UNSW SYDNEY | Australia's Global University

School of Computer Science & Engineering

Trustworthy Systems Group

The seL4 Microkernel: Provable Security for the Real World

Gernot Heiser

UNSW and seL4 Foundation

gernot@unsw.edu.au

Cyberattacks Are Everywhere

BITSIGHT

Report Shows Cyber Attacks on Cloud Services Have Doubled

Cyberattacks on Automated Vehicles Rise by 99%: Report

By CISOMAG - June 9, 2020

Cyber Attacks That Target Electrical Devices and Equipment: What Engineers Should Know

February 10, 2020 by Ikimi .O

News / World

'Most serious cyberattack of the Ukraine war': Tens of thousands modems crippled

AP By Associated Press | 5:38pm Mar 31, 2022

Increasingly used by
- organised crime
- state actors

NEWS | February 7, 2022

Ransomware attack on Swissport causes delay at Zurich Airport

Core Problem: Complexity

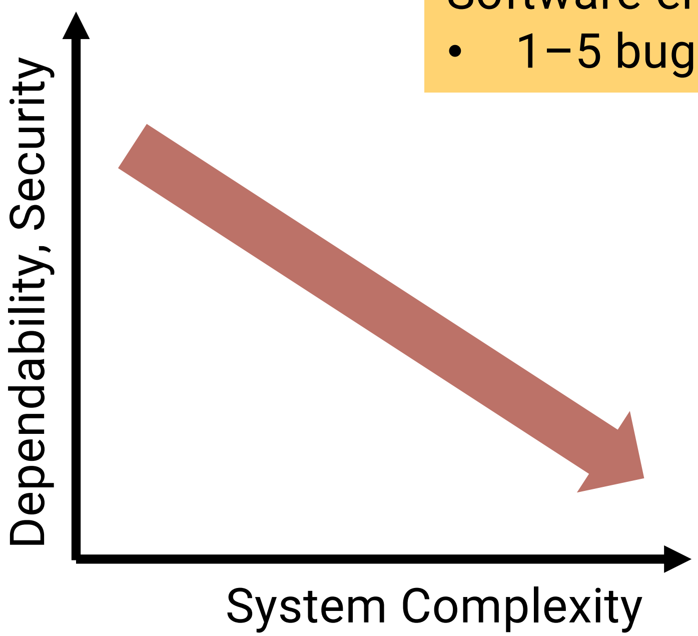

Software-engineering rule of thumb:
- 1–5 bugs per 1,000 lines of quality code

Bluetooth protocol stack:
100s kSLOC

Linux/Windows kernel:
10s MSLOC

Standard Approach: Patch-and-Pray

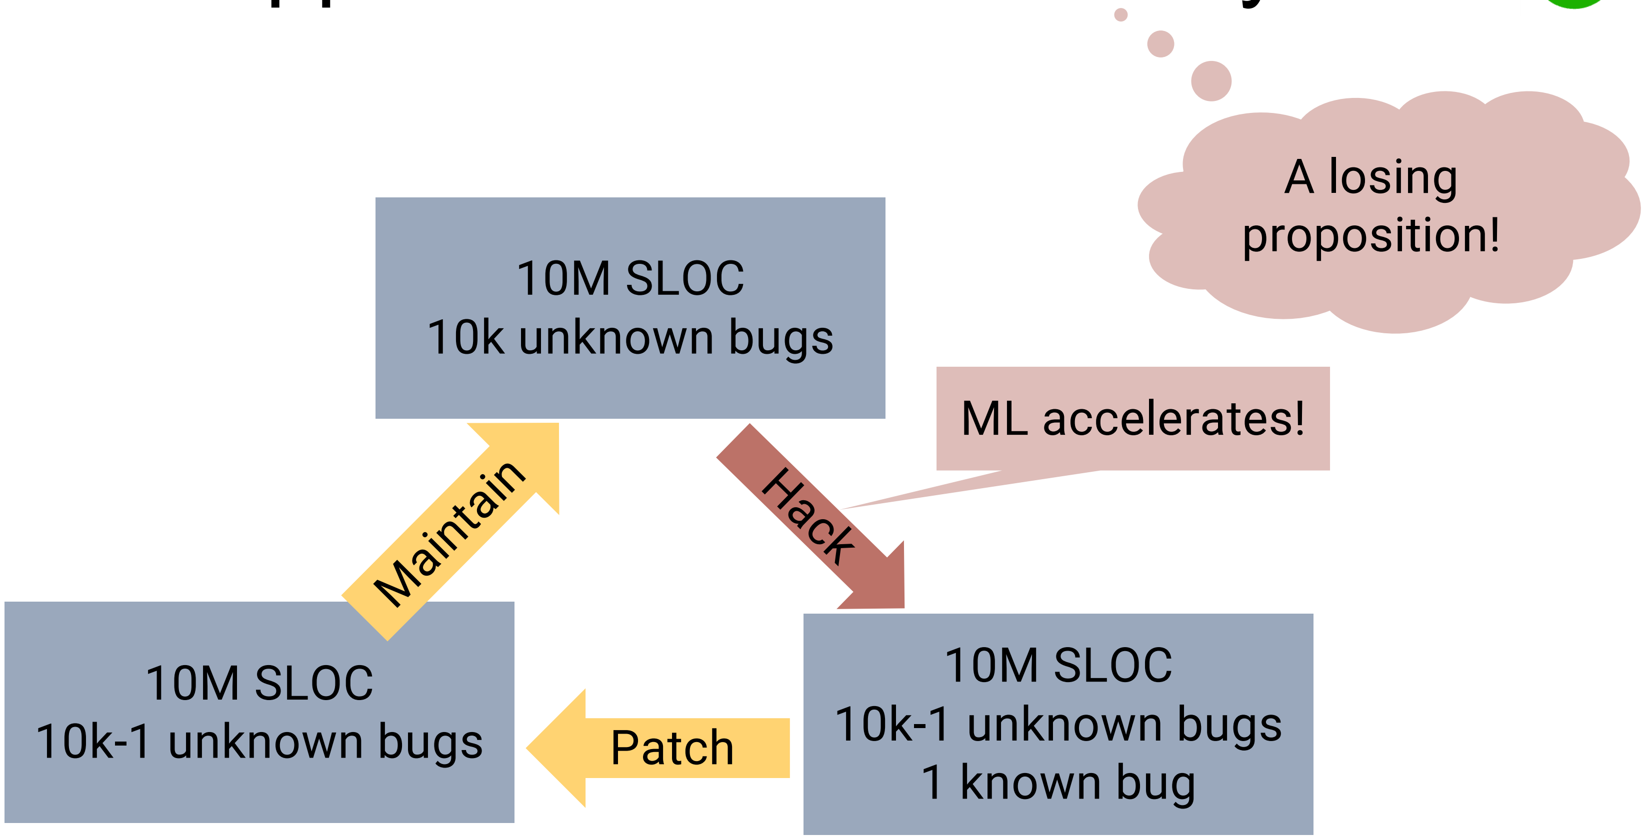

How Can We Do Better?

Step 1: Minimise Trusted Computing Base

Modularisation: Separate functions
- operating-system services
- applications

Microkernel enforces isolation
- kernel code reduced to minimum
- mediates hardware resources
- performance critical

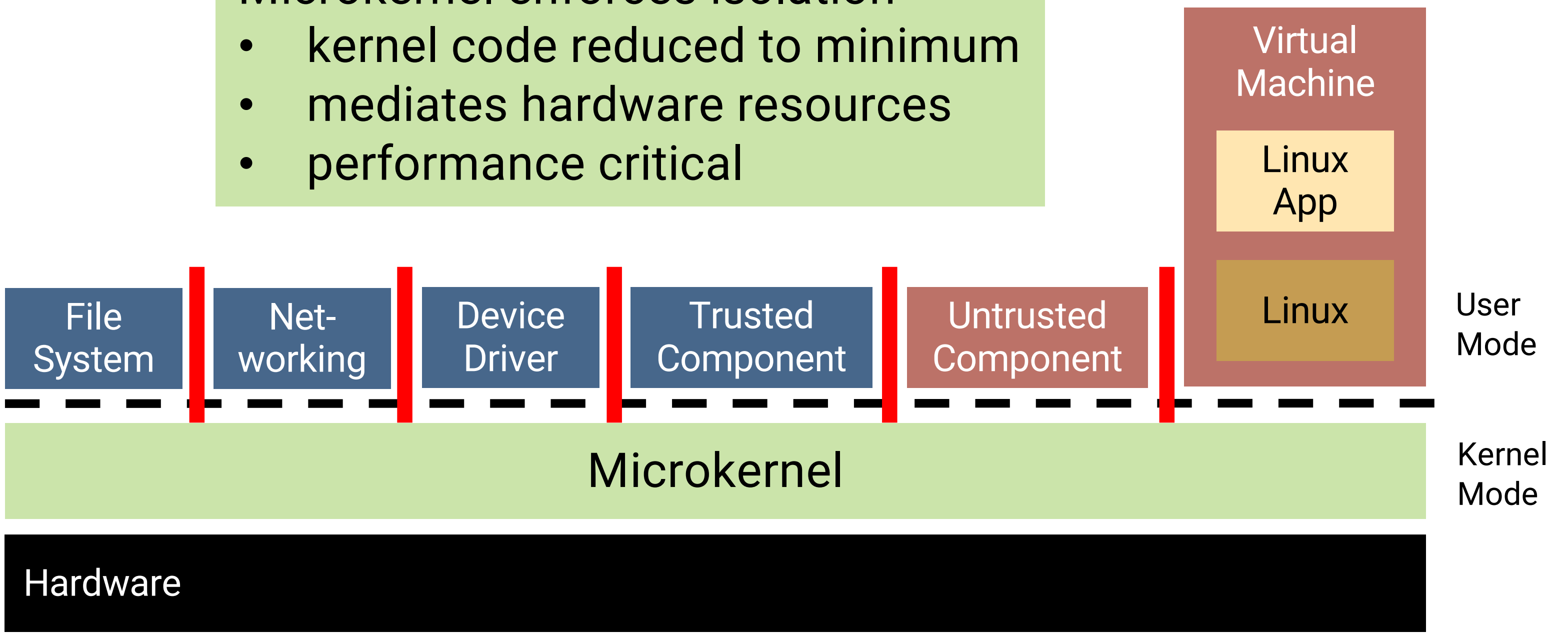

Step 2: Mathematical Proof

- First OS with proof of implementation correctness
- Only verified OS with fine-grained protection (capabilities)
- Only protected-mode RTOS with sound and compete WCET analysis
- World's fastest microkernel

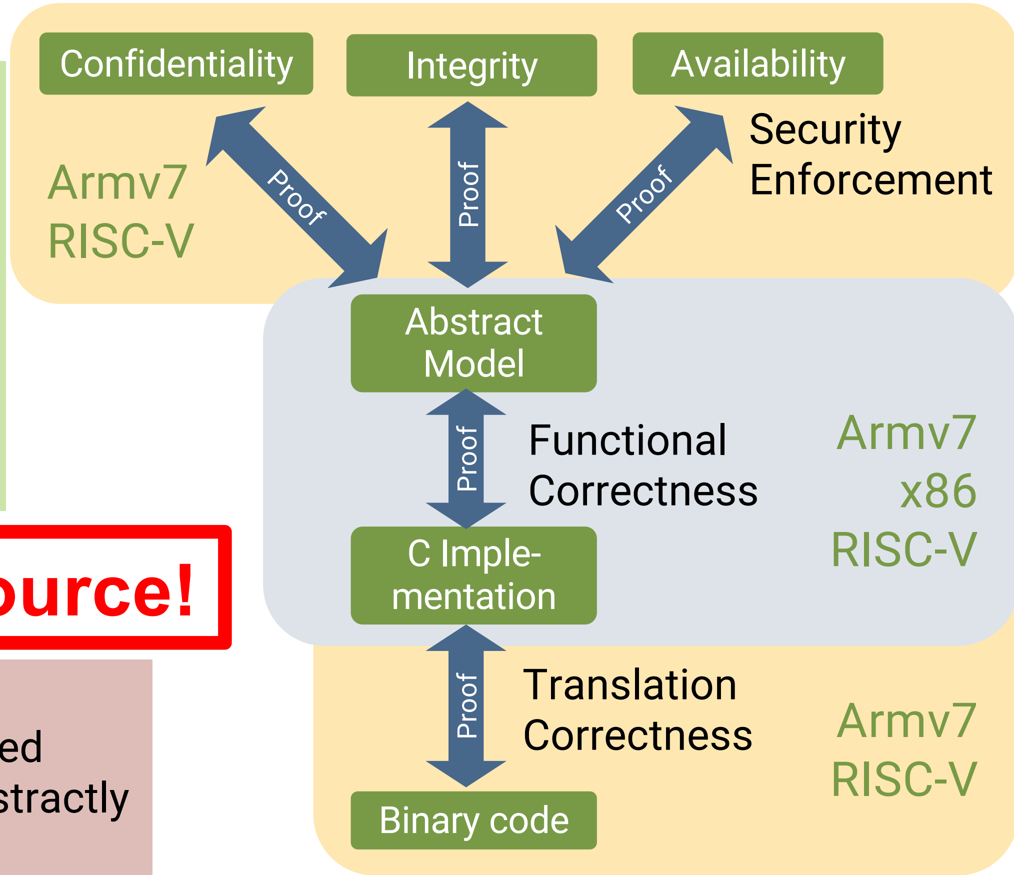

Open Source!

Present limitations

- initialisation code not verified
- MMU, caches modelled abstractly
- Multicore not yet verified

How Do Refinement Proofs Work?

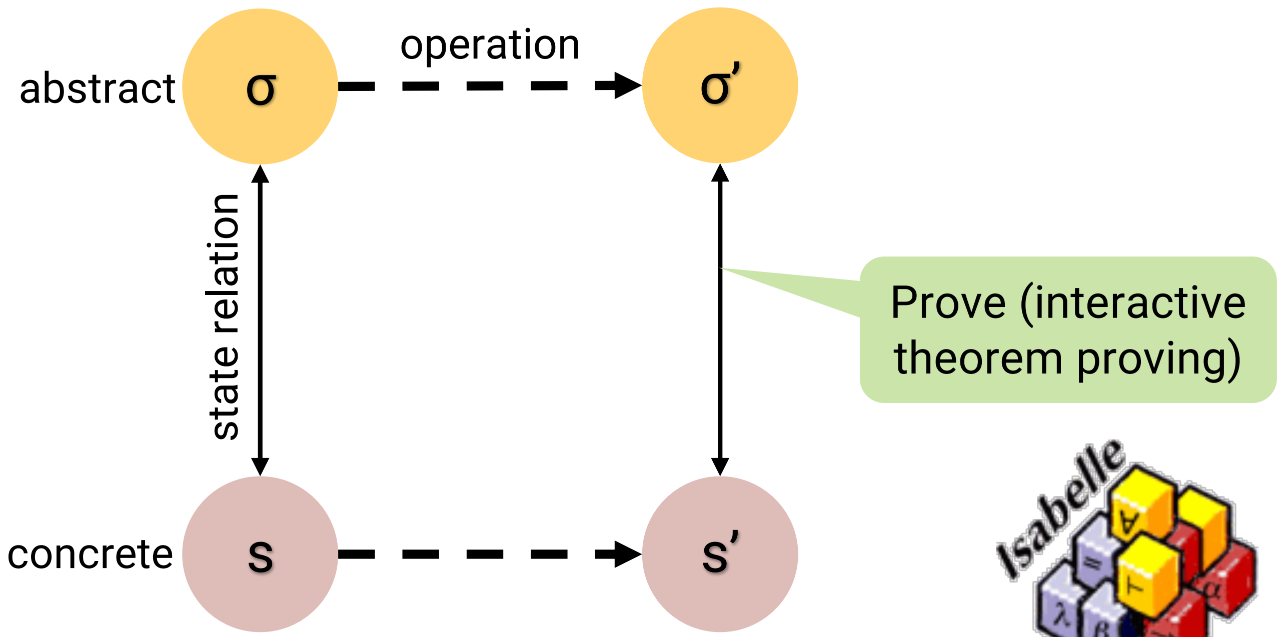

"Forward simulation":
Prove state correspondence of abstract and concrete levels

What Does Verification Mean?

Kinds of properties proved for functional correctness

- Behaviour is fully captured by abstract 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 mis-formed input
  - ✓ …

Can prove further properties on abstract level!

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

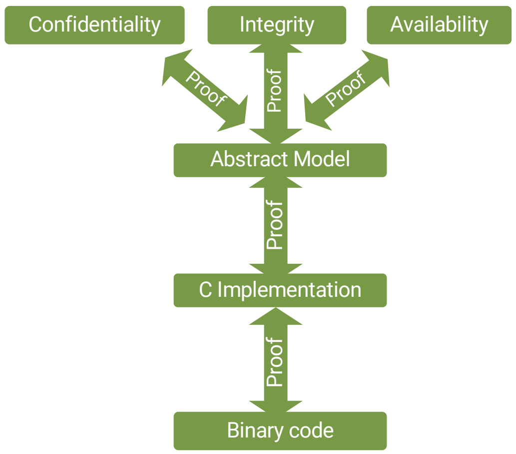

With binary verification do not need to trust the C compiler!

Minimise Trusted Computing Base

Modularisation: Separate components
- operating-system serv
- applications

Microkernel enforces isolation
rnel code reduced to minimum
tes hardware resources
rformance critical

Verification makes isolation bullet-proof!

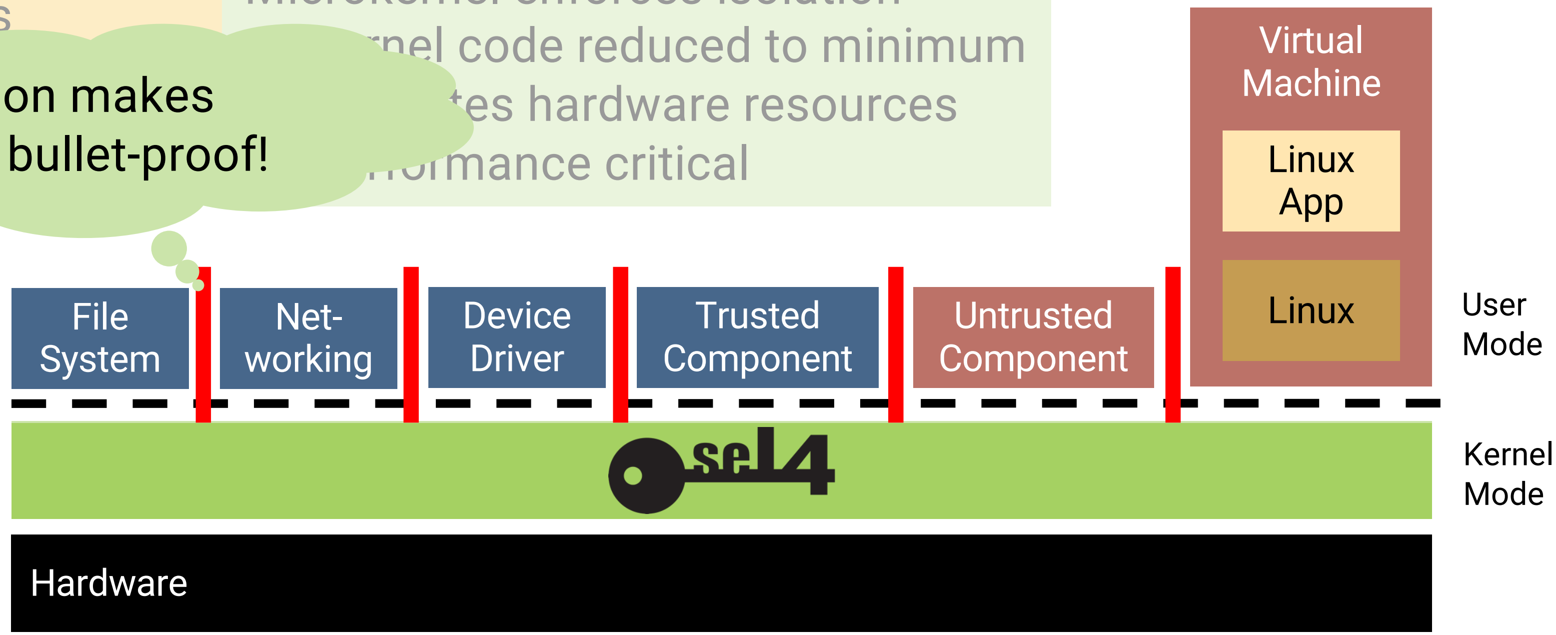

Capabilities: Fine-Grained Protection

- Enforce *least privilege*
- No communication unless explicitly authorised!

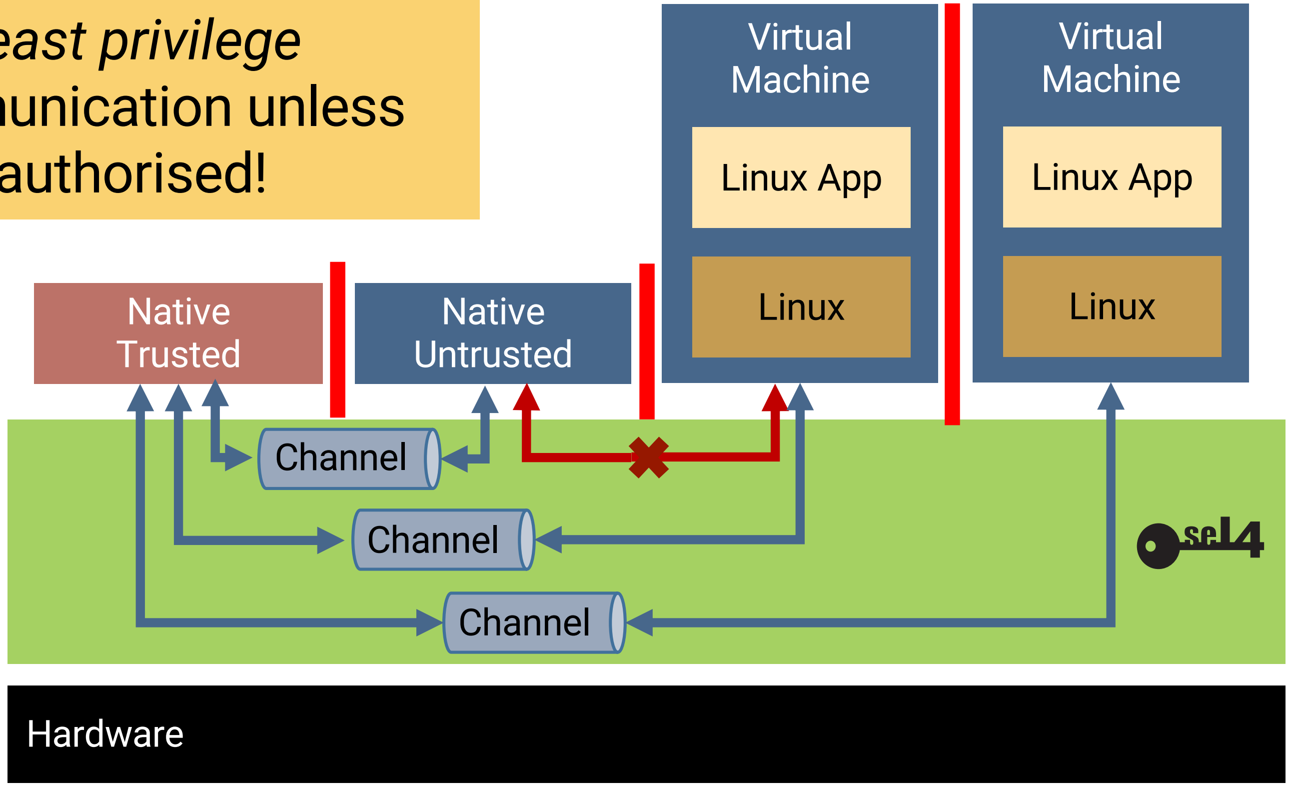

Real-Time: Capabilities for Time

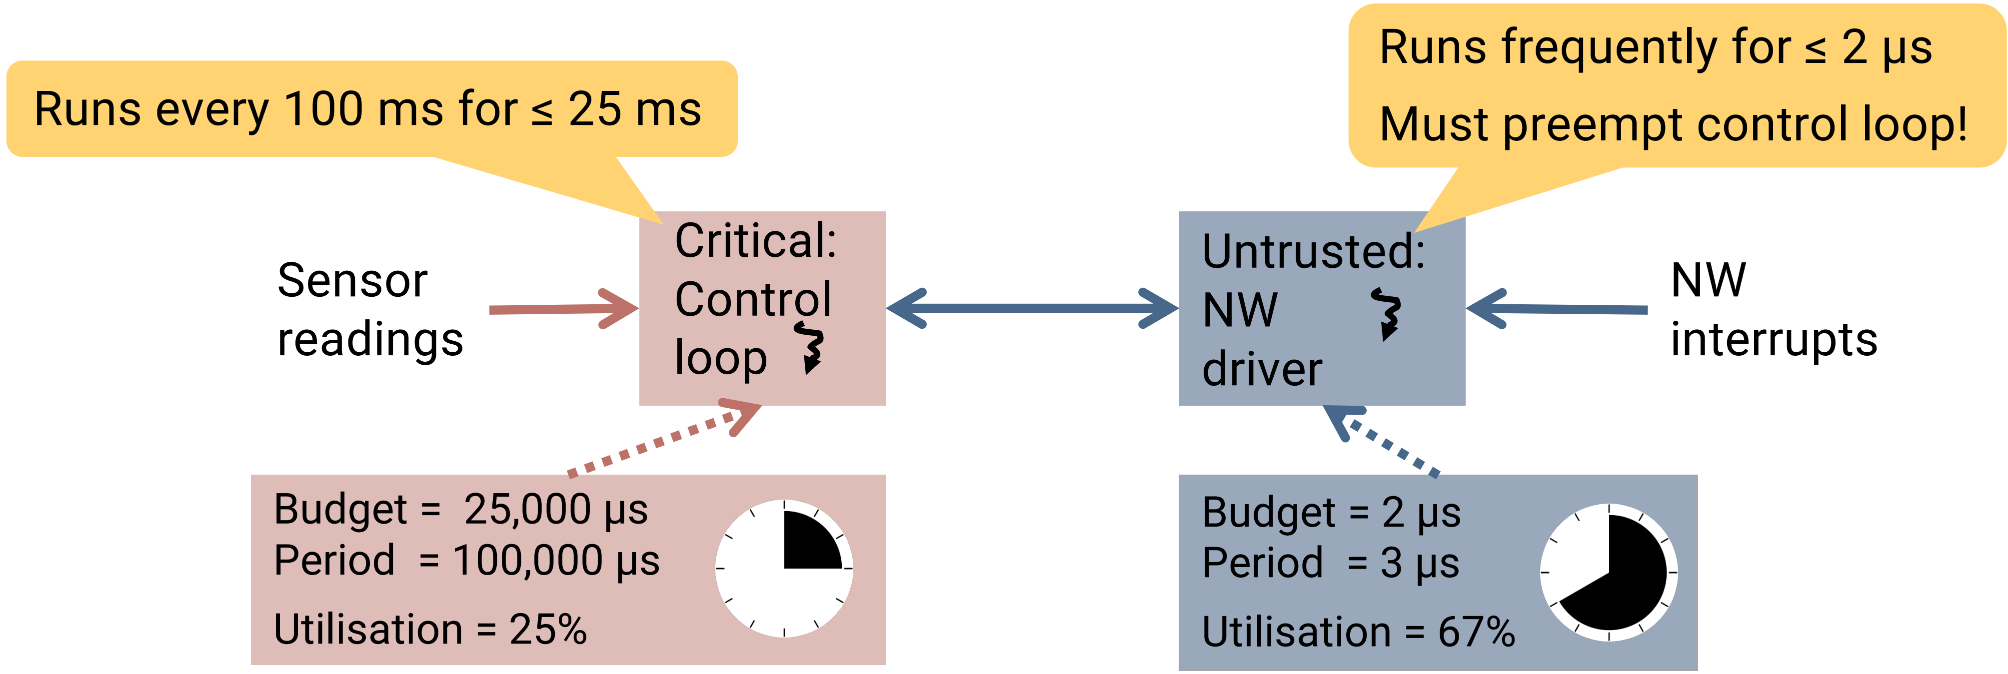

Time as first-class resource: capabilities provide bounded access to CPU

Worst-Case Execution-Time Analysis

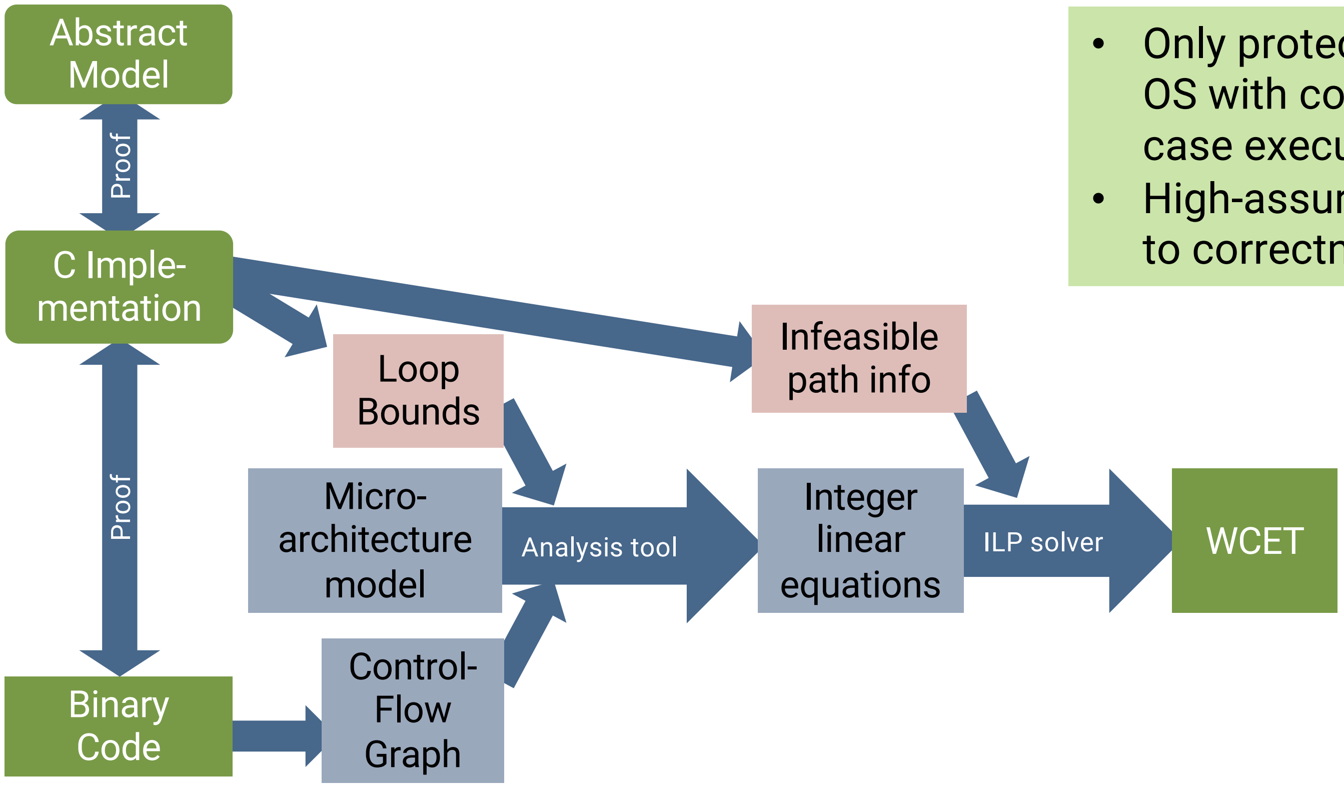

- Only protected-mode real-time OS with complete, sound worst-case execution-time analysis
- High-assurance by connecting to correctness proofs

Note: Armv6 only
- insufficient timing info for modern processors
- Open RISC-V implementations should enable it again!

The Benchmark for Performance

Round-trip cross-address-space IPC on 64-bit Intel Skylake

	seL4	Fiasco.OC L4Re	Zircon
Latency (cycles)	986	2717	8157
Mandatory HW cost* (cycles)	790	790	790
Overhead absolute (cycles)	196	1972	7367
Overhead relative	25%	240%	930%

Smaller is better

World's fastest microkernel!

***:** The Cost of `SYCALL` + 2 × `SWAPGS` + `SYSRET` = 395 cycles, times 2 for round-trip
Source:
Zeyu Mi, Dingji Li, Zihan Yang, Xinran Wang, Haibo Chen: "SkyBridge: Fast and Secure Inter-Process Communication for Microkernels", EuroSys, April 2019

Made For Real-World Use

Autonomous vehicles

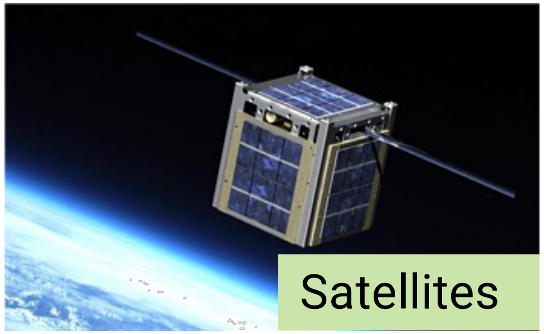

Satellites

Secure communication device
In use in multiple defence forces

Laot: Critical infrastructure protection

“World’s Most Secure Drone”

Tweet

DARPA
@DARPA

We brought a hackable quadcopter with defenses built on our HACMS program to @defcon #AerospaceVillage. As program manager @raymondrichards reports, many attempts to breakthrough were made but none were successful. Formal methods FTW!

Using seL4 in Cyberphysical Systems

seL4 Principles

Result: High barrier to uptake!

Security:
- Suitable base for security-critical systems
- Provably correct and secure

Proper microkernel:
- Minimal
- Provides policy-free mechanisms only
- Single access-control mechanism: Capabilities

Performance:
- Security is no excuse for poor performance!
- Don't pay for what you don't use

Anti-Principles:
- Hardware abstraction
- Prevent foot guns
- Usability

User-level issue!

The microkernel is the assembly language of operating systems!

Taming seL4: The seL4 Core Platform

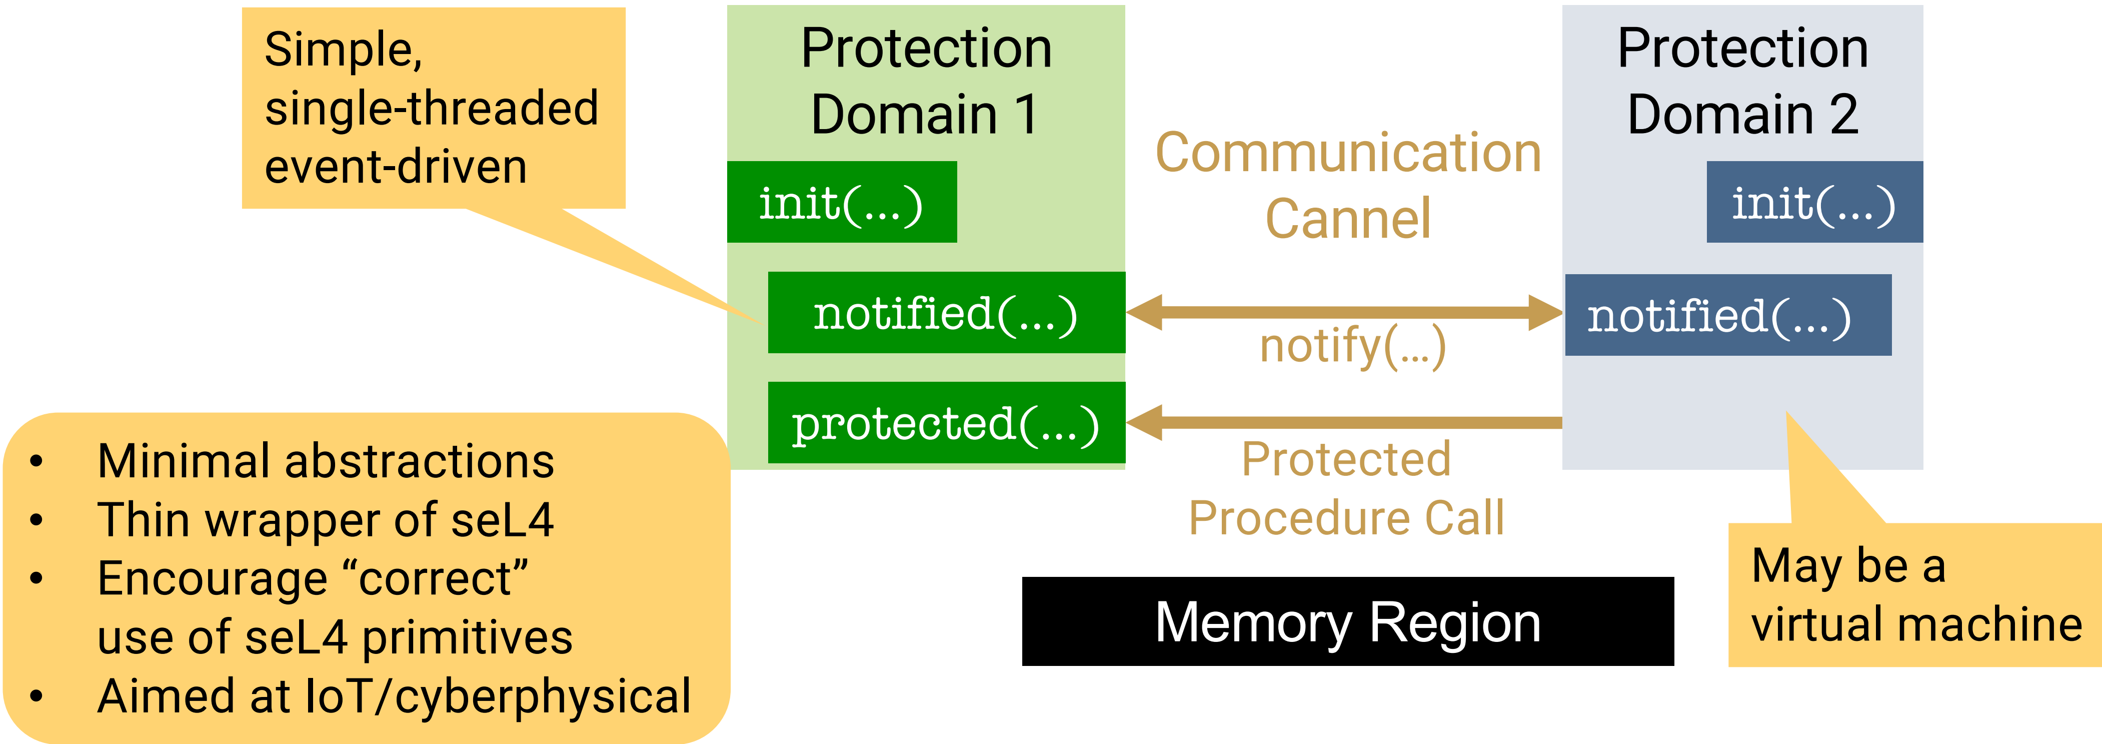

seL4CP Verification

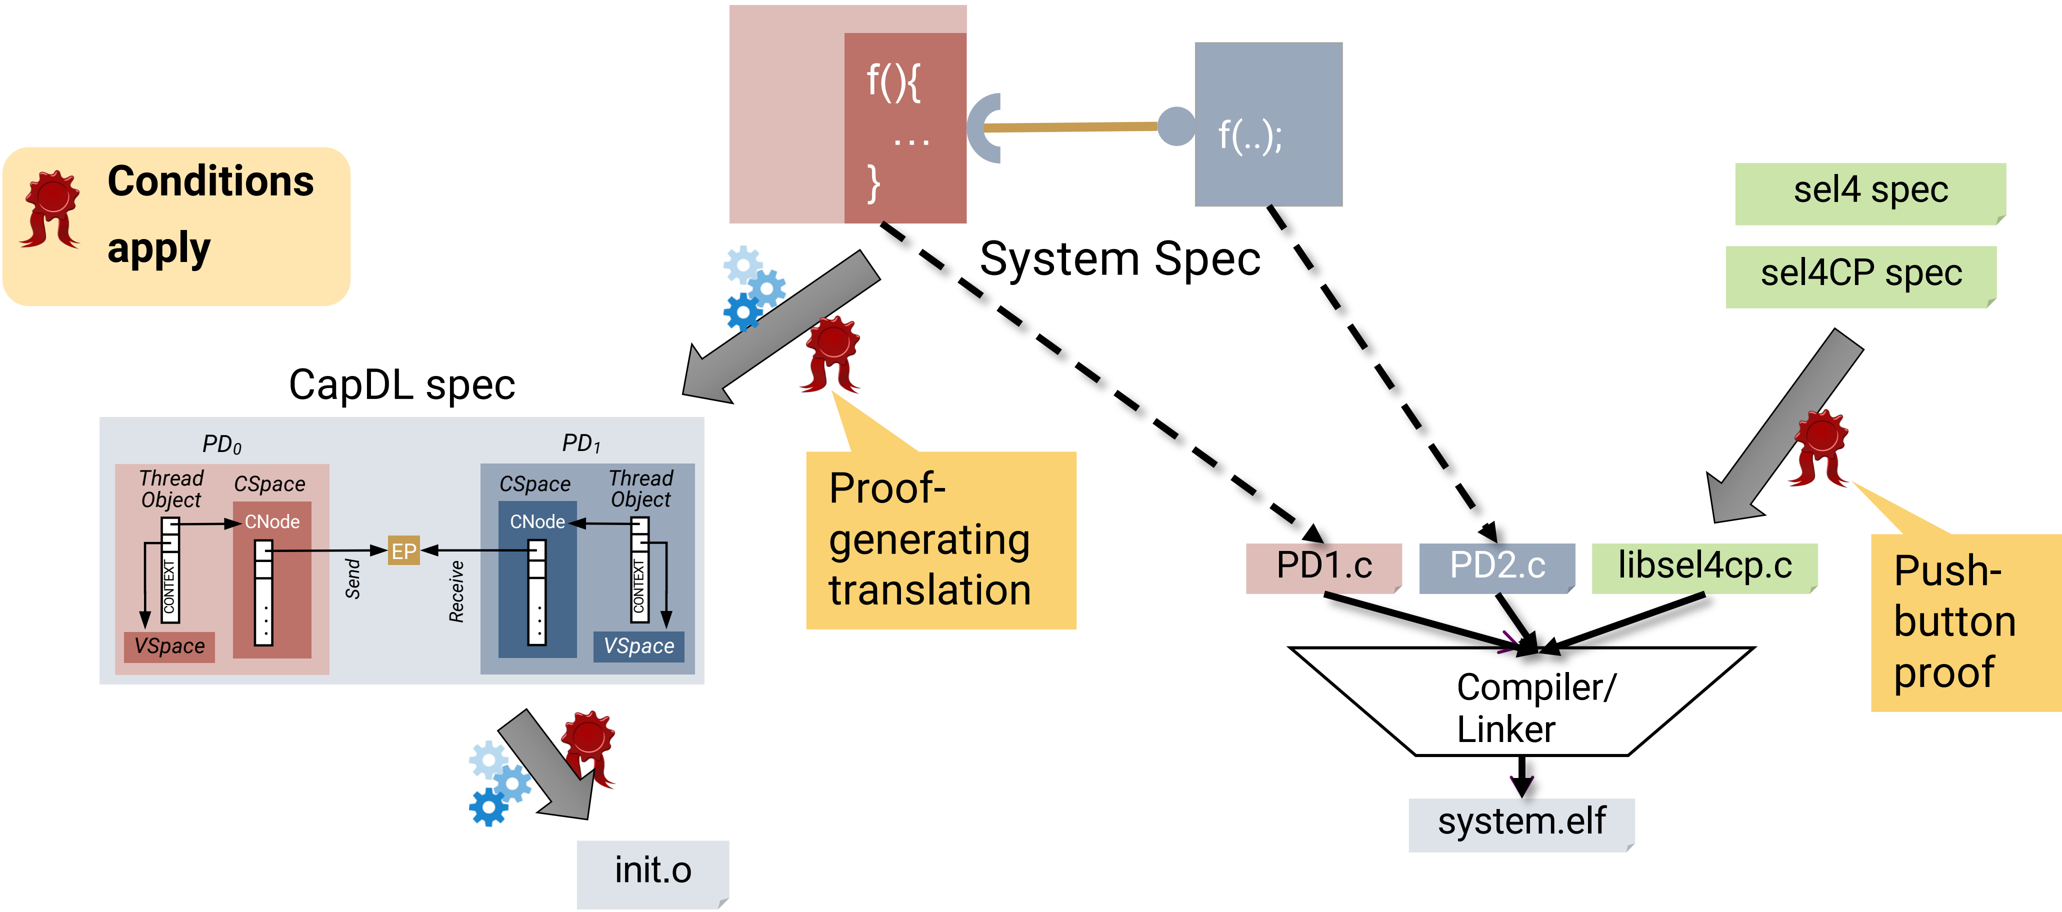

seL4CP-based Highly Modular OS

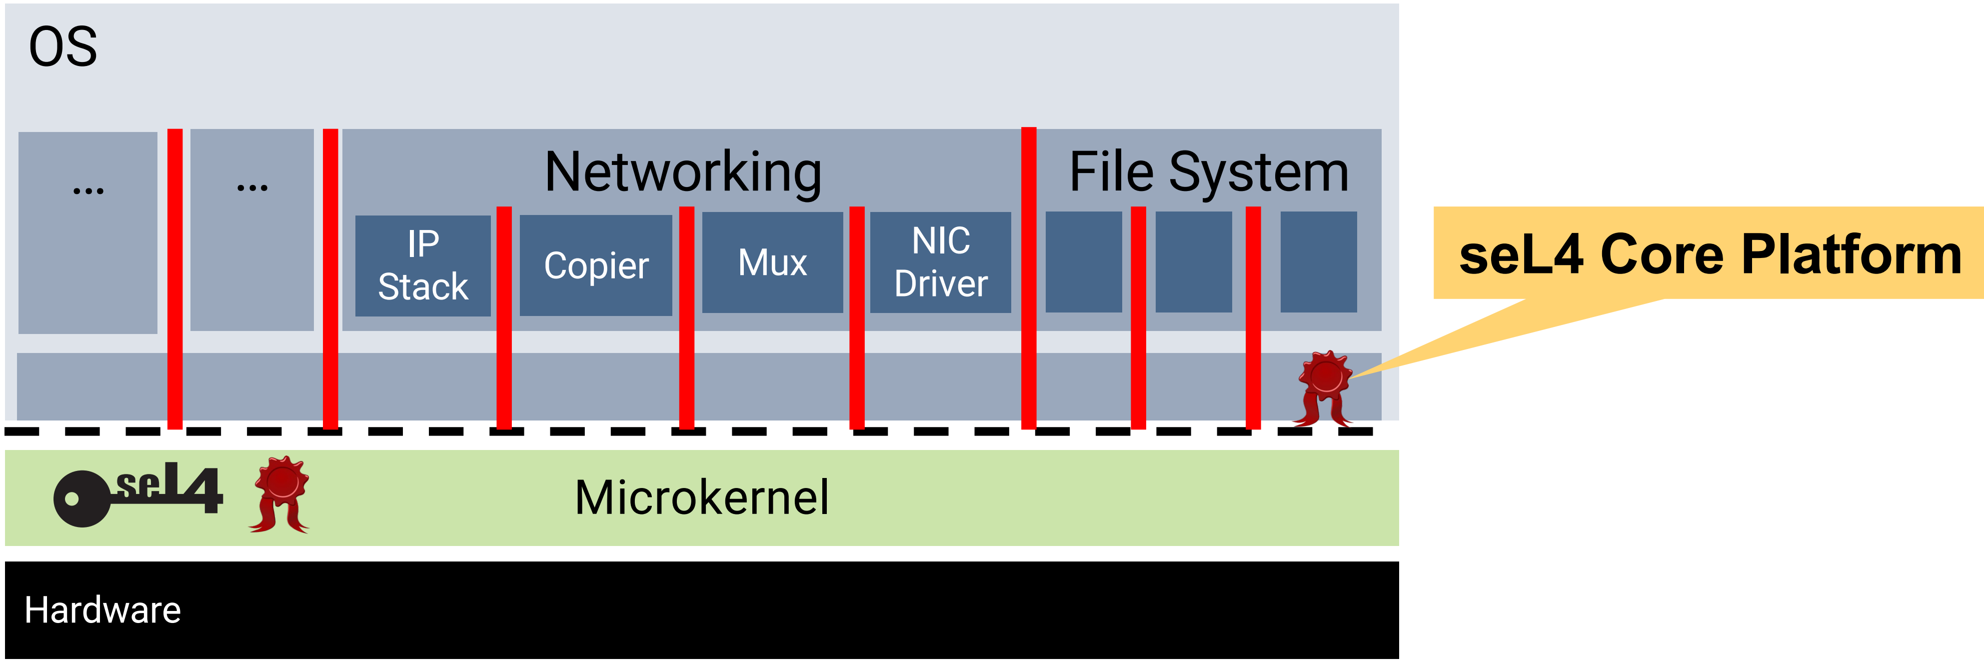

Example: Networking

Strict separation of concerns: Large number of extremely simple components

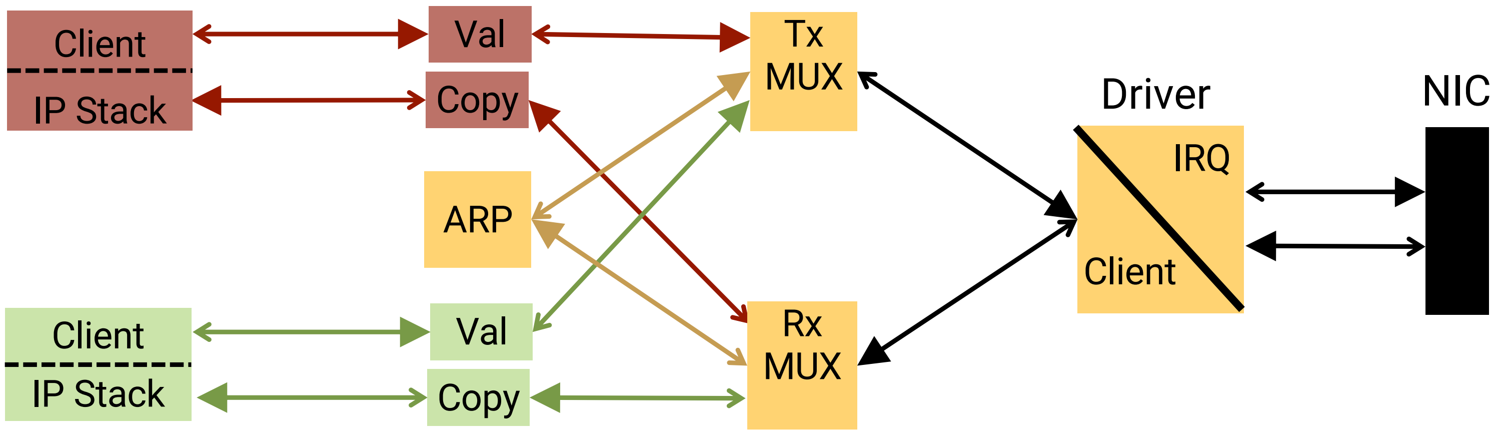

Comparison to Linux (i.MX8)

Linux:
- NW driver: 4k lines
- NW system total: 1M lines

Performance?

seL4 design:
- NW driver: 700 lines
- MUX: 400 lines
- Copier: 200 lines
- IP stack: much simpler, client library
- shared NW system total < 2,000 lines

Written by second-year student!

Evaluation Setup

- External load generator
- Client echoes packets

2 context switches per packet

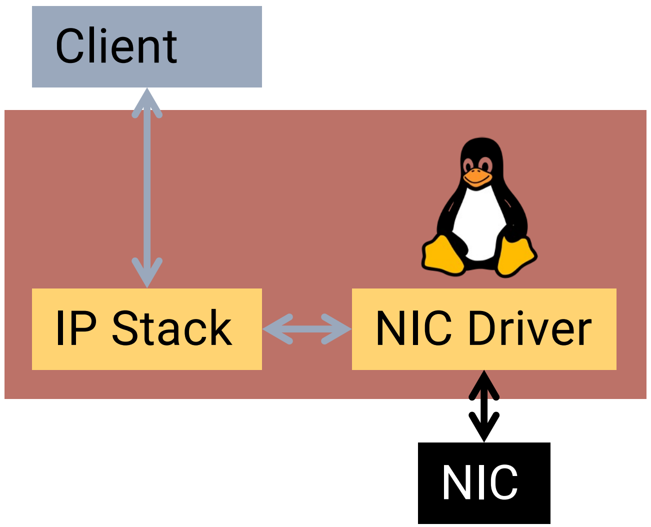

10 context switches per packet

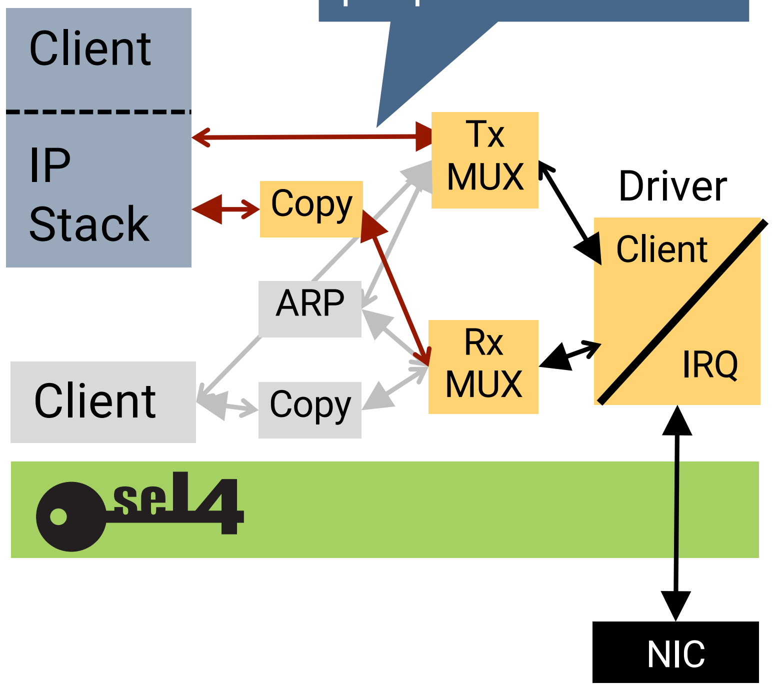

Achieved Performance: i.MX8

- Gigabit Ethernet
- single core

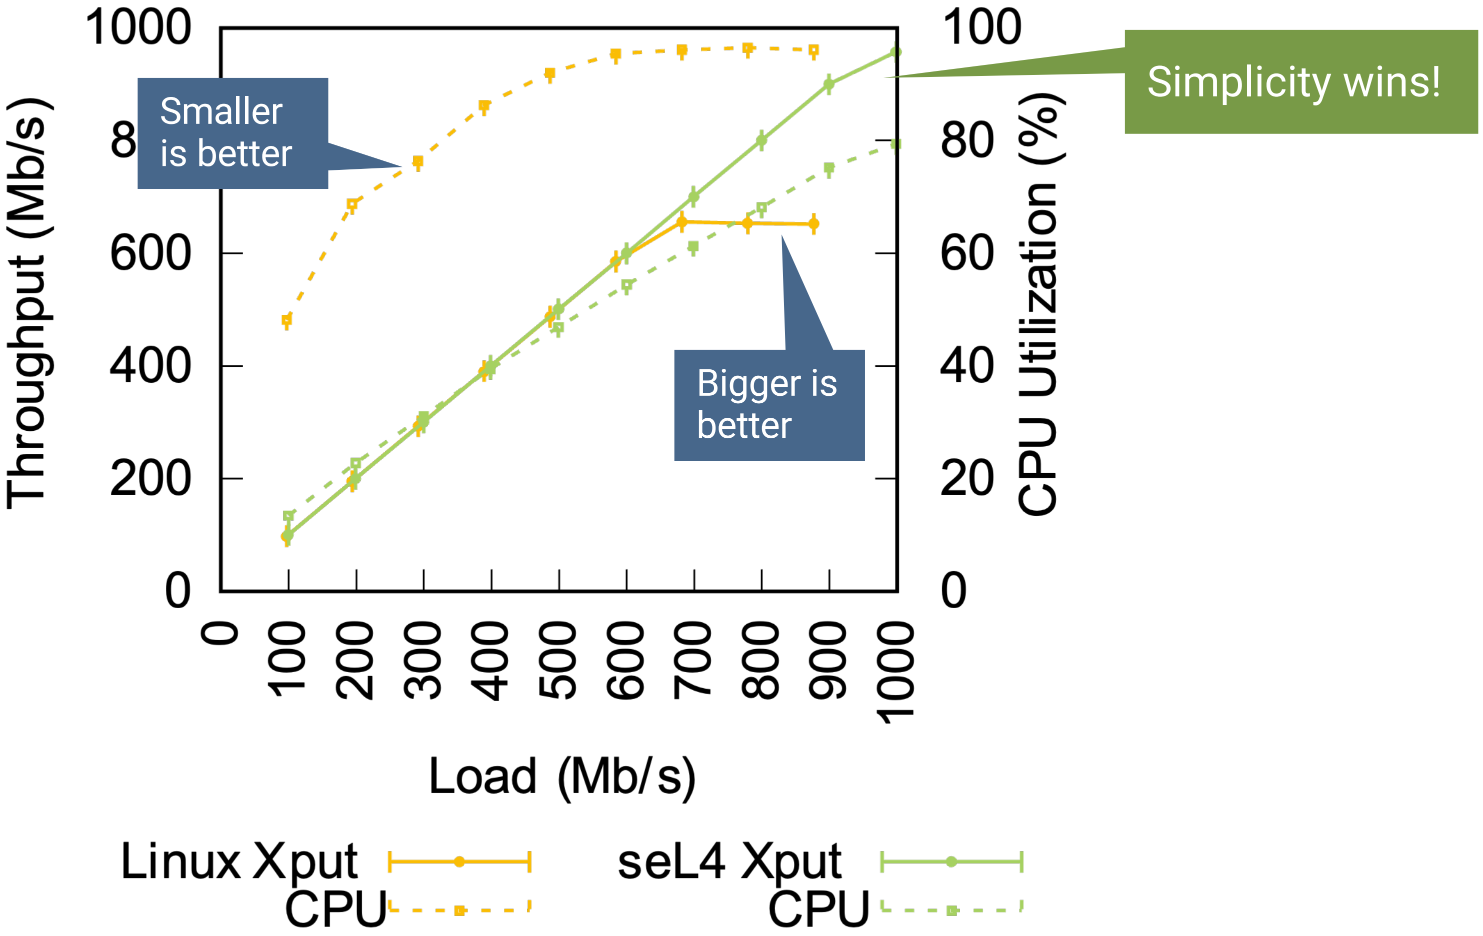

Highly Modular OS: Timeline

- Q4’23: First release of OS
  - with point-of-sale reference system
- Q2’24: Release of matured, documented, OS & PoS system
  - including performance analysis
- Q4’24: Verification of key components of OS

The seL4 Foundation

Premium Members

地 平 线 Horizon Robotics

jumptrading

NIO

UNSW SYDNEY

General Members

Cog

DORNERWORKS

galois

Google

HENSOLDT Detect and Protect.

KRY10

LOTUS

penten

Proofcraft

Raytheon Technologies

SPACEMIT 进迭时空

TII Technology Innovation Institute

xcalibyte

Associate Members

ETH zürich

KANSAS STATE UNIVERSITY

in association with National Cyber Security Centre

THE AUTOWARE FOUNDATION

RISC-V®

TUM

Security is no excuse for bad performance!

https://trustworthy.systems