



NICTA

# Persistence with Asynchronous I/O

Gernot Heiser

NICTA and University of New South Wales, Sydney

IFIPS WG 10.4, January 2011



Australian Government

Department of Broadband, Communications  
and the Digital Economy

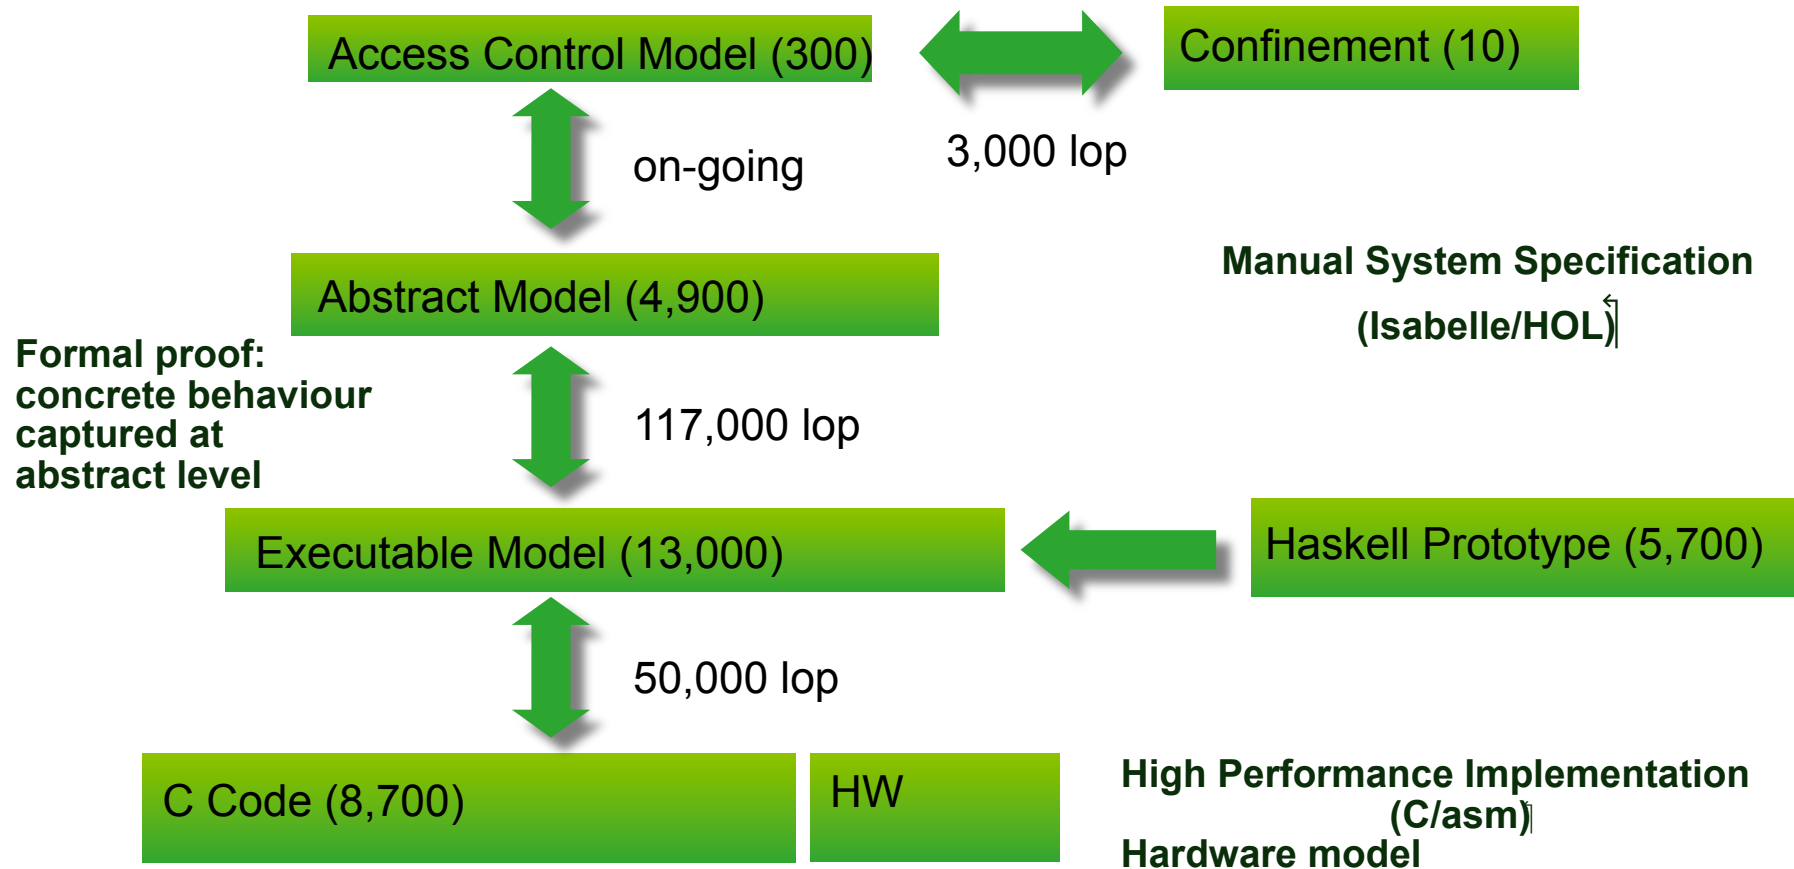
Australian Research Council

## NICTA Funding and Supporting Members and Partners




- Databases are concerned with dependability
  - ACID properties: atomicity, concurrency, integrity, durability
  - Ensured typically by write-ahead logging:
    - At transaction commit, write log data to disk
    - Block further processing until data is stable
- This puts slow I/O devices on the critical path of transactions – why?
- DBMS protects against system failure:
  - OS crash
  - Power outage
- If those could be ruled out, blocking could be avoided
- Requires a truly dependable system: seL4
  - Formally-verified OS microkernel

# seL4 Proof of Functional Correctness



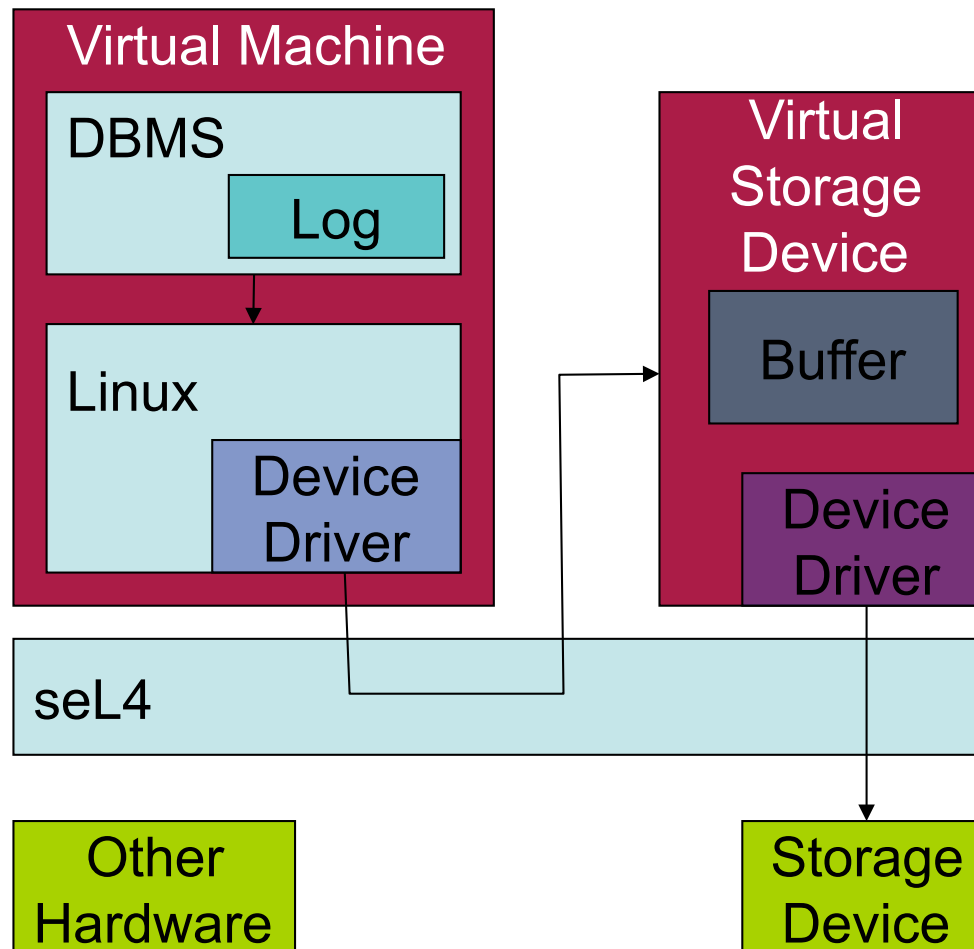
# How?

- Could port DBMS to run directly on seL4



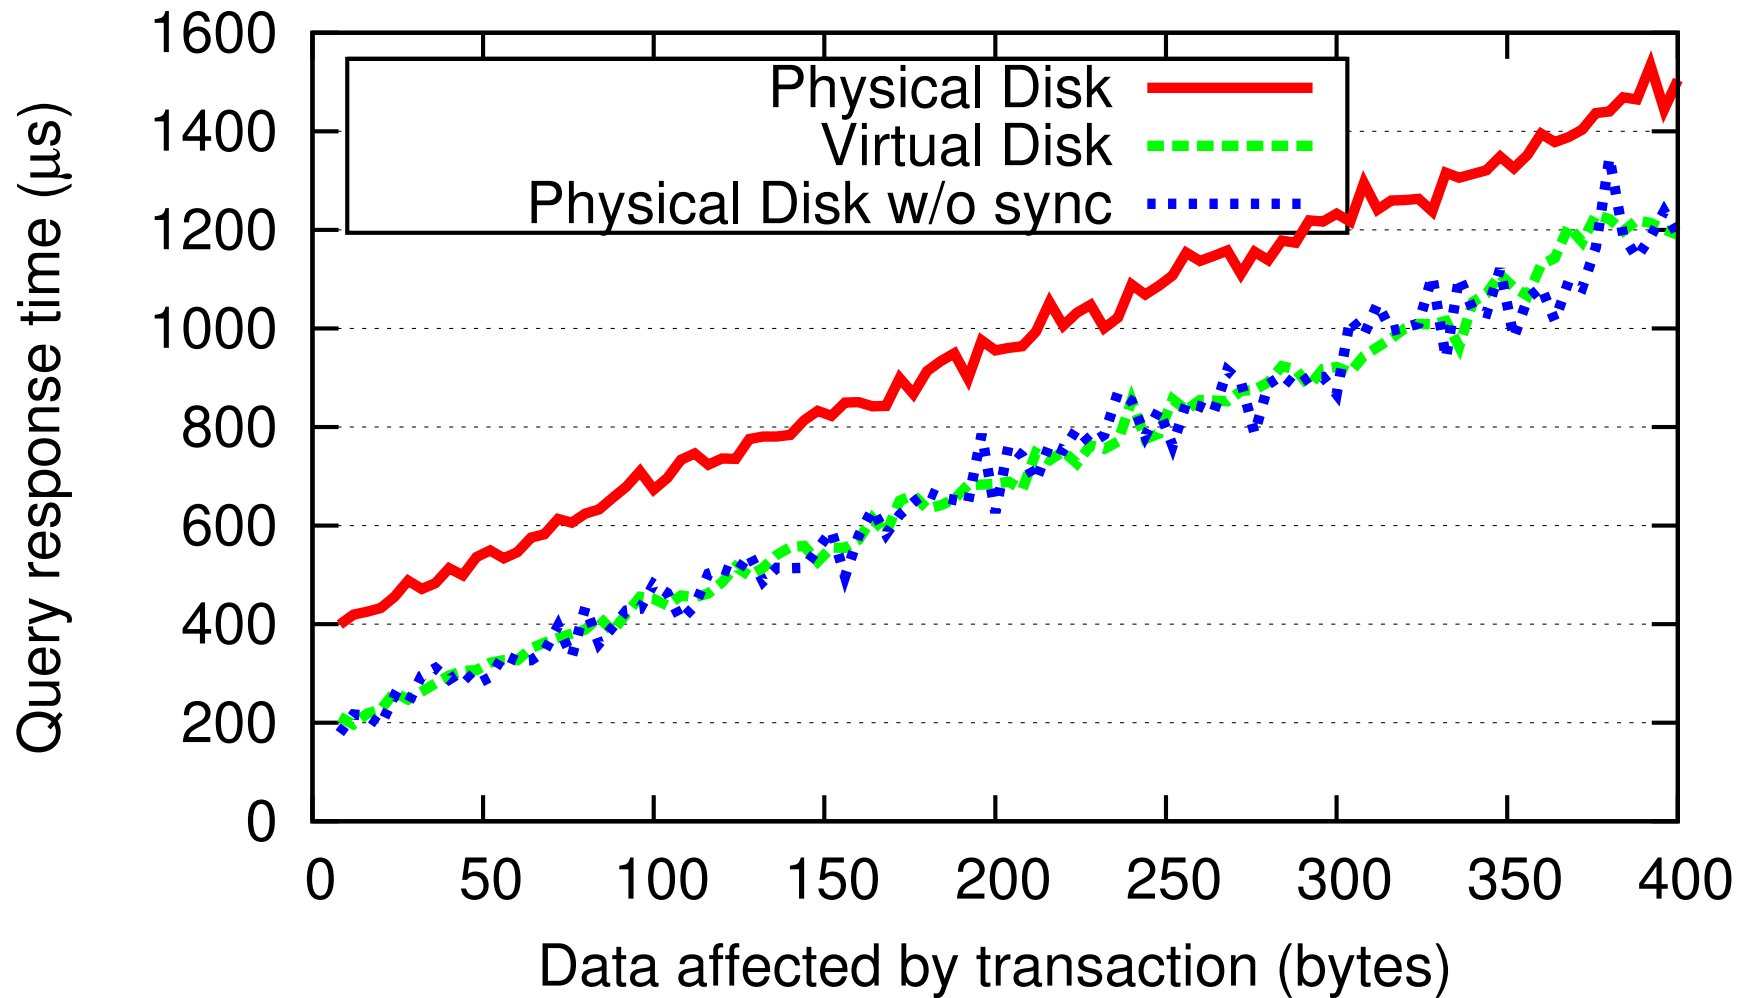
- Problem: costly, legacy issues, etc  $\Rightarrow$  not very attractive

# Alternative: Use Virtualization



- No changes to DBMS or OS!

# Performance



# Promising Initial Results!

---



- This is work in progress, stay tuned!

<mailto:gernot@nicta.com.au>

Google: "ertos"

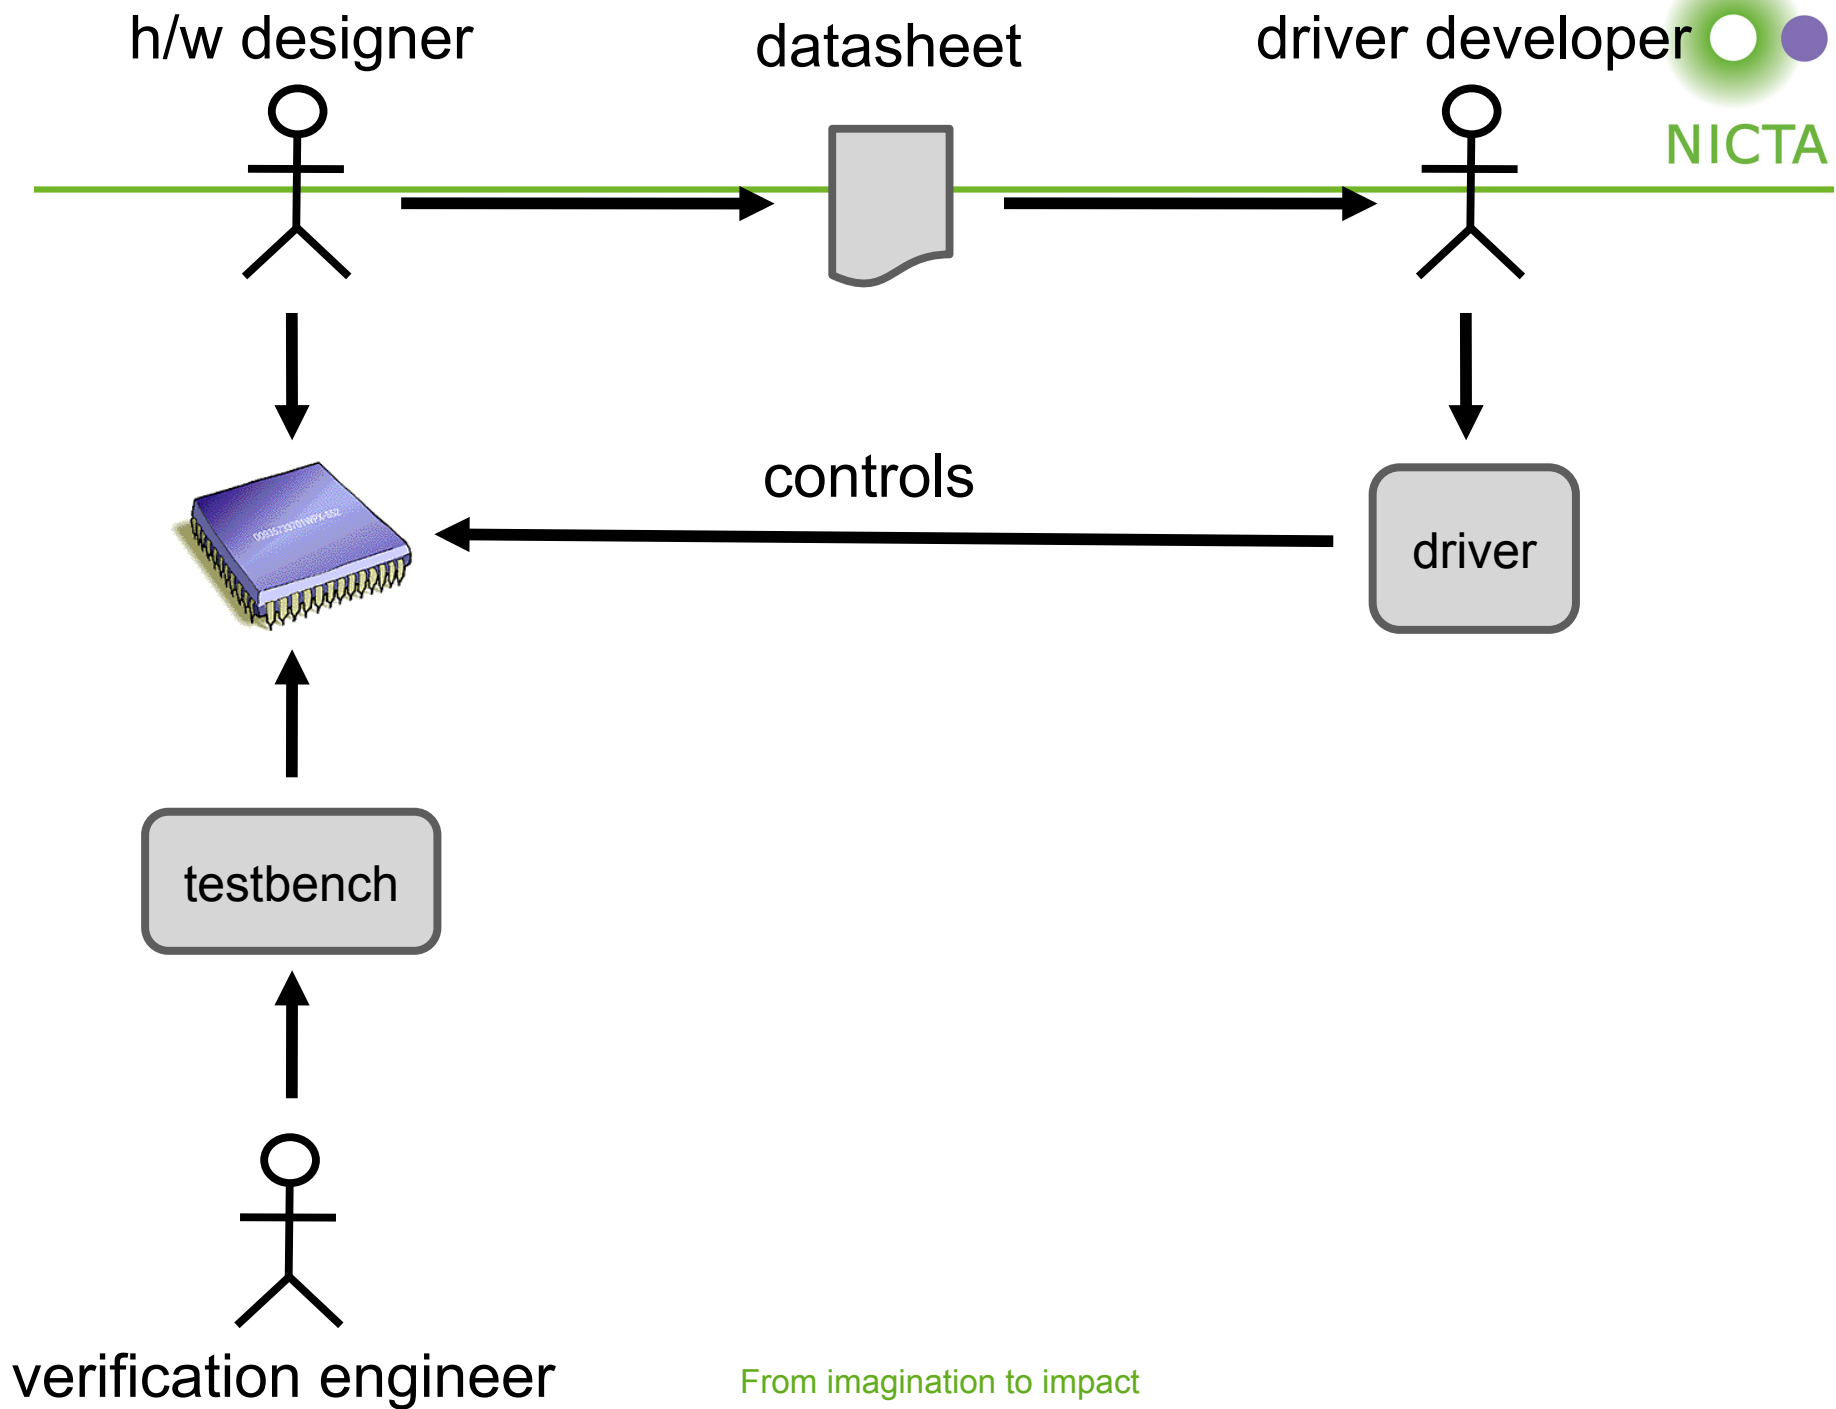
# Improving Device Driver Reliability

---



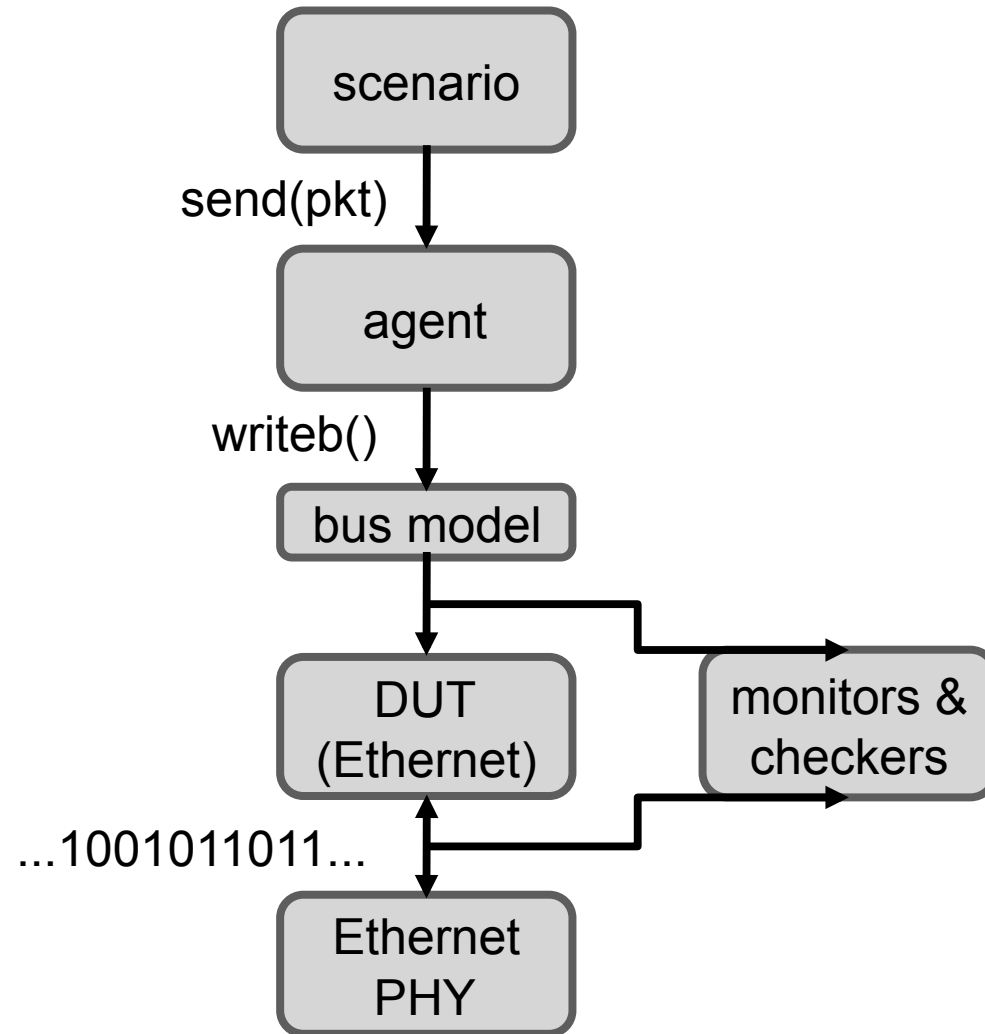
- Device drivers are the biggest source of OS crashes
- Problem is how they are produced



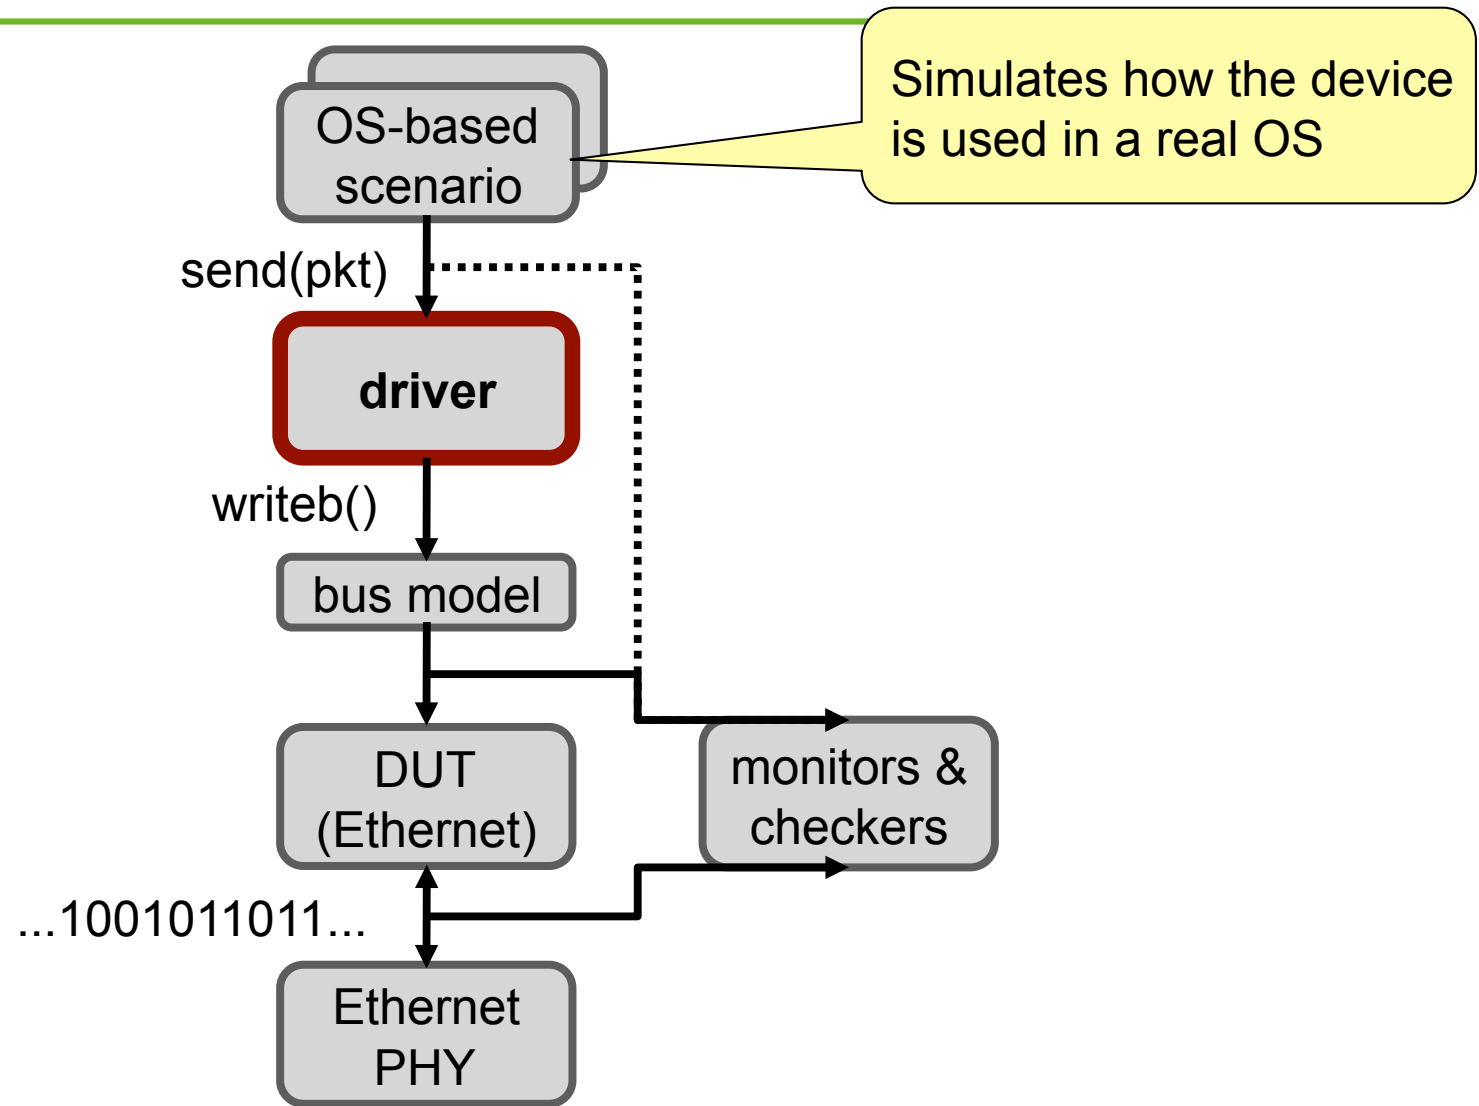


From imagination to impact

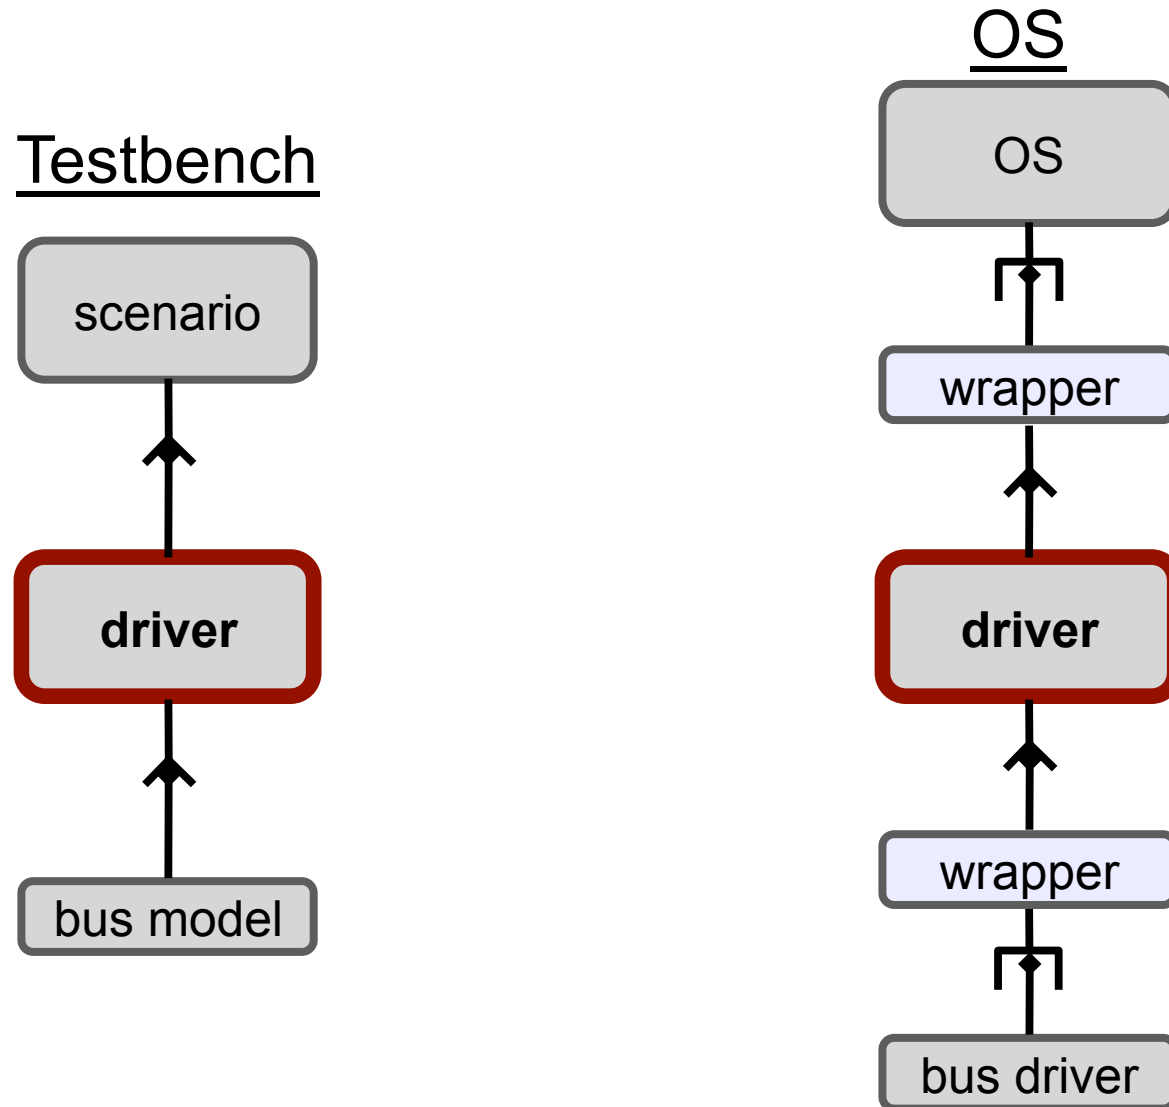
# Testbench architecture



# Hardware/software co-verification



# Driver interface unification



From imagination to impact

# Paper will be given at ASPLOS

---



- This is work in progress, stay tuned!

<mailto:gernot@nicta.com.au>

Google: "ertos"