# Securing Functional Programs with Hardware Support WG 2.8, Utrecht, 2024

**KC Sivaramakrishnan** 









## **Voter Absenteeism in India**

- 1/3 of eligible Indian voters do not cast their votes<sup>1</sup>
  - ~300 million in 2019
- **Reasons:** voter apathy, *internal migration*
- 450 million internal migrants in India<sup>2</sup>
  - Work and family reasons
- Constituency can be updated
  - But not no emotional connection, transient residence
- ECI  $\rightarrow$  IIT Madras, "How to enable remote voting?"

[1] "Discussion on improving voter participation of domestic migrants using remote voting", Election Commission of India [2] "Internal Migration in India Grows, But Inter-State Movements Remain Low", Census of India 2011

# **Electronic Voting Machines (EVM) in India**

- EVMs used in India since 1990s
- EVM unit has
  - Physical candidate list + Physical buttons
  - Voter-verified paper audit trail (VVPAT)
    - Small fraction per constituency cross-checked
- Technology<sup>1</sup>
  - Processor
    - NXP MK61FX512VMD12 (ARM Cortex-M4) and others
  - Software built by ECIL and BEL
    - Closed source

[1] Patrick Jones, "Delivering Democracy: The History and Deployment of Electronic Voting Machines in India and the United States", PhD thesis



**Ballot Unit** 

**VVPAT** 



# **Remote Voting Machine (RVM)**

- Process<sup>1</sup>
  - Remote voters pre-register to vote from a remote constituency
  - Ballot list pre-loaded onto the machine
- Technology
  - No network connection
  - **Ballots for different constituencies**
  - Electronic display

### **Public Ballot Display**



# **Remote Voting Machine (RVM)**

- Long road ahead
  - Need amendments in the Constitution to allow remote voting
  - Social challenges Building trust
  - Technological challenge more complex than EVM

[1] "Discussion on improving voter participation of domestic migrants using remote voting", Election Commission of India, 2022

**Public Ballot Display** 



# An idealised RVM stack

- Desiderata
  - **Open-source** secrecy is not security
  - Hardware and software security features
  - Formally-verified software, where possible
  - Small trusted computing base (TCB)
- Out of scope<sup>1</sup> lacksquare

  - End-of-poll audit
- Shakti RISC-V processor running bare-metal MirageOS unikernels

[1] "Citizens' Commission on Elections' Report on EVMs and VVPAT", IIT Delhi

E2E cryptographic verifiability: cast-as-intended, recorded-as-cast, counted-as-recorded



## Hardware and Software

- Shakti RISC-V processor family
  - https://gitlab.com/shaktiproject
  - developed at IIT Madras
  - Bluespec SystemVerilog
  - C-class: 5-stage, in-order; can boot Linux
  - 2 tape-outs: 22nm Intel, 180nm SCL
- MirageOS Unikernels
  - https://mirageos.org/
  - A library operating system in OCaml
  - Small TCB
  - Runs on Xen and KVM hypervisors, baremetal





## Challenge



- - C vulnerabilities can compromise OCaml's safety
- **Goal:** Safe hardware/software stack for constrained environments



• C code is unavoidable — OCaml runtime, mirage runtime, drivers, sqlite, crypto



# **Compartments / SFI — overview**

- Compartments offer *intra-process* isolation
  - Functions mapped to compartments
  - Restrict control flow and data access across security boundaries
- Control flow restricted by
  - Whitelisted PC ranges
  - Shadow stack to prevent ROP attacks
- Data access restricted by
  - VMM tricks (or) fat pointers (or) capabilities (à la CHERI)



## FIDES – Secure compartments

- Security-hardened Shakti RISC-V processor + MirageOS unikernels
- Intra-process compartments
  - Vulnerabilities in C do not affect OCam
- Compartment access matrix defined at *link time*
  - Run *unmodified* OCaml and C code
- Small extension to hardware and software
  - Two new instructions added to RISC-V ISA: Val and Checkcap
  - Modification to LLVM and OCaml compiler to emit these instructions



### **Access Matrix**

## Threat model

- Source code is untrusted
  - Inline assembly and use of Obj.magic trusted
- All code is compiled with FIDES C and OCaml compiler
  - Compiler instrumentation added by FIDES is correct
  - OCaml runtime is trusted
- Binary executable cannot be tampered with
- Hardware attacks rowhammer, fault attacks, side-channels are out of scope

### Guarantees

- Control-flow integrity
  - The control flow in every execution of the program respects the compartment access matrix
- Memory safety
  - No memory errors; all references point to valid memory
  - Pointers cannot be forged

# **FIDES** — Challenges and opportunities

- OCaml offers memory safety
  - Hardware-accelerated fat pointers only for C code
    - Fine-grained data compartments
    - No fat pointers for OCaml code
  - Pay attention to FFI boundaries

### FIDES code compartment must now handle FP features!

Higher-order functions, tail calls, exceptions

# **Compartments for RVM**



### indicates call is allowed



Highly secure



### indicates call is allowed

let res\_tab = Array.make num\_candidates 0 (\* candidate\_id->num\_votes \*)

let count\_votes votes\_arr (\* decrypted votes array \*) = let inc\_vote candidate\_id = res\_tab.(candidate\_id) <- res\_tab.(candidate\_id) + 1</pre>

Array.iter inc\_vote votes\_arr





## Higher-order functions — Idea 1



### indicates call is allowed

let res\_tab = Array.make num\_candidates 0 (\* candidate\_id->num\_votes \*)

let count\_votes votes\_arr (\* decrypted votes array \*) =
 let inc\_vote candidate\_id =
 res\_tab.(candidate\_id) <- res\_tab.(candidate\_id) + 1</pre>

Array.iter inc\_vote votes\_arr

"Confused Deputy" attack



## **Higher-order functions – Idea 2**



### indicates call is allowed

let res\_tab = Array.make num\_candidates 0 (\* candidate\_id->num\_votes \*)

```
let count_votes votes_arr (* decrypted votes array *) =
  let inc_vote candidate_id =
    res_tab.(candidate_id) <- res_tab.(candidate_id) + 1</pre>
```

Array.iter inc\_vote votes\_arr

X Limited compartment resource Shared state?





## Fluid compartments



### indicates call is allowed

let res\_tab = Array.make num\_candidates 0 (\* candidate\_id->num\_votes \*)

let count\_votes votes\_arr (\* decrypted votes array \*) =
 let inc\_vote candidate\_id =
 res\_tab.(candidate\_id) <- res\_tab.(candidate\_id) + 1</pre>

Array.iter inc\_vote votes\_arr

A fluid compartment inherits the policies of the caller



## Shadow stack

- Stores the return addresses for inter-compartment calls
- Inaccessible from user-code
  - Maintained and validated by hardware





## Non-call-return control flow

- Typical compartment schemes handle only call-return sequence
- OCaml has several non-call-return control-flow operations
  - Tail calls, exceptions, effect handlers!
  - Need to manage the shadow stack carefully



- Exceptions may be thrown across compartments
  - Need to unwind shadow stack appropriately
  - Challenge: Detect when intra-compartment exceptions are raised
- Solution: Security monitor (SM) updates last exn\_pc to a special routine

# Tail calls and shadow stacks

### Intra call then inter tailcall



### At g to h call, return address at f pushed onto the shadow stack

### Inter call then inter tailcall



At g to h call, nothing pushed onto the shadow stack

## **Fat pointers**



- Fat pointers into the stack have frame scope
  - Each frame has a cookie freshened at call and return
- val instruction validates fat pointer before access
- OCaml does not use fat pointers •
  - At FFI, use OCaml object header info to create fat pointer
  - Use a special cookie that skips temporal validation

### Hardware changes

# for each instruction hw\_check: jmp sm\_entry

sm\_exn: checkcap

sm\_ret: checkcap

```
if (not (pc in pc_range(cur_comp) or
         pc in pc_range(fluid_comp))):
    assert (*pc == checkcap)
```

## Hardware changes

```
14 sm_entry:
     caller = ra # caller's return address
15
     callee = csr_comepc # callee's entry address
16
     if (callee == sm_ret):
17
18
       # inter-compartment return
19
       jmp sm_exit
     if (callee == sm_exn):
20
       # handling inter-compartment exception
21
22
       jmp sm_exn_unwind_shadow_stack
     if (caller == sm_ret):
23
       # inter-compartment tail call preceded by
24
       # inter-compartment call
25
       jmp l_tail_call
26
     if (in_C(callee)):
27
       save(sm_stack,C_callee_saved_registers)
28
     else:
29
       # used during GC for finding roots on the stack across
30
31
       # compartments
       push(sm_gc_stack,caller)
32
```

```
# Update the PC of the latest exception handler so
33
     # that inter-compartment exceptions can be detected
34
     exn_handler_pc = get_latest_exn_handler_pc()
35
     if (exn_handler_pc != sm_exn):
36
       set_lastest_exn_handler_pc(sm_exn)
37
       push(sm_stack, exn_handler_pc)
38
     push(sm_stack, caller)
39
     push(sm_stack, sp)
40
     push(sm_stack, cur_comp)
41
     target_comp = compartment_id(callee)
42
     l_tail_call:
43
      if (not (access_allowed(cur_comp, target_comp))):
44
45
           access_error()
     zero_out(temp_and_callee_saved_registers)
46
47
     cur_comp = target_comp # switch to the target compartment
48
     ra = sm_ret
     jmp callee
49
```



## **Evaluation**

- Compiler changes
  - ~300 lines for OCaml, ~2300 lines for LLVM
- Protyped on Xilinx Artix-7 AC701 FPGA
  - ► 38.2K LUTs (+6.1% over base)
  - 17.4K registers (+6.0% over base)
- Performance on voting application
  - 4% increase in code size
  - 23% increase in instruction cycle count





## Limitations

- **Features:** Effect handlers, parallelism
- OCaml runtime is trusted
  - WIP: Verified garbage collector for OCaml
- Data compartments are too weak
  - Objects shared across compartments remain accessible forever
  - Revocation through ownership and borrowing à la Rust
    - modal types in OCaml
- Hardware is exotic
  - Arm MTE for fat pointers in C?

### References

- Panel on Indian Electronic Voting Machines (EVMs), EVT/WOTE '10
- Patrick Jones, <u>"Delivering Democracy: The History and Deployment of</u> Electronic Voting Machines in India and the United States", PhD thesis
- <u>"Internal Migration in India Grows, but Inter-State Movements Remain Low"</u> Census of India 2011
- <u>"Citizens' Commission on Elections' Report on EVMs and VVPAT"</u>, IIT Delhi
- <u>"Discussion on improving voter participation of domestic migrants using</u> <u>remote voting</u>", Election Commission of India
- <u>"FAQs on Security Features of The ECI-EVMs"</u>, ECI