
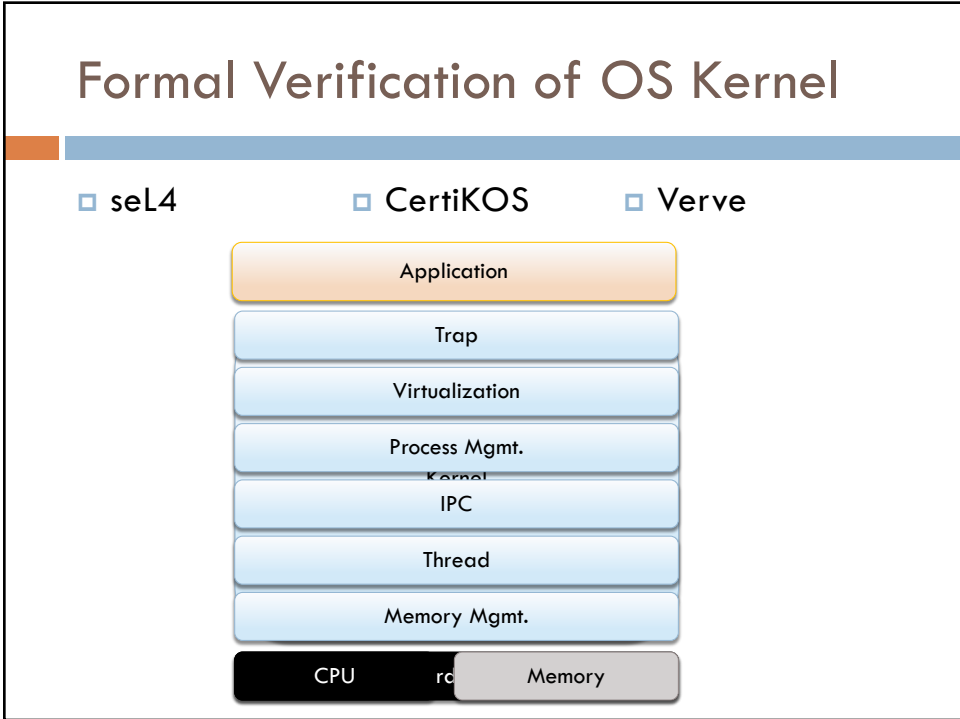


CERTIFIED INTERRUPTIBLE OS KERNELS AND DEVICE DRIVERS

Zhong Shao
Joint work with Hao Chen, Newman Wu, Joshua Lockerman, and Ronghui Gu
Yale University

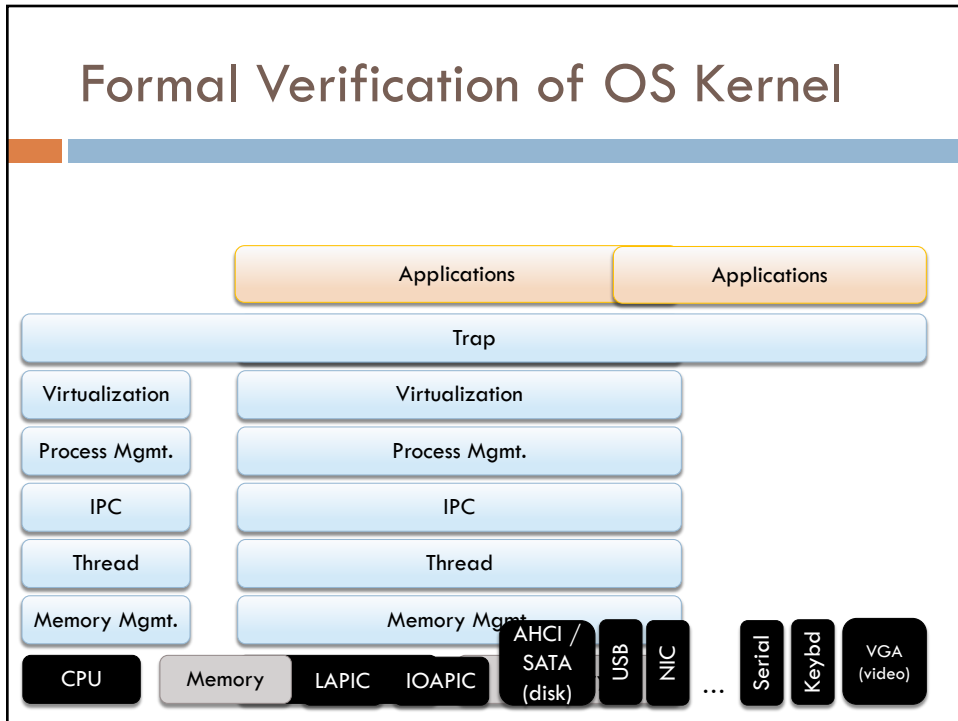

CPSC 422 Operating Systems
October 17, 2022

1



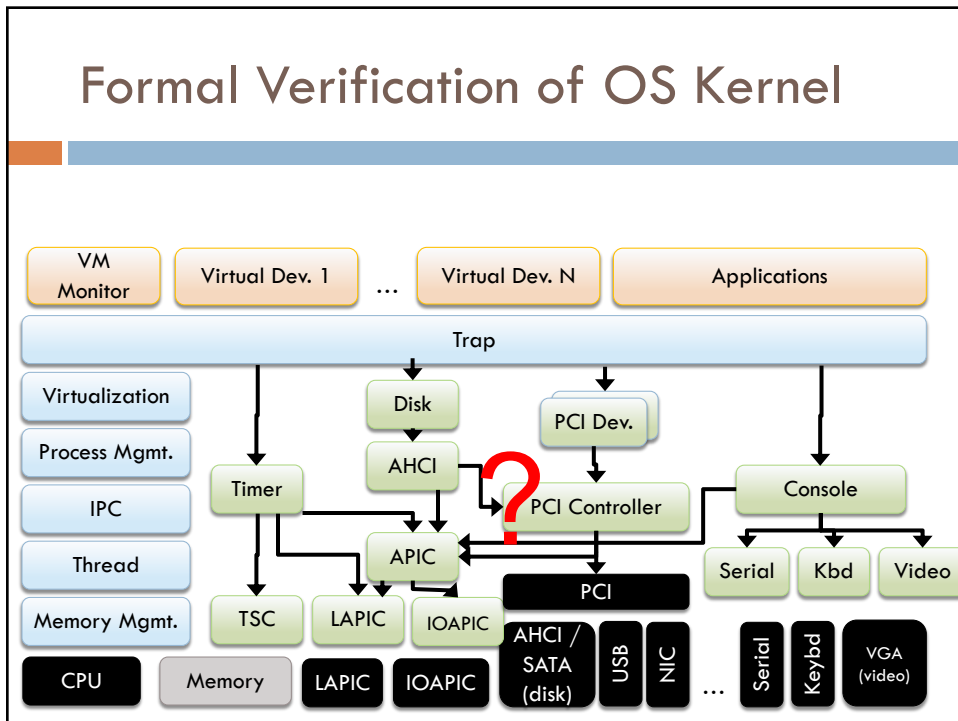
2

Formal Verification of OS Kernel



3

Formal Verification of OS Kernel



4

Device Drivers in Mainstream OS


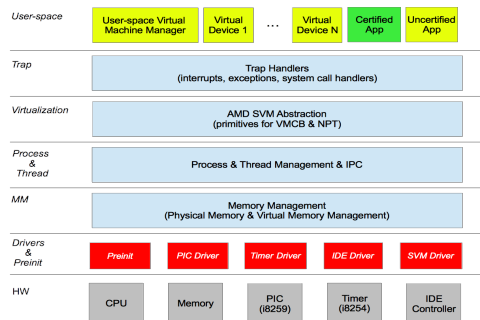
- 70% of Linux 2.4.1 kernel are device drivers.
- 70% of Windows crash are caused by third-party driver code.



5

mCertiKOS Overview [POPL'15]

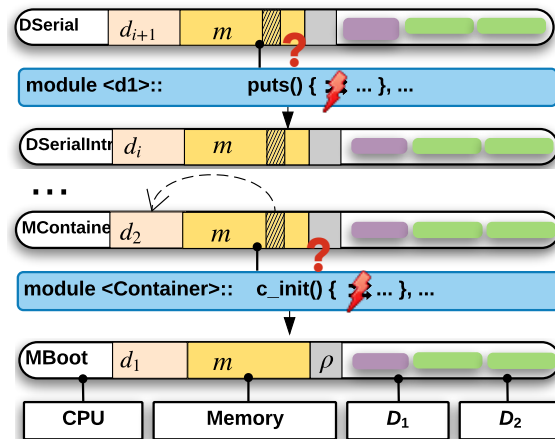
- Single-core version of CertiKOS.
- 3k LOC, can boot Linux as guest.
- Aggressive use of abstraction over deep specification (37 layers).



6

Main Challenge

Every fine-grained processor step could be interrupted.



7

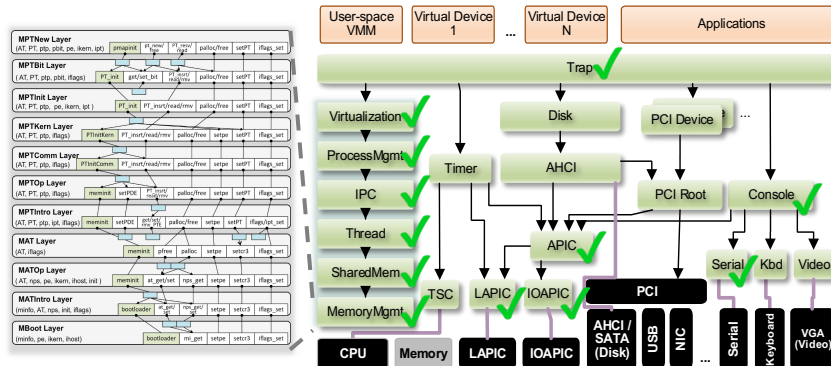
Other Challenges

- ▣ Interrupt hardware can be *dynamically* configured.
- ▣ Devices and CPU run in *parallel*.
- ▣ Device drivers are written in both C and *assembly*.
- ▣ The correctness results of different components should be *linked formally*.

8

Our Contributions [PLDI'16]

The *first* formally verified *interruptible* OS kernel with *device drivers*.



9

Our Contributions [PLDI'16]

- New techniques for certifying abstraction layers with multiple *logical CPUs* and devices.
- New techniques for building formal *certified device hierarchies*.
- An abstraction-layer-based approach for reasoning about *interrupts*.
- **Case study:** interruptible mCertikOS with device drivers.

10

Our Contributions

- ▣ New techniques for certifying abstraction layers with multiple *logical CPUs* and devices.
- ▣ New techniques for building formal *certified device hierarchies*.
- ▣ An abstraction-layer-based approach for reasoning about *interrupts*.
- ▣ **Case study:** interruptible mCertiKOS with device drivers.

13


Hardware Device Model

- ▣ Devices are modeled as transition systems parameterized by all possible lists of external events.
- ▣ Example external events:
 - ▣ Recv (s: list char)
 - ▣ KeyPressed (c: Z)
- ▣ State: observable registers.
- ▣ Transition:
 - ▣ environmental transition: δ^{env}
 - ▣ I/O transition: δ^{CPU}

14

Raw Device Object

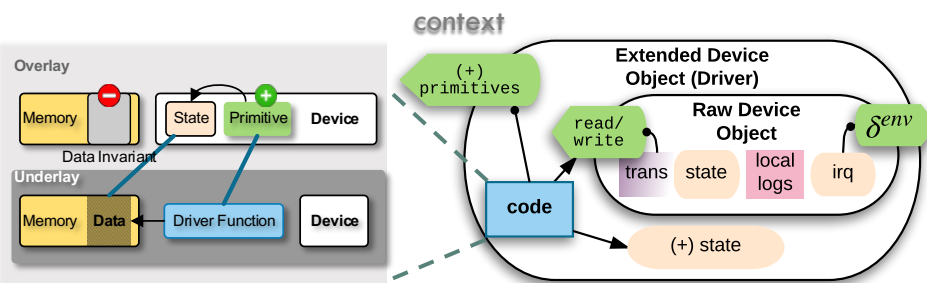
- Local log for the list of observed external events.
- Multiple local logs to handle disjoint set of external events asynchronously.
- Read/Write instructions: IN/OUT, memory mapped I/O, etc.



15

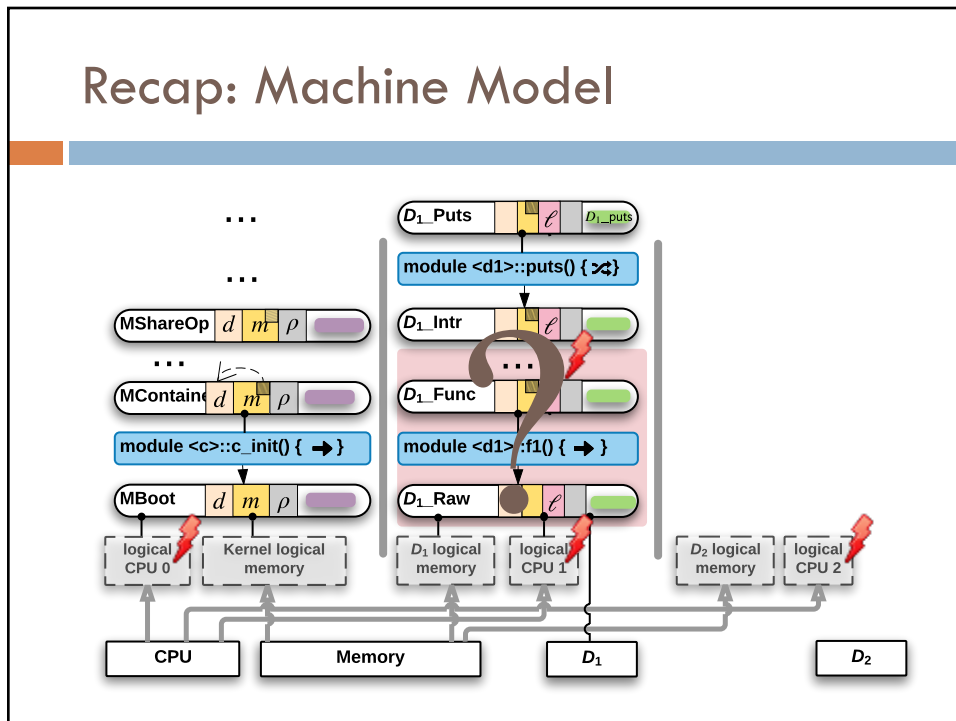
Extended Device Object

Driver as a logical device.



16

Recap: Machine Model



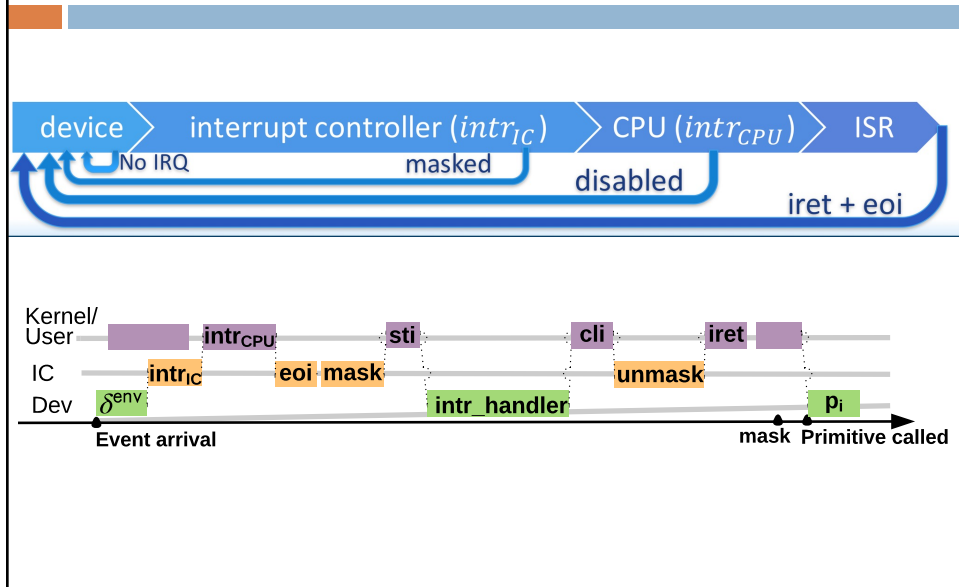
17

Our Contributions

- ▣ New techniques for certifying abstraction layers with multiple **logical CPUs** and devices.
- ▣ New techniques for building formal **certified device hierarchies**.
- ▣ An abstraction-layer-based approach for reasoning about **interrupts**.
- ▣ **Case study:** interruptible mCertikOS with device drivers.

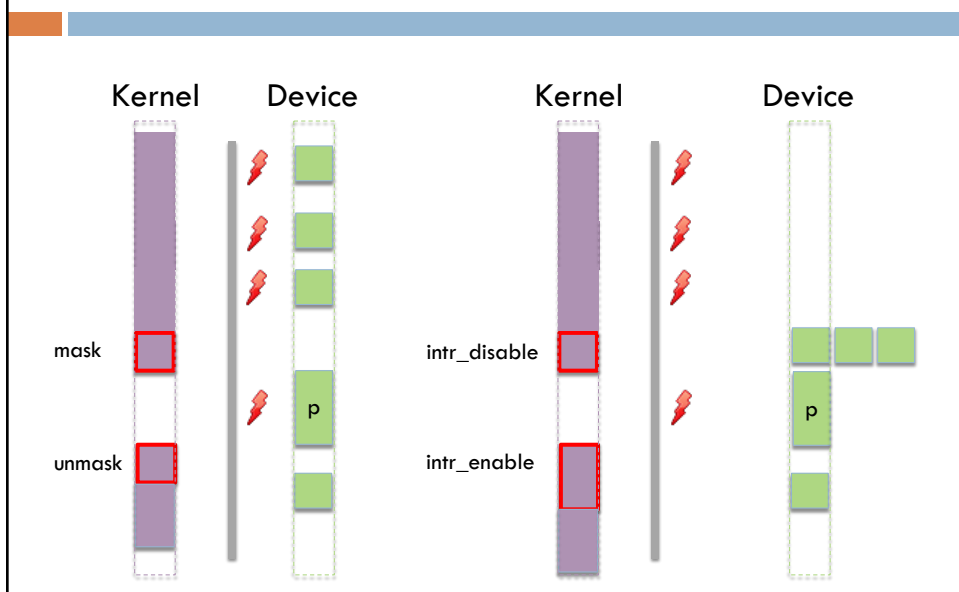
18

Interrupt Models



19


New HW Interrupt Model



20

Semantics of *intr_disable*

- Scans external events.
- Recursively performs the environmental transition.
- Synchronizes unhandled interrupts.



details in the paper

DISABLENOINTR: Disable with no unhandled interrupt

$$\frac{(e, \ell_i) = \text{next}(\ell^{env}, \ell_i) \quad s_{\text{tmp}} = \delta^{\text{env}}(s, e) \quad s'.\text{irq} = \text{false} \quad s' = s[\text{Flag} \leftarrow 0]}{\text{intr_disable}(s, \ell_i, \ell^{env}) = (s', \ell_i)}$$


DISABLEINTR: Disable with unhandled interrupts

$$\frac{(e, \ell_i) = \text{next}(\ell^{env}, \ell_i) \quad s' = \delta^{\text{env}}(s, e) \quad s'.\text{irq} = \text{true} \quad (s'', \ell_i'') = \text{intr_handler}(s', \ell_i, \ell^{env}) \quad (s''', \ell_i''') = \text{intr_disable}(s'', \ell_i'', \ell^{env})}{\text{intr_disable}(s, \ell_i, \ell^{env}) = (s''', \ell_i''')}$$

21

Semantics of *intr_enable*

- Recursively discharges pending interrupts.
- Delayed interrupts that occur while the interrupt is disabled.



details in the paper

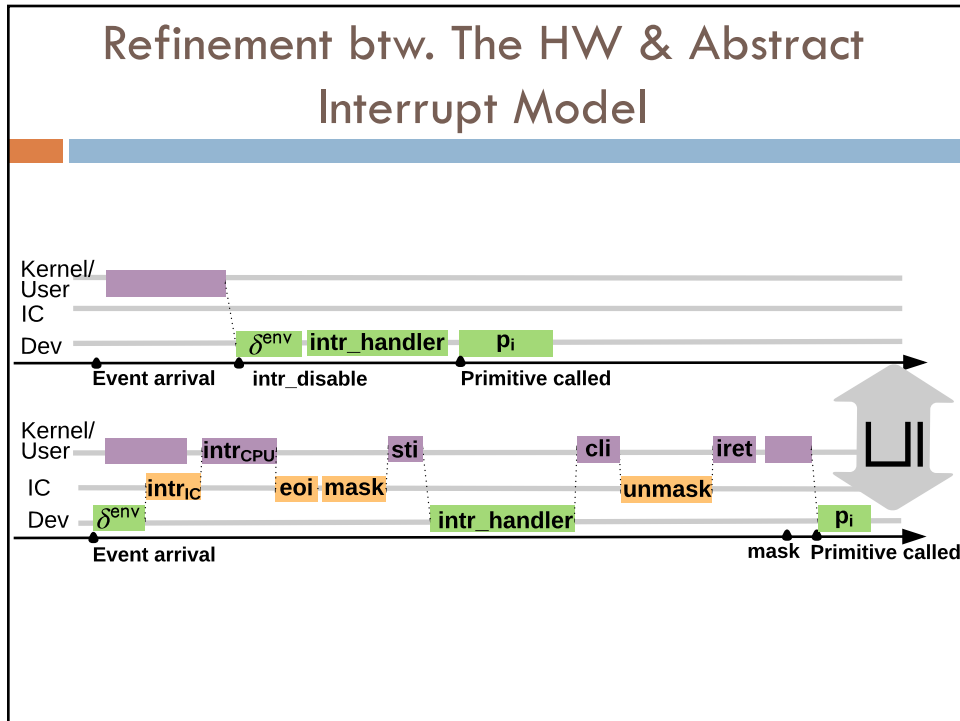
ENABLENOINTR: Enable with no pending interrupt

$$\frac{s.\text{irq} = \text{false} \quad s' = s[\text{Flag} \leftarrow 1]}{\text{intr_enable}(s, \ell_i, \ell^{env}) = (s', \ell_i)}$$

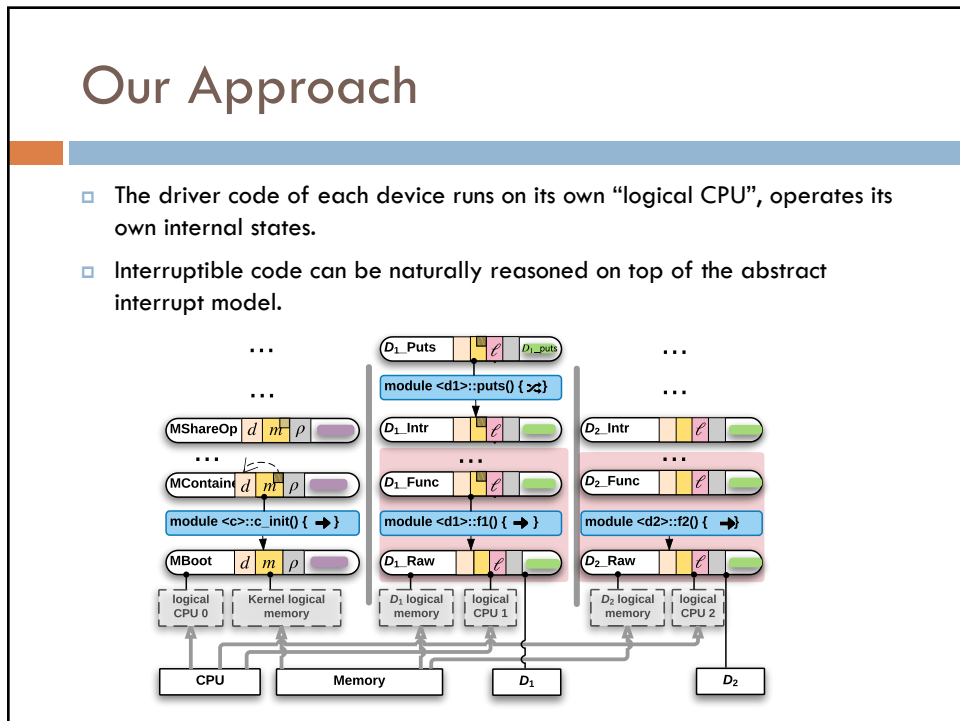
ENABLEINTR: Enable with pending interrupts

$$\frac{s.\text{irq} = \text{true} \quad (s', \ell_i') = \text{intr_handler}(s, \ell_i, \ell^{env}) \quad (s'', \ell_i'') = \text{intr_enable}(s', \ell_i', \ell^{env})}{\text{intr_enable}(s, \ell_i, \ell^{env}) = (s'', \ell_i'')}$$

22



23



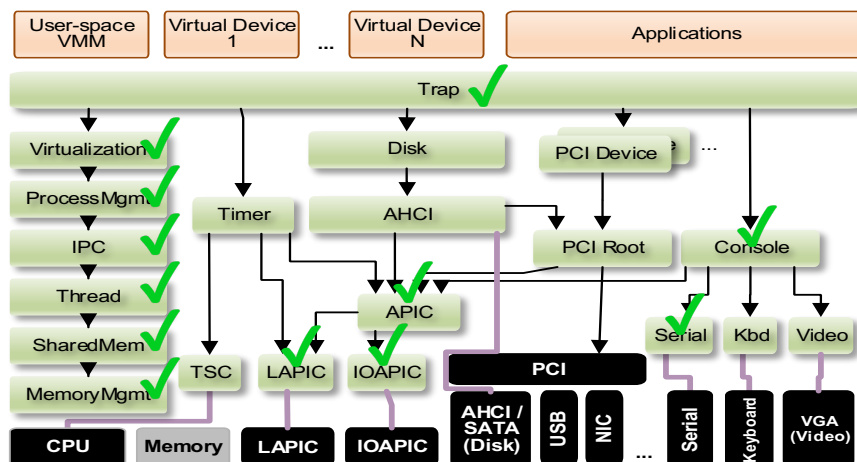
24

Our Contributions

- ▣ New techniques for certifying abstraction layers with multiple *logical CPUs* and devices.
- ▣ New techniques for building formal *certified device hierarchies*.
- ▣ An abstraction-layer-based approach for reasoning about *interrupts*.
- ▣ **Case study:** interruptible mCertikOS with device drivers.

25

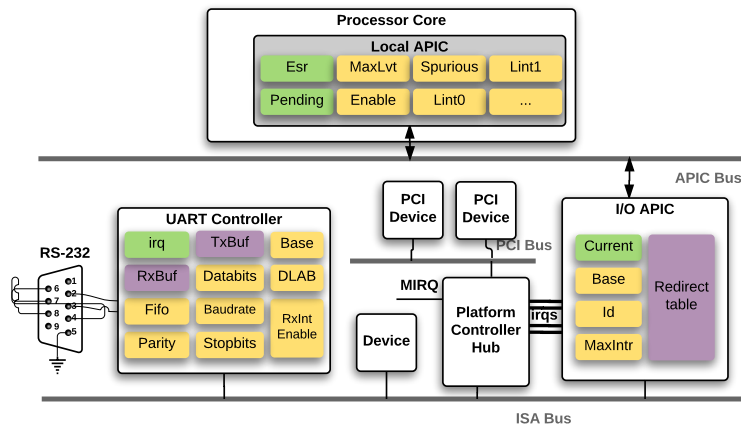
Interruptible mCertikOS with Drivers



26

Case Study: Modeling HW Devices

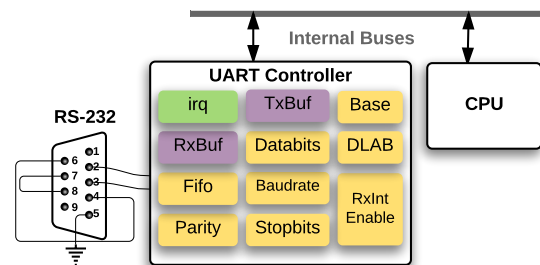
- Serial Port, I/O APIC, Local APIC, CPU interrupt handling.



27

Case Study: Serial Device

- States: see figure
- Transitions: `serial_trans_env` + `serial_trans_IO`
- Read/Write primitives: `serial_read` / `serial_write`



28

Serial Interrupt Handler

```

1 void serial_intr () {
2   unsigned int hasMore;
3   int t = 0;
4   hasMore = serial_getc ();
5   while (hasMore && t < CONSOLE_BUFFER_SIZE) {
6     hasMore = serial_getc ();
7     t++;
8   }
9 }

1 unsigned int serial_getc () {
2   unsigned int rv = 0;
3   unsigned int rx;
4   if (serial_exists()) {
5     if (serial_read(COM1 + COM_LSR, BIT1) % 2 == 1)
6       {
7         rx = serial_read(COM1 + COM_RX, M_ALL);
8         cons_buf_write(rx);
9         rv = 1;
10        }
11    }
12   return rv;
13 }

```

29

Serial Driver

```

1 void serial_puts(char * s, int len) {
2   int i = 0;
3   while (i < len && s[i] != 0) {
4     serial_intr_disable ();
5     serial_putc (s[i]);
6     serial_intr_enable ();
7     i++;
8   }
9 }

1 void serial_putc (unsigned int c) {
2   unsigned int lsr = 0, i;
3   if ( serial_exists() ){
4     for (i = 0; !lsr && i < 12800; i++) {
5       lsr = serial_read(0x3FD) & 0x20;
6       delay();
7     }
8     serial_write (0x3F8, c);
9     ...

```

30

What We Have Proved

- Total functional correctness.
- Safety.
- *Contextual refinement* between the lowest and the top level abstract machine:

$$\forall P, \llbracket K \bowtie P \rrbracket_{x86} \sqsubseteq \llbracket P \rrbracket_{mCertiKOS}$$
- Data invariants:
 - Console's circular buffer is always well-formed.
 - Interrupt controller states are always consistent.
- The framework also ensures that:
 - No code injection attacks, buffer overflow, integer overflow, null pointer access, etc.

31

Size of TCB and Spec/Proof

- In the TCB
 - X86 hardware model
 - Hardware device/interrupt model (510 LOC)
 - System call specification (126 LOC)
 - Bootloader
 - Coq proof checker
 - Pretty-printing phase of the CompCert compiler
- Rest of the spec/proof (about 20k LOC)
 - Intermediate and auxiliary specifications and definitions
 - Coq proof scripts

32

Conclusion

- Compositional framework for building certified interruptible kernel with device drivers.
 - Certified abstraction layers with multiple logical CPUs.
 - An abstraction-layer-based approach for expressing interrupts.
- The first formally verified interruptible OS kernel with device drivers.
- Extensions:
 - Other drivers
 - Concurrency
 - Larger kernel

