

Proof Engineering Considered Essential

FM'14 Singapore

Gerwin Klein



Australian Government

Department of Broadband, Communications and the Digital Economy

Australian Research Council





NICTA Funding and Supporting Members and Partners

















Queensland

Windows

An exception 06 has occured at 0028:C11B3ADC in VxD DiskTSD(03) + 00001660. This was called from 0028:C11B40C8 in VxD voltrack(04) + 00000000. It may be possible to continue normally.

Press any key to attempt to continue.

 Press CTRL+ALT+RESET to restart your computer. You will lose any unsaved information in all applications.

Press any key to continue

A process or thread crucial to system operation has unexpectedly exited or been first his is the first time you've seen this stop error screen, instant your computer. If this screen appears again, follow these stops: The see to make sure any new hardware or software is properly installed. This is a new installation, ask your hardware or software manufered. To this is a new installation, ask your hardware or software manufered. This windows updates you wight need. To replace continue, disable or needed any newly firstalled manufered is software. Siable stops memory options such as components, restart is software, press fit to select advanced startup options, and the you computer, press fit to select advanced startup options, and the is stop: 0x0000004 (0x0000003, 0x0500000, 0x05000000, 0x0500000, 0x050000, 0x050000, 0x050000, 0x050000, 0x0500000, 0x050000, 0x05000, 0x05000, 0x050000, 0x05000, 0x0

A problem has been detected and windows has been shut down to prevent damage to your computer.



	-	
Sec. 1	-	
-		the state of
1917-24		-
Zearch 1	-	
1.1		1
-		

•





Isolation

Isolation is the Key

Trustworthy Computing Base

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

Applications

- fault isolation
- fault identification
- IP protection
- modularity

Trusted next to Untrusted





NICTA

Isolation is the Key

Trustworthy Computing Base

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

Applications

- fault isolation
- fault identification
- IP protection
- modularity

Trusted next to Untrusted





NICTA





SOL

Functional Correctness Possible





Functional Correctness Possible





Functional Correctness Possible





*conditions apply









Proof Architecture [SOSP'09]



NICTA

Proof Architecture Now





Proof Architecture Now









First and only kernel with proof of integrity and confidentiality enforcement – at binary level



First and only kernel with proof of integrity and confidentiality enforcement – at binary level

World's fastest microkernel on ARM architecture



First and only kernel with proof of integrity and confidentiality enforcement – at binary level

World's fastest microkernel on ARM architecture

Predecessor deployed on 2 billion devices



First and only kernel with proof of integrity and confidentiality enforcement – at binary level

World's fastest microkernel on ARM architecture

Predecessor deployed on 2 billion devices

First and only protected-mode operating-system with complete and sound timing analysis

seL4: Unique Assurance



First and only protected-mode operating-system with complete and sound timing analysis



- Functional correctness: 12 person years
- Integrity+Confinement: 10 person months
- Non-interference: 48 person months
- Binary Verification: automatic





• Industry Best Practice:

- High assurance (Common Criteria EAL 6+):
 \$1,000/LOC, model verification + testing, unoptimised
- Low assurance (traditional embedded kernels): \$100–200/LOC, 1–5 faults/kLOC, optimised



• Industry Best Practice:

- High assurance (Common Criteria EAL 6+):
 \$1,000/LOC, model verification + testing, unoptimised
- Low assurance (traditional embedded kernels): \$100–200/LOC, 1–5 faults/kLOC, optimised

State of the Art – seL4:

- \$400/LOC, binary-level formal proof, optimised
- Estimate repeat would cost half
 - about as much as unverified predecessor Pistachio!
- Aggressive optimisation [APSys'12]
 - much faster than traditional high-assurance kernels
 - as fast as best-performing low-assurance kernels



timised

• Industry Best Practice:

- High assurance (Common Criteria EAL 6+):
- \$1,000/' • Low assurar

\$100–20

State of th

• \$400/LOC,

Formal verification getting close to traditional kernel development.

- Estimate repeat would cost half
 - about as much as unverified predecessor Pistachio!
- Aggressive optimisation [APSys'12]
 - much faster than traditional high-assurance kernels
 - as fast as best-performing low-assurance kernels



Industry Best Practice:

• High assurance (Common Criteria EAL 6+):



Still too expensive for large-scale user code development.

Automation, Synthesis, Proof Generation

Next Step: Full System Assurance

DARPA HACMS Program:

• Provable vehicle safety

Boeing Unmanned

Little Bird (AH-6)

NICTA

• Red Team must not be able to divert vehicle



NICTA

SMACCMcopter Research Vehicle

> UNIVERSITY OF MINNESOTA

Scale 500 k -400 k -300 k -200 k -100 k -0 AFP entries by submission date ■ four-color theorem, Isabelle/HOL, CompCert Odd Order Theorem, L4.verified, Verisoft тΟ

Proof Introspection

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

NICTA Copyright 2014

19

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.



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 400kloc proof jungle? **Hard!**





Development of seL4 code + spec artefacts (sloc)





NICTA Copyright 2014

• automating mechanical tasks, custom tactics

• proof craft

- proof development

Proof Development

• custom proof calculus,

• decomposition of proofs over people,





- proof development

- decomposition of proofs ov
- custom proof calculus,
- automating mechanical task
- 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 ov
- custom proof calculus,
- automating mechanical task
- 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

- decomposition of proofs ov
- custom proof calculus,
- automating mechanical task
- 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 editcheck cycle; for that, I am even willing to (temporarily) sacrifice soundness.



23

– proof maintenance

Problems of Scale

- changes, updates, new proofs, new features
- automated regression, keep code in sync
- refactoring
- simplification





23

Problems of Scale

- proof maintenance

- changes, updates, new proofs, new features
- automated regression, keep code in sync
- refactoring
- simplification

Dan's Conclusion

Verification is fast, maintenance is forever.





Research Challenges

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







• But: New Properties and Problems





• But: New Properties and Problems

- Results are checkable
 - You know when you are done!
 - No testing
 - 95% proof: no such thing





- 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





- 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



- 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





Deliver within Time and Budget



• Estimation:

- time and effort
- how precisely, with which confidence?
- how early?





NICTA Copyright 2014

Deliver within Time and Budget

• Estimation:

- time and effort
- how precisely, with which confidence?
- how early?
- Size of artefacts
 - easier to predict?
 - related to effort?







Deliver within Time and Budget

• Estimation:

- time and effort
- how precisely, with which confidence?
- how early?
- Size of artefacts
 - easier to predict?
 - related to effort?
- Complexity
 - from initial artefacts?
 - which influence?







• User Interface

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

Example.thy	
Example.thy (~/)	isabelle 🗘
theory Example	Filter: 🏷
imports Base	Example.thy
begin	theory Example theory Example
inductive path for R :: "'a \Rightarrow 'a \Rightarrow bool" where	 Inductive path for K :: "a ⇒ a ⇒ theorem example:
base: "path R x x"	end
step: "R x y \implies path R y z \implies 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 Pyz	-
ultimately show "P x z" by auto	
qed	
end	
 Output Prover Session Raw Output 	
,1 (35/405) (isab	elle,sidekick,UTF-8-Isabelle)NmroUG46/120Mb 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

O O Example.thy		
Example.thy (~/)	🗘 🗿 🖪 🛛 isabelle	\$
theory Example	Filter:	
imports Base	Example thy	V
begin	▼ theory Example	
	theory Example	-
inductive path for R :: "'a \Rightarrow 'a \Rightarrow bool" where	 theorem example: 	a
base: "path R x x"	end	
step: "R x y \implies path R y z \implies path R x z"		
theorem example:		
<pre>fixes x z :: 'a assumes "path R x z" shows "P x z"</pre>		
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		
1		
 Output Prover Session Raw Output 		• •
(isabelle	lle sidekick UTE-8-Isabelle)Nm r o UC 46/120Mb 3:3	(8 P



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?



• Are there Proof Engineering Laws?





- Are there Proof Engineering Laws?
 - Proofs always become larger and more complex over time. (from Cope's rule)





- Are there Proof Engineering Laws?
 - Proofs always become larger and more complex over time. (from Cope's rule)
 - Adding manpower to a late proof project makes it later. (from Brooks' law)





- Are there Proof Engineering Laws?
 - Proofs always become larger and more complex over time. (from Cope's rule)
 - Adding manpower to a late proof project makes it later. (from Brooks' law)
 - You cannot reduce the complexity of a given proof beyond a certain point. Once you've reached that point, you can only shift the burden around.
 (from Tesler's law)





- Are there Proof Engineering Laws?
 - Proofs always become larger and more complex over time. (from Cope's rule)
 - Adding manpower to a late proof project makes it later. (from Brooks' law)
 - You cannot reduce the complexity of a given proof beyond a certain point. Once you've reached that point, you can only shift the burden around.
 (from Tesler's law)
 - Are they true?



Summary



seL4

- Full verification. Full performance.
- Already cost effective for high assurance.
- Going open source and open proof in 2014.



Summary



seL4

- Full verification. Full performance.
- Already cost effective for high assurance.
- Going open source and open proof in 2014.



Proof Engineering

- Should become a research discipline.
- Work has started. A lot more to be done.


http://sel4.systems





From imagination to impact