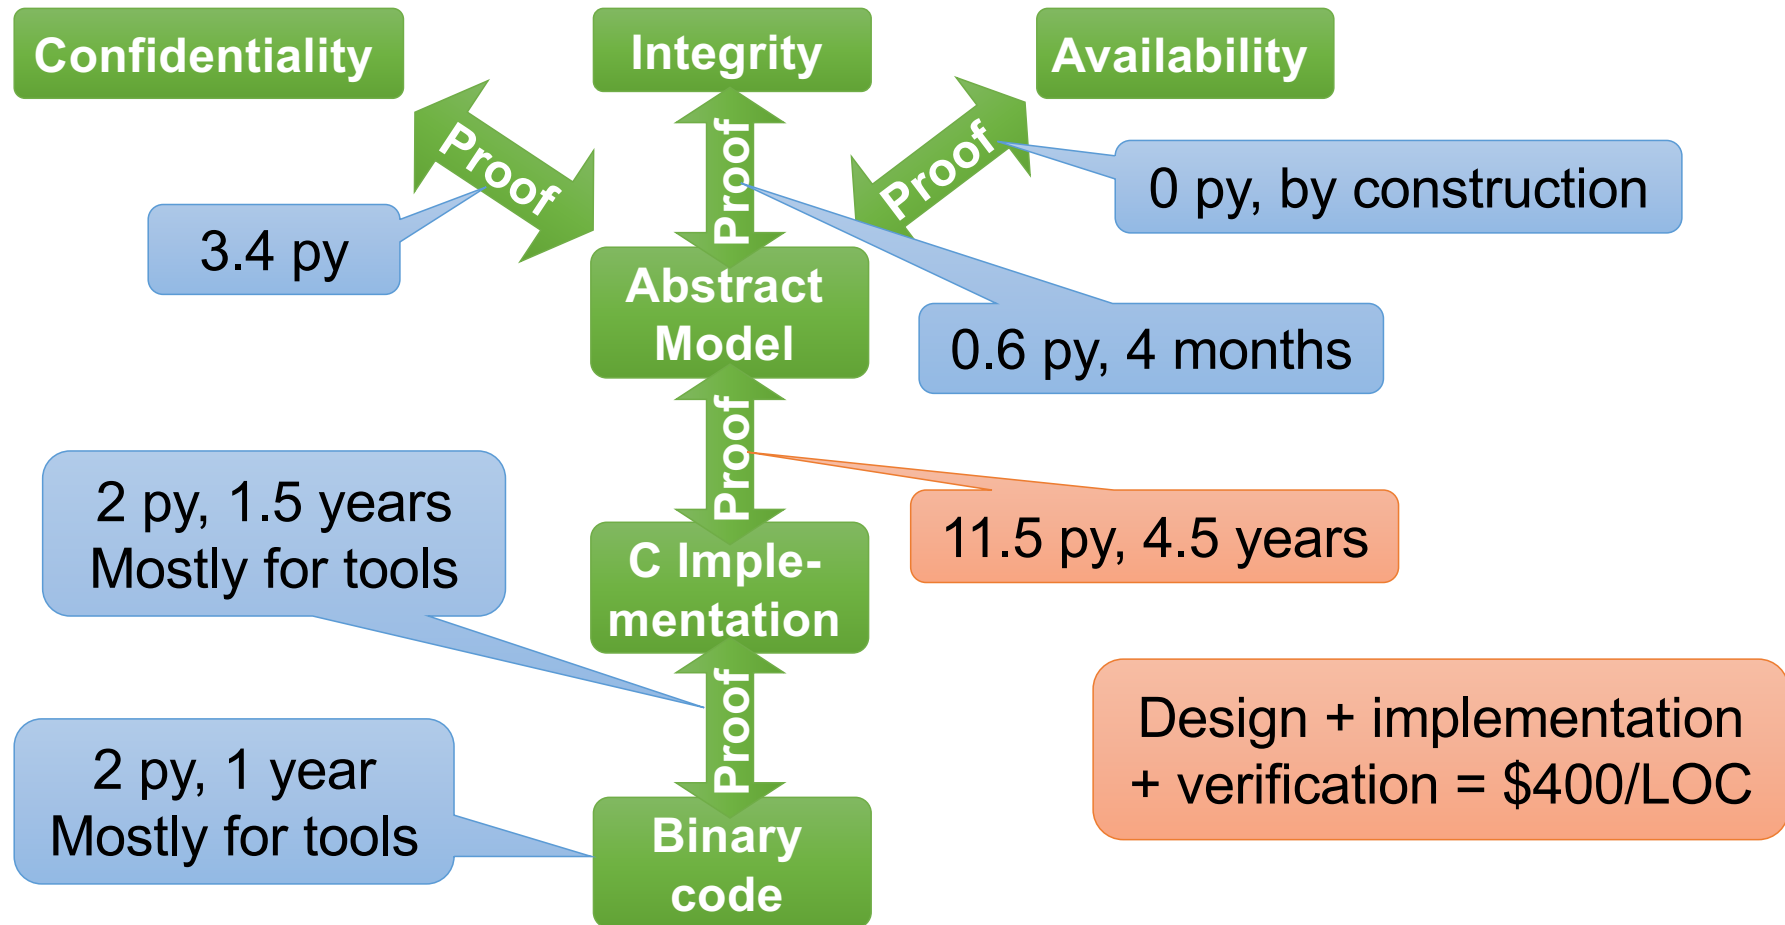


seL4 Why So Hard for 9,000 LOC?

seL4 call
graph



seL4 Verification Cost



seL4 Microkernel Life-Cycle Cost in Context

