

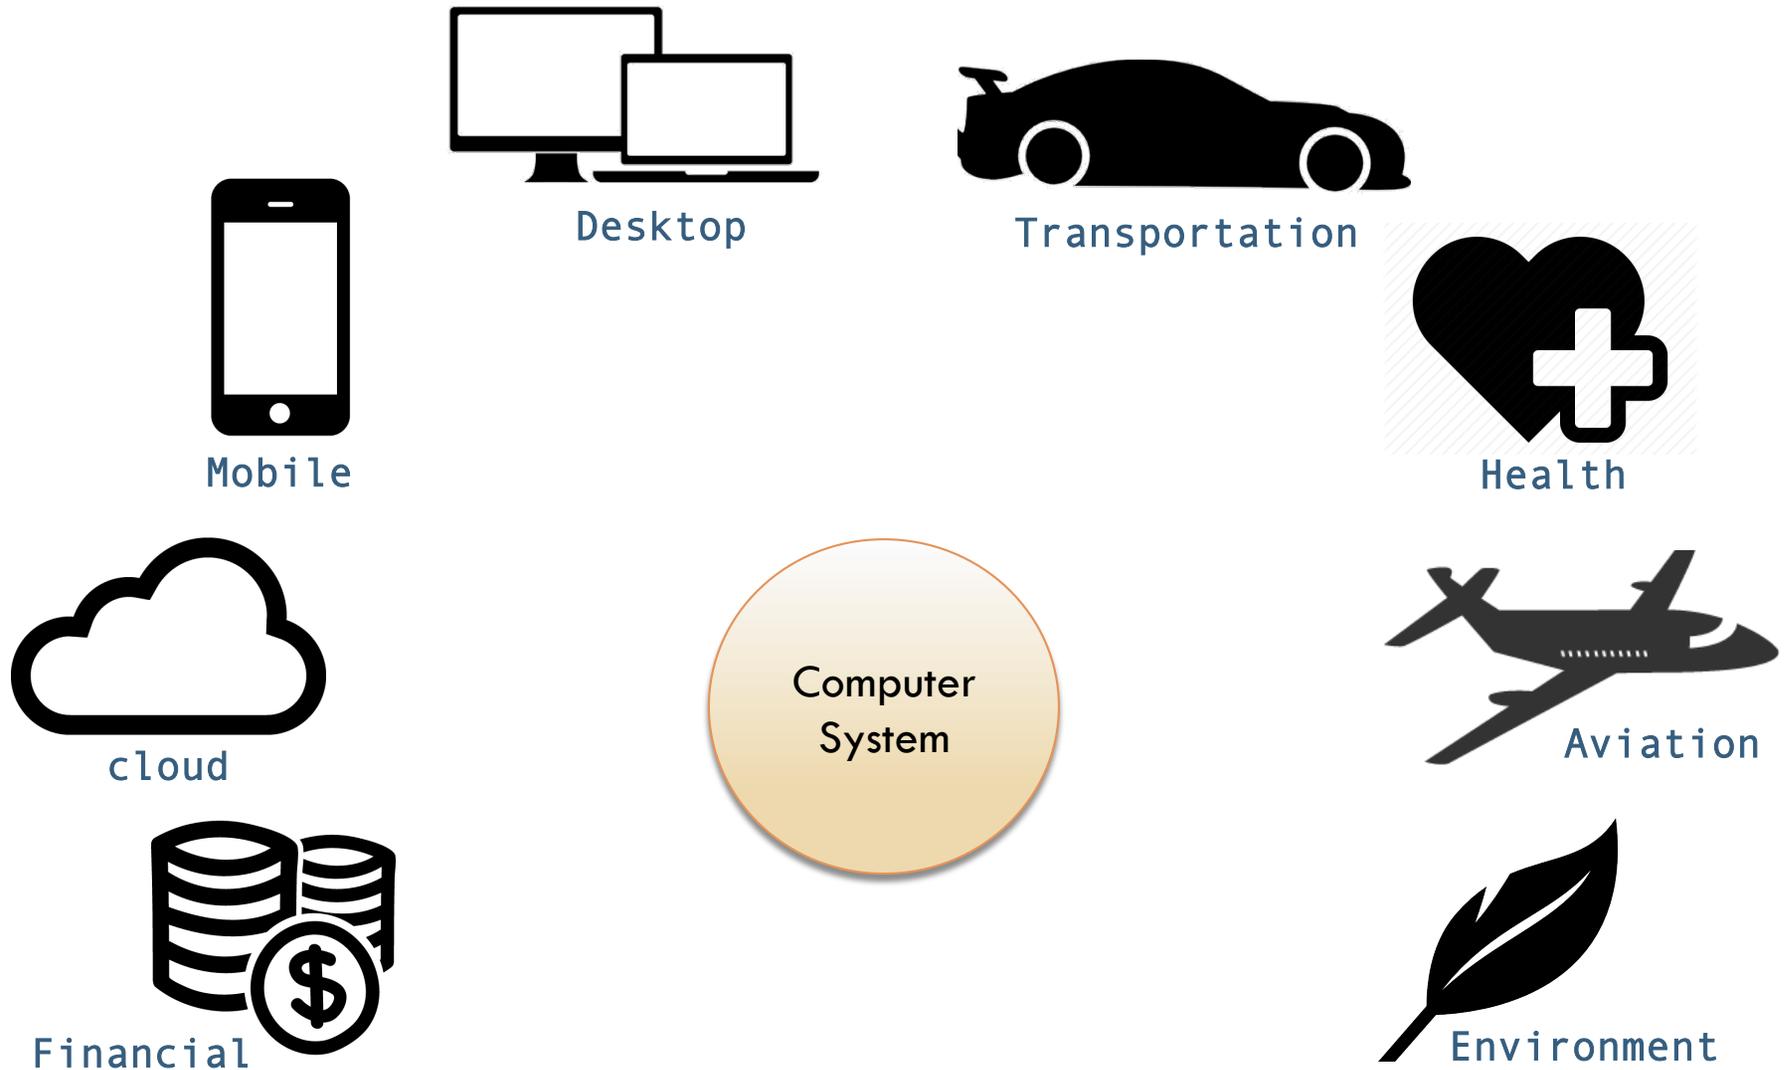
TOWARD COMPOSITIONAL VERIFICATION OF INTERRUPTIBLE OS KERNELS AND DEVICE DRIVERS

Xiongnan (Newman) Wu

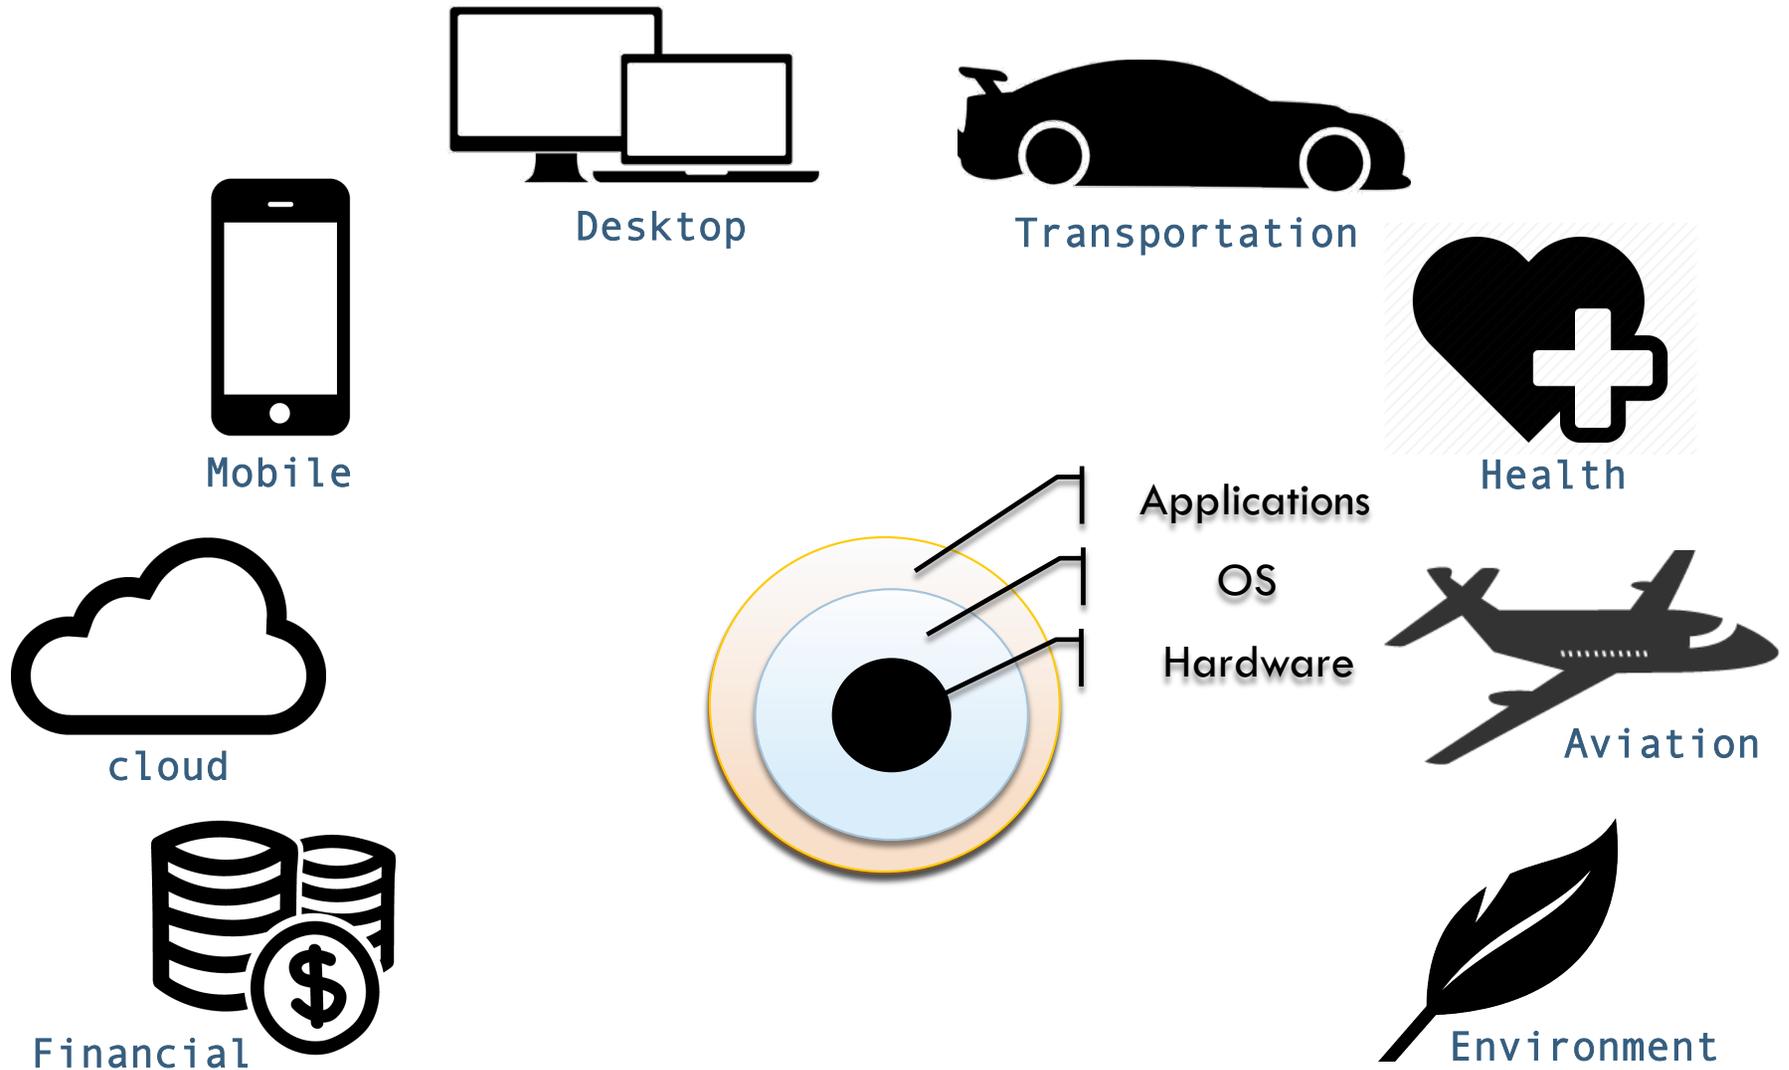
Joint work with Hao Chen, Zhong Shao, Joshua Lockerman, and Ronghui Gu

Yale University

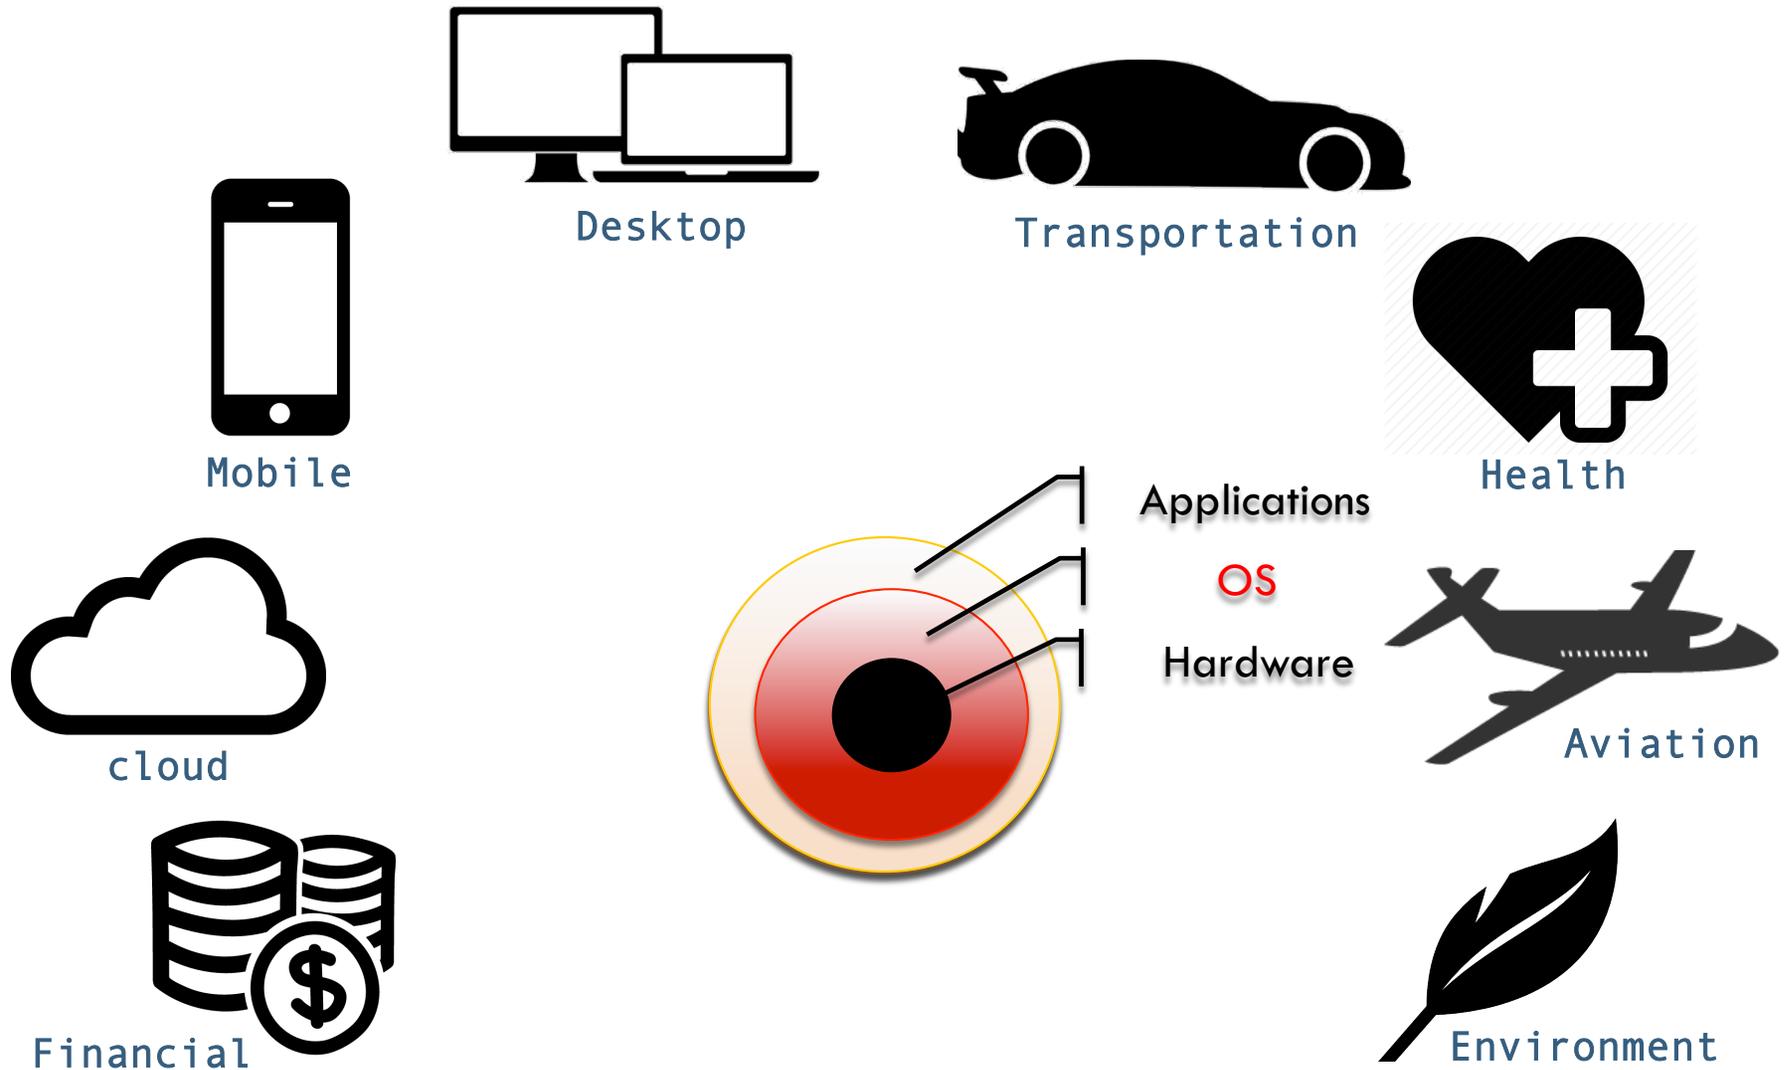
Do we really need high-assurance OS?



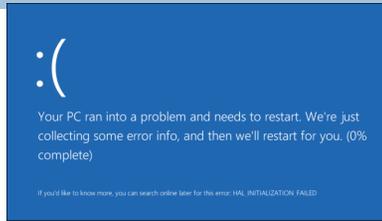
Do we really need high-assurance OS?



Do we really need high-assurance OS?



Do we really need high-assurance OS?



Crash



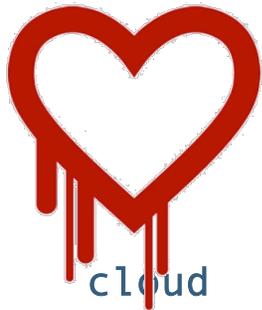
Accident



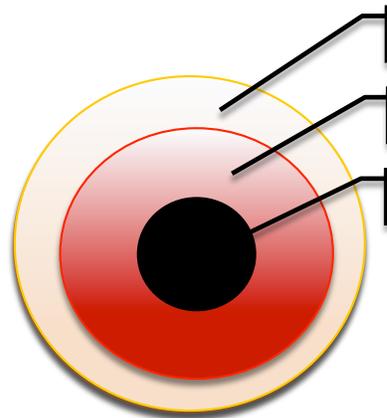
Mobile



Life



cloud



Applications

OS

Hardware



Loss



Financial



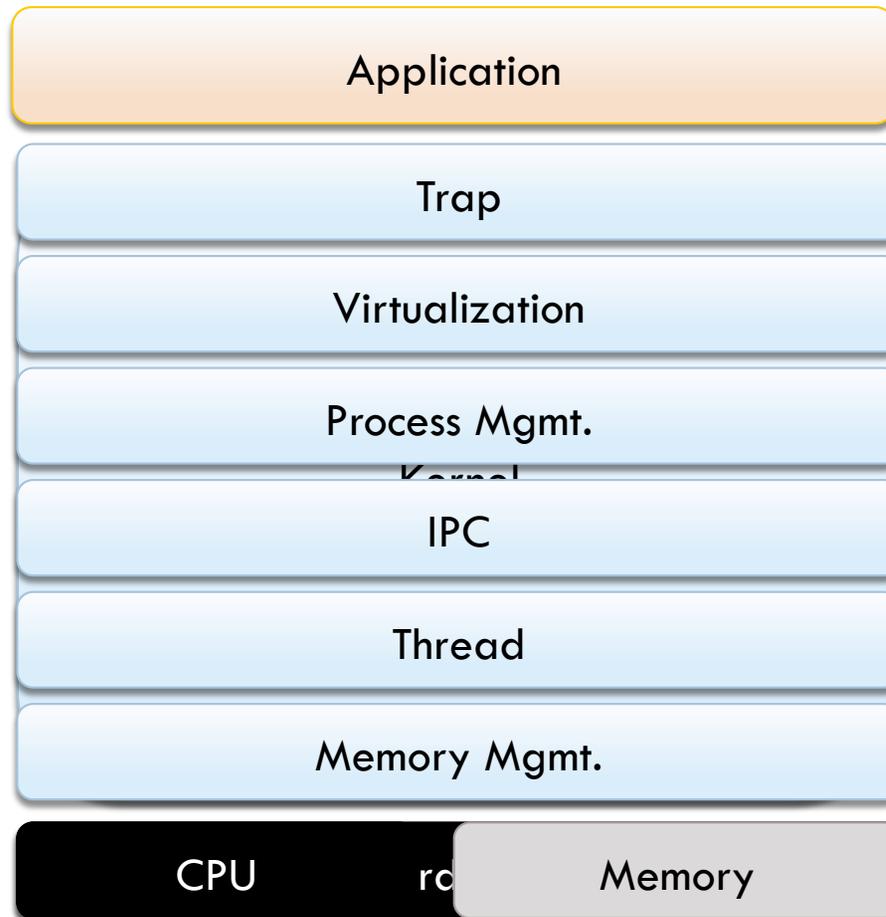
Environment

Formal Verification of OS Kernel

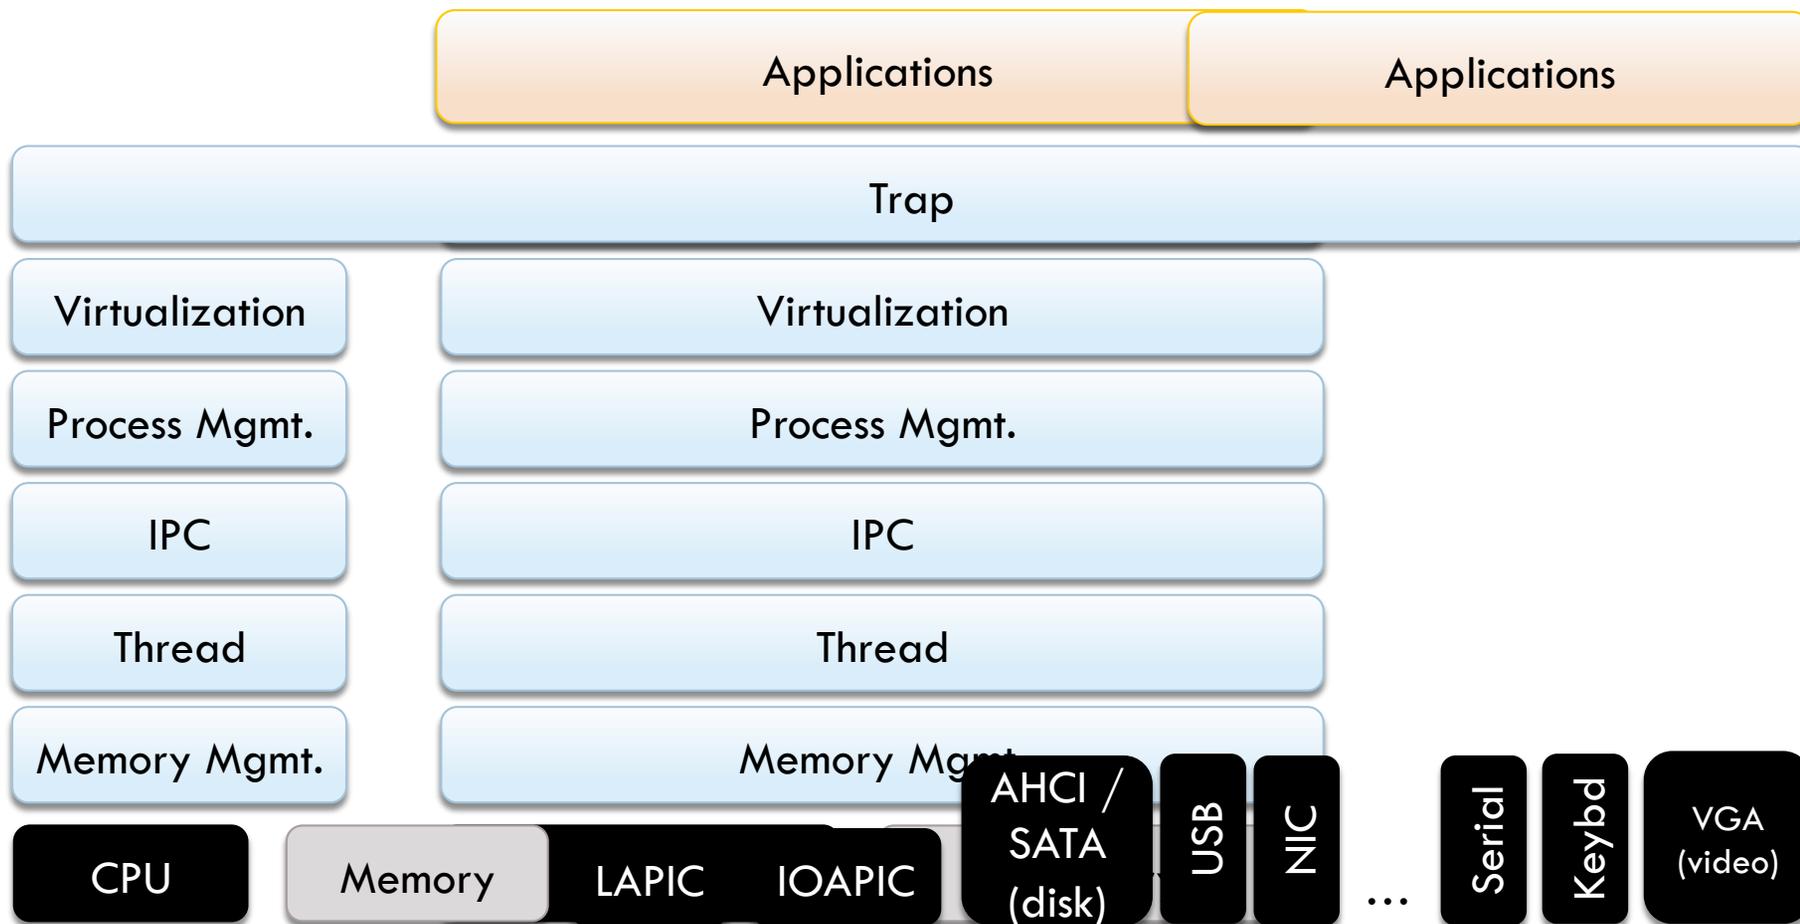
□ seL4

□ CertiKOS

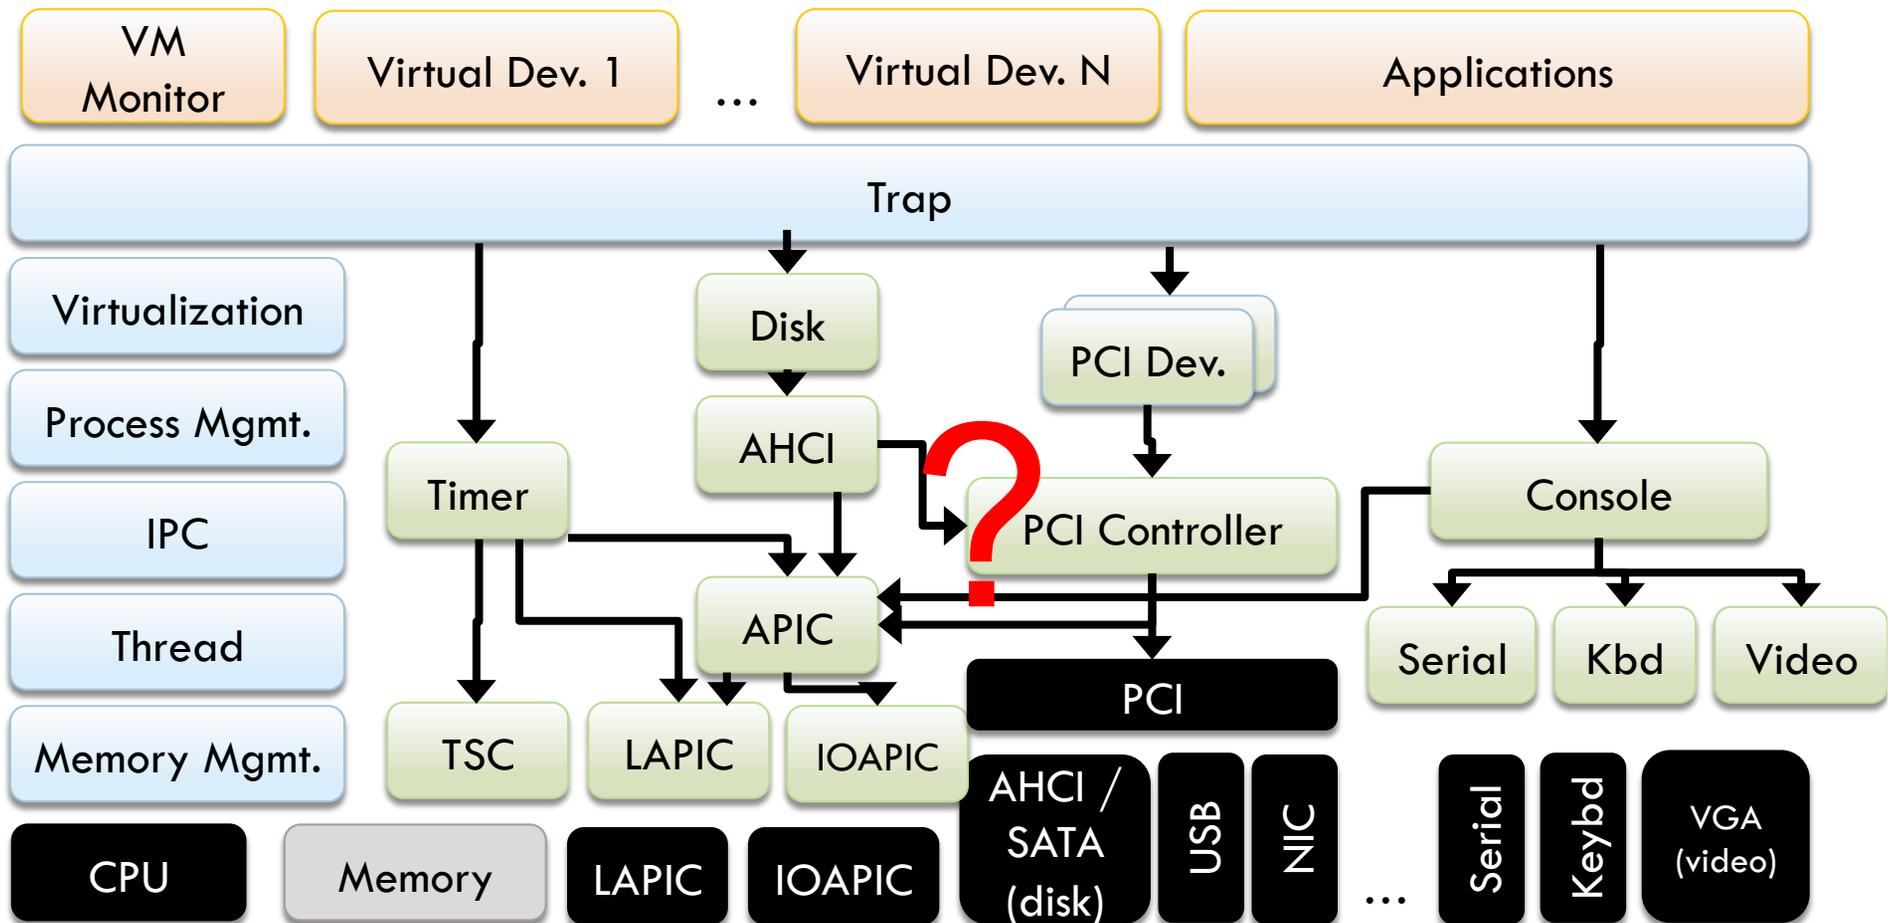
□ Verve



Formal Verification of OS Kernel



Formal Verification of OS Kernel



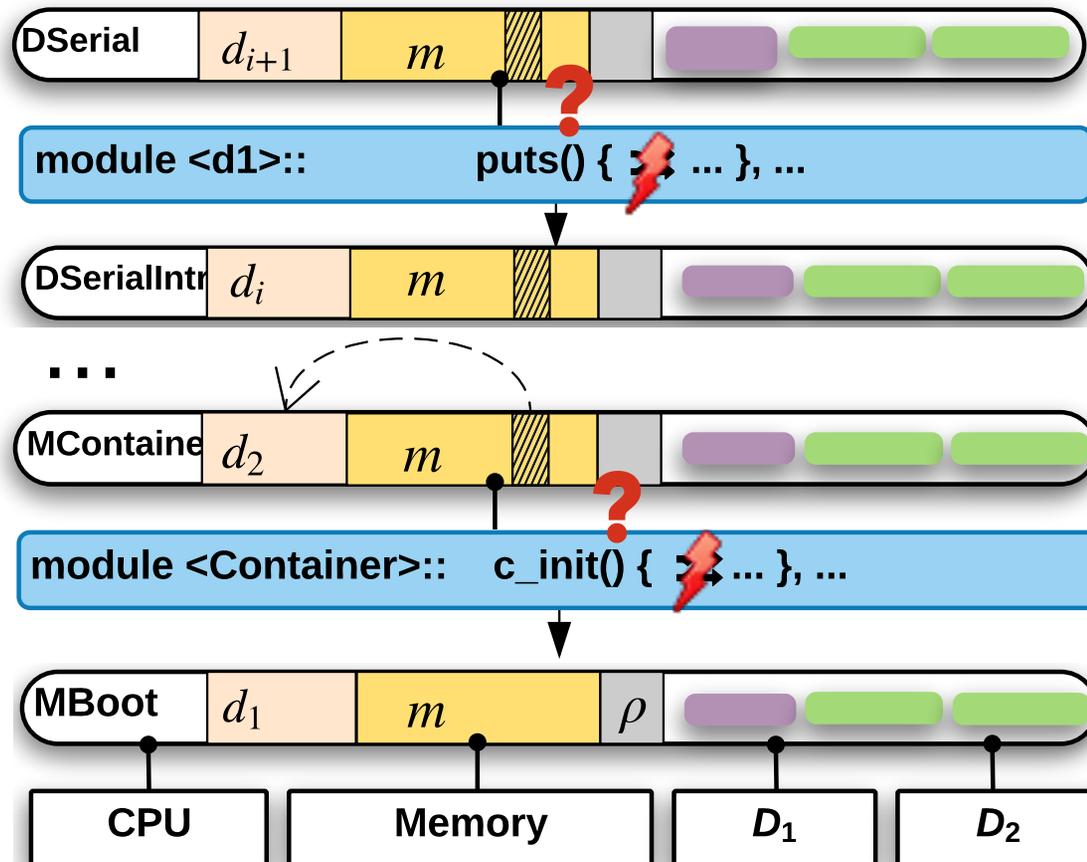
Device Drivers in Mainstream OS

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



Main Challenge

Every fine-grained processor step could be interrupted.



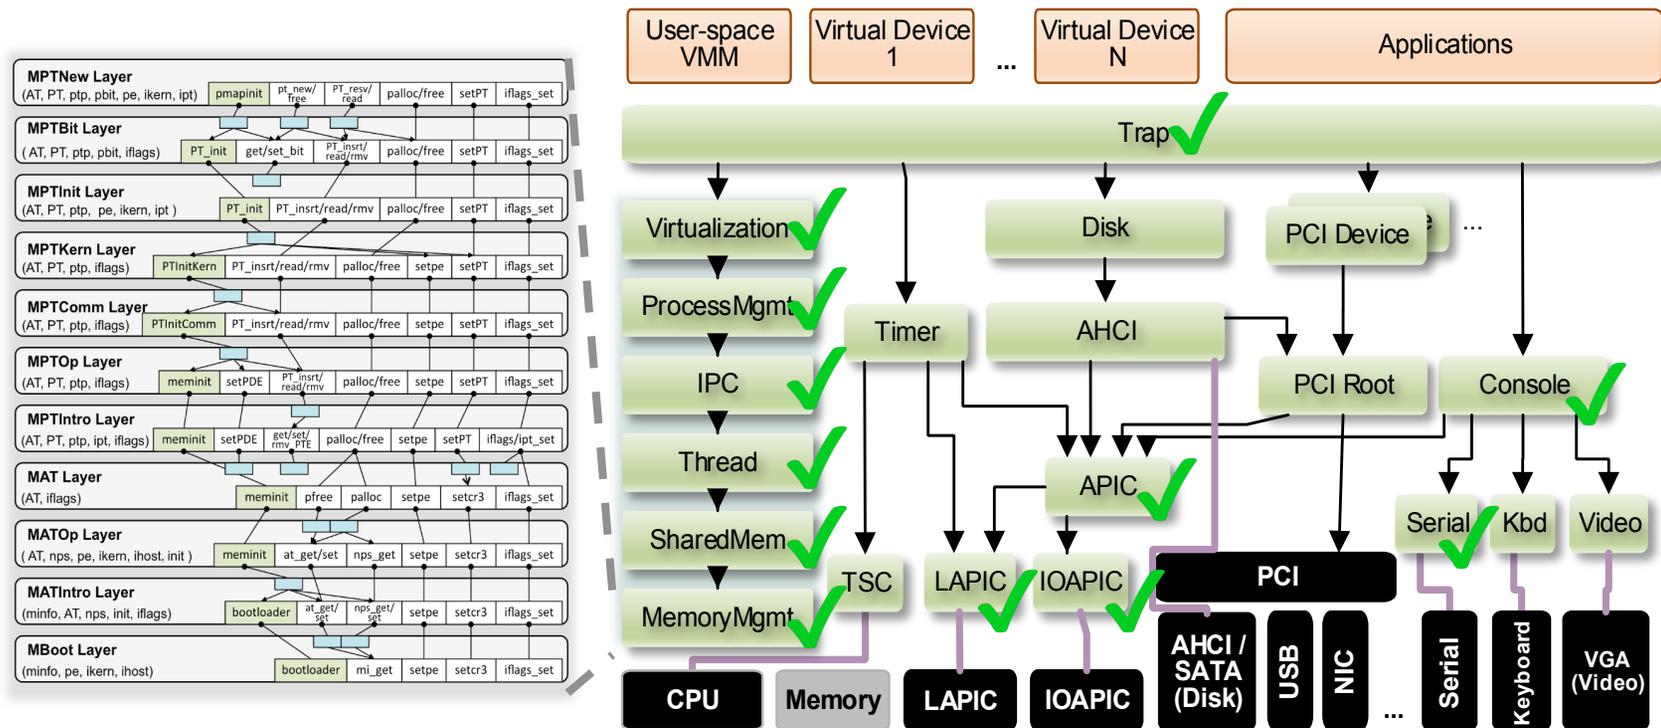
Other Challenges



- ❑ 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*.

Our Contributions

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



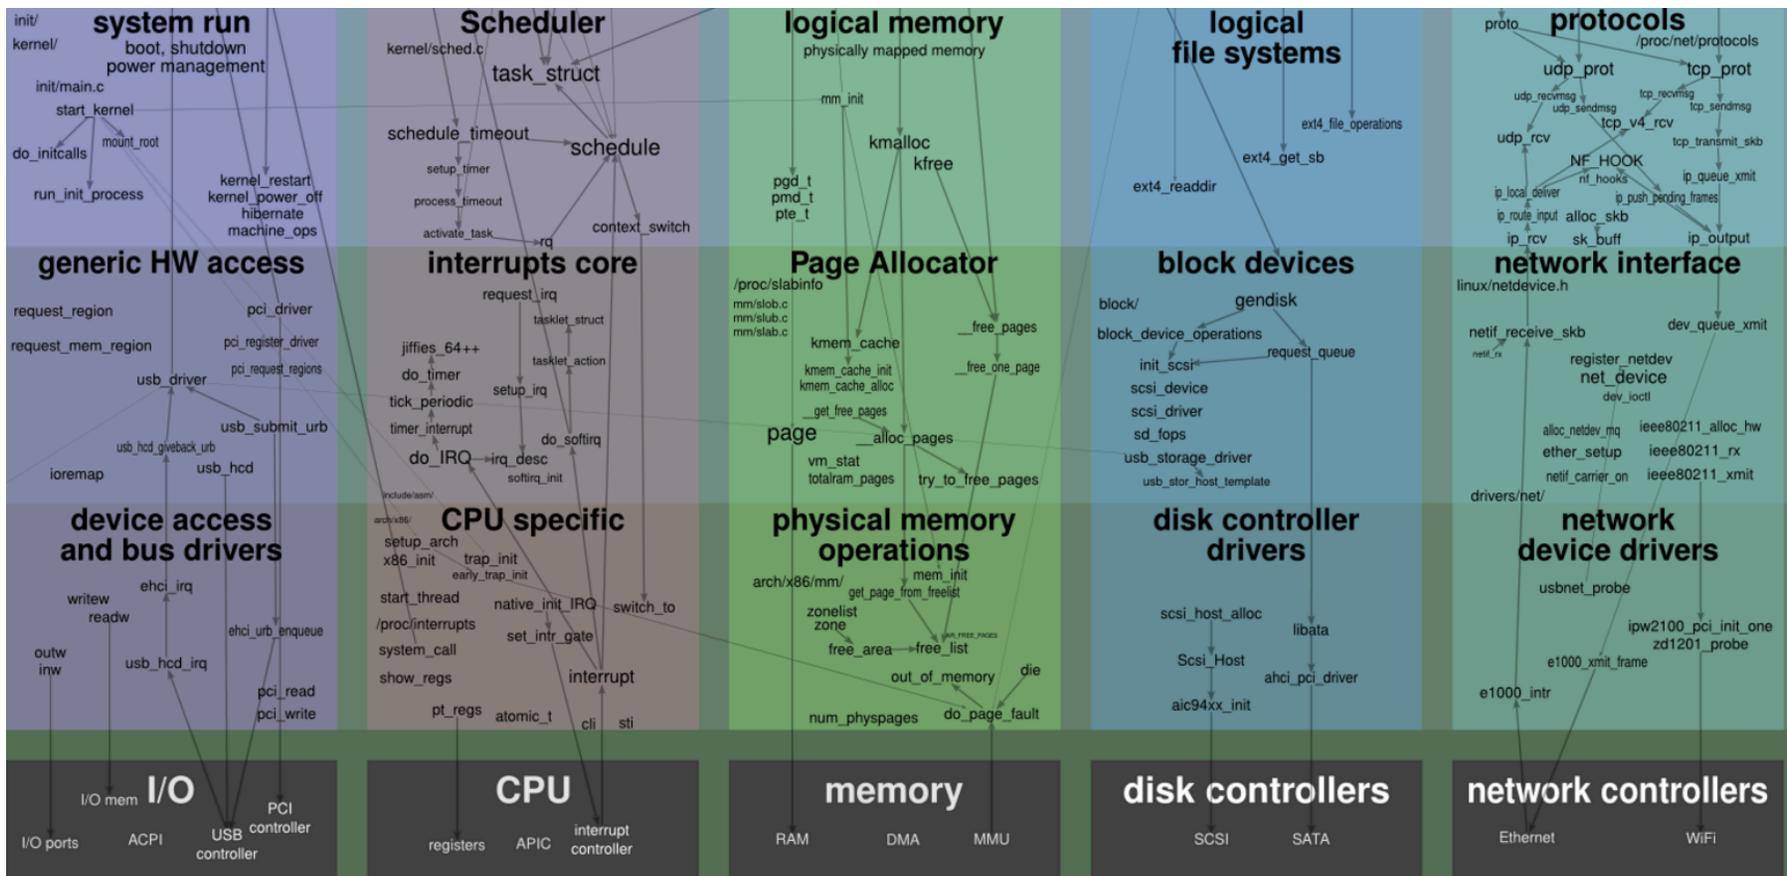
Our Contributions



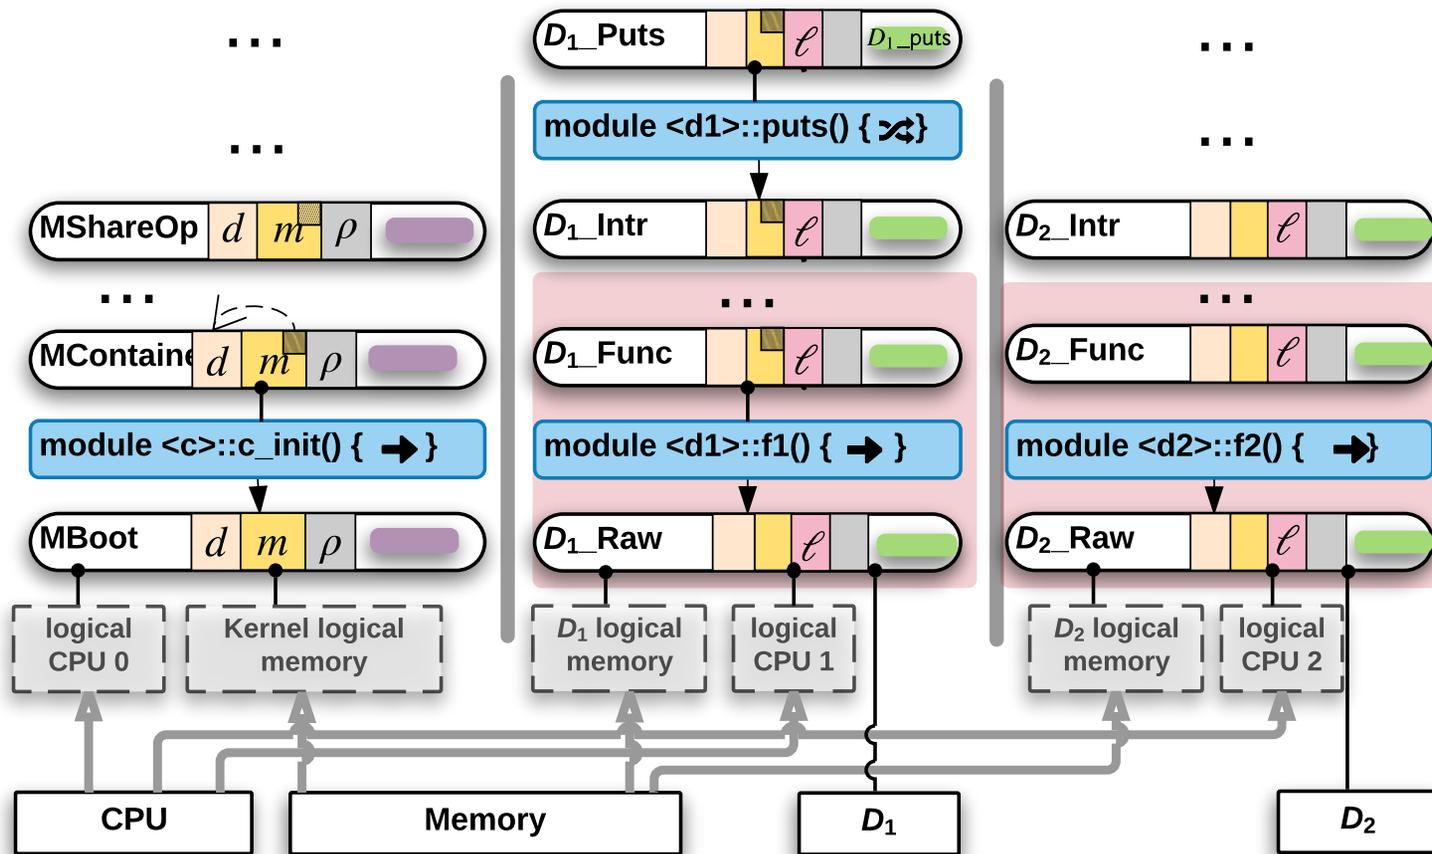
- New techniques for certifying abstraction layers with multiple *logical CPUs* and devices.
- New techniques for building formal *certified device hierarchies*.
- An abstraction-layer-based approach for reasoning about *interrupts*.
- **Case study:** interruptible mCertiKOS with device drivers.

Linux Kernel Map

Kernel components are sorted into different stacks of abstraction layers based on their underlying hardware device.



New Machine Model



Our Contributions



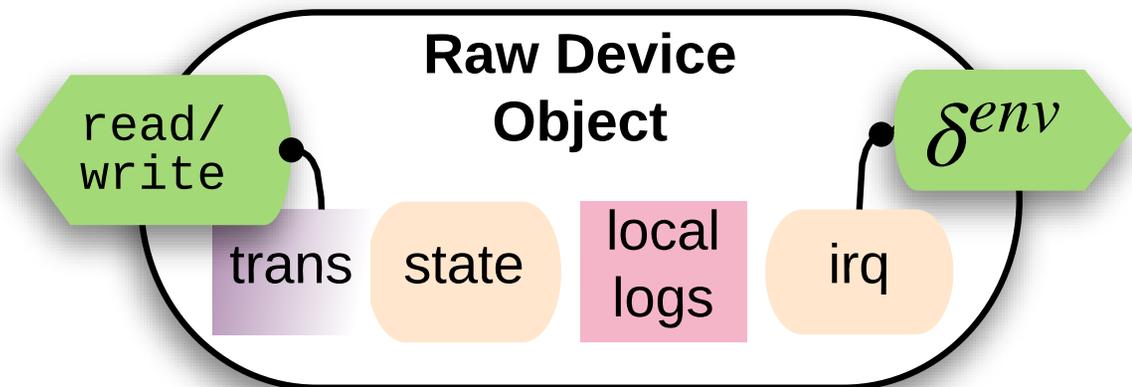
- New techniques for certifying abstraction layers with multiple *logical CPUs* and devices.
- New techniques for building formal *certified device hierarchies*.
- An abstraction-layer-based approach for reasoning about *interrupts*.
- **Case study:** interruptible mCertikOS with device drivers.

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}

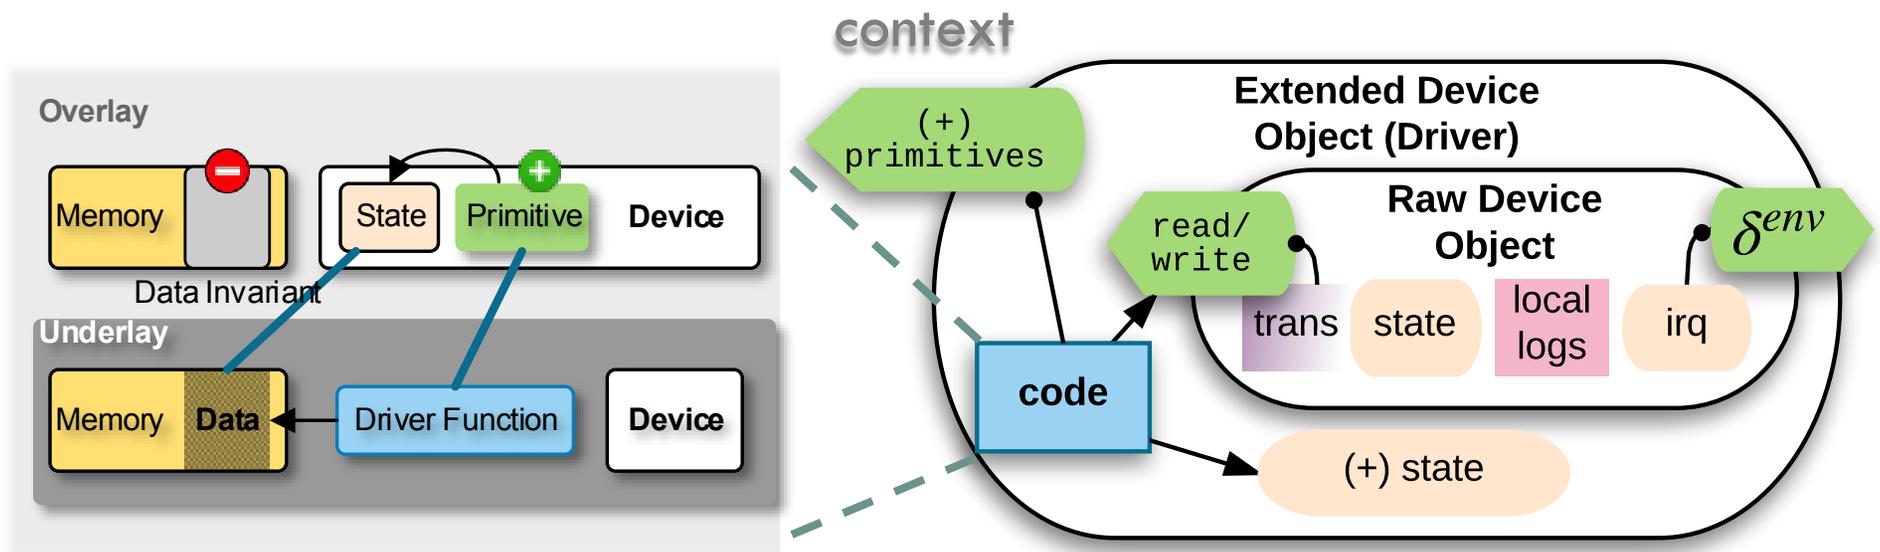
Raw Device Object

- ▣ 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.

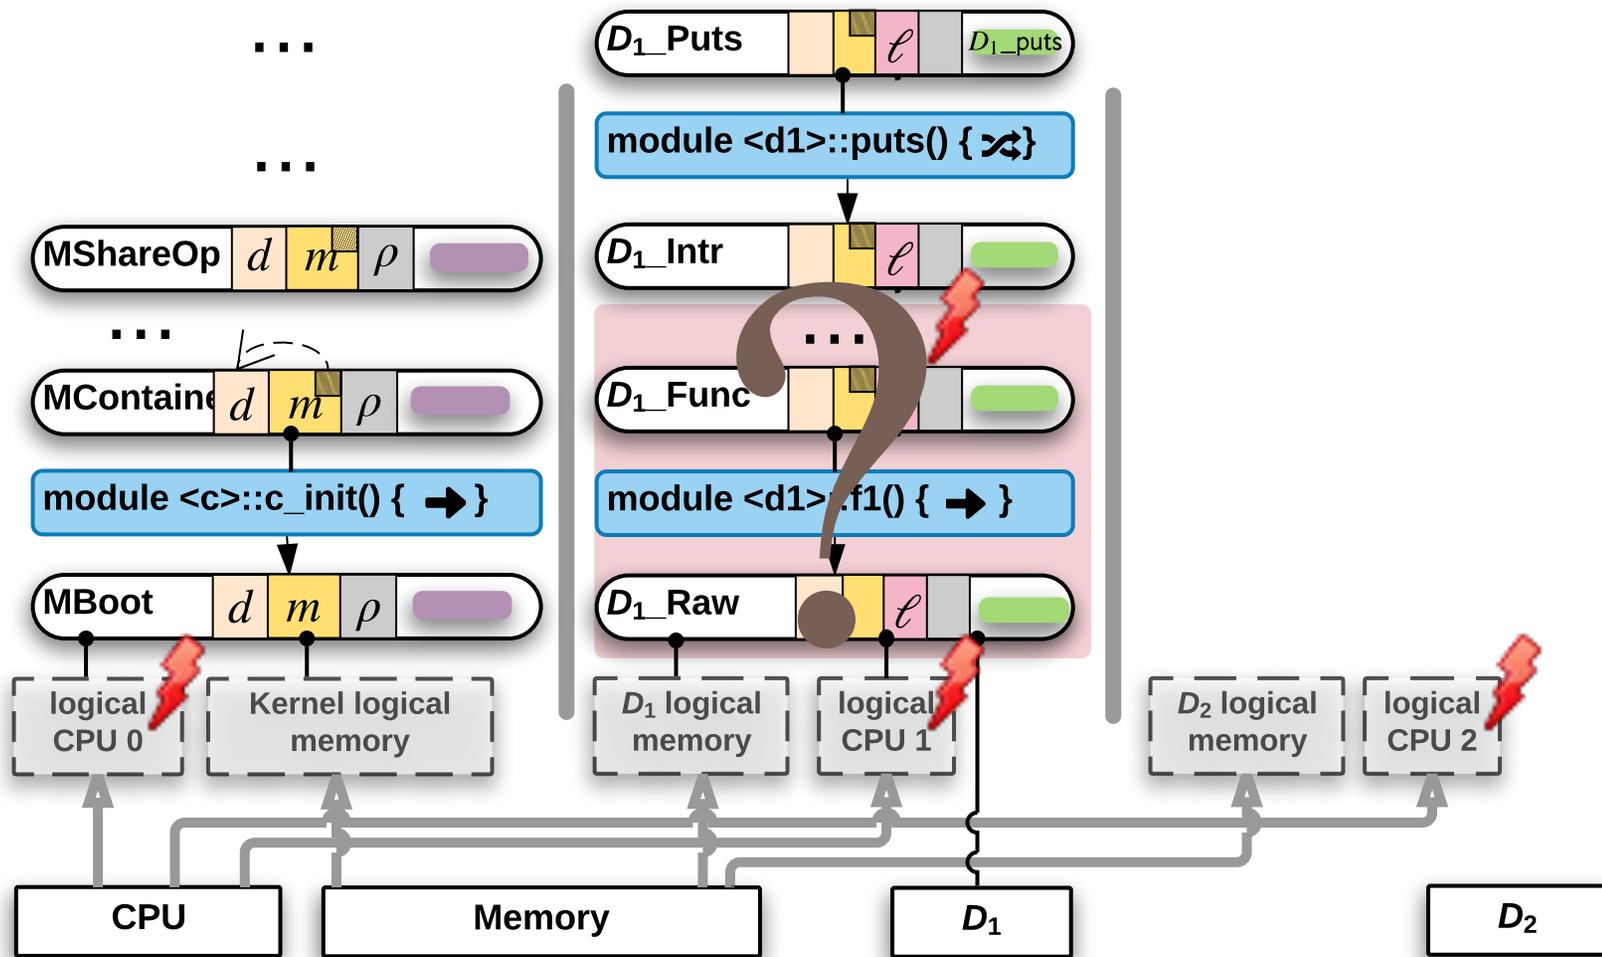


Extended Device Object

Driver as a logical device.



Recap: Machine Model

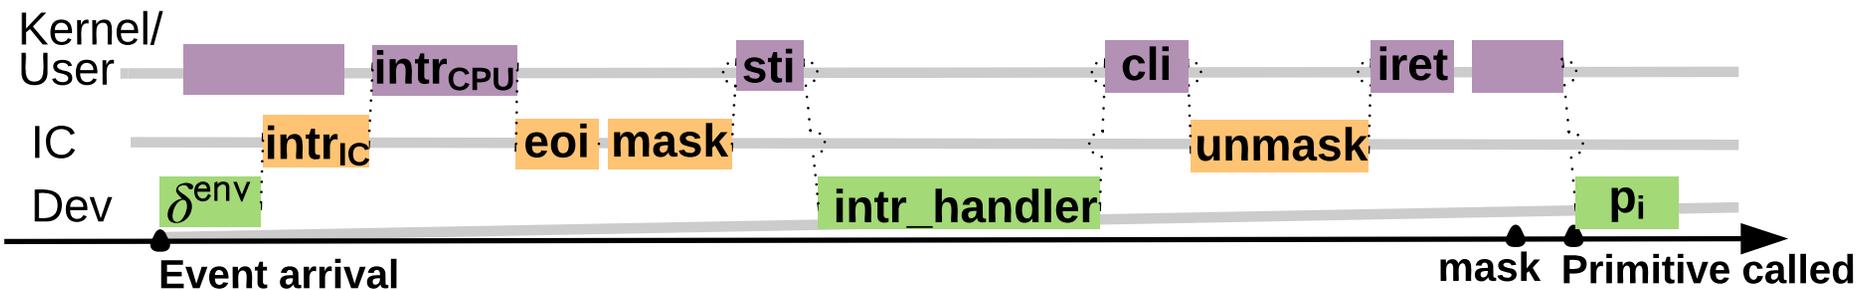


Our Contributions

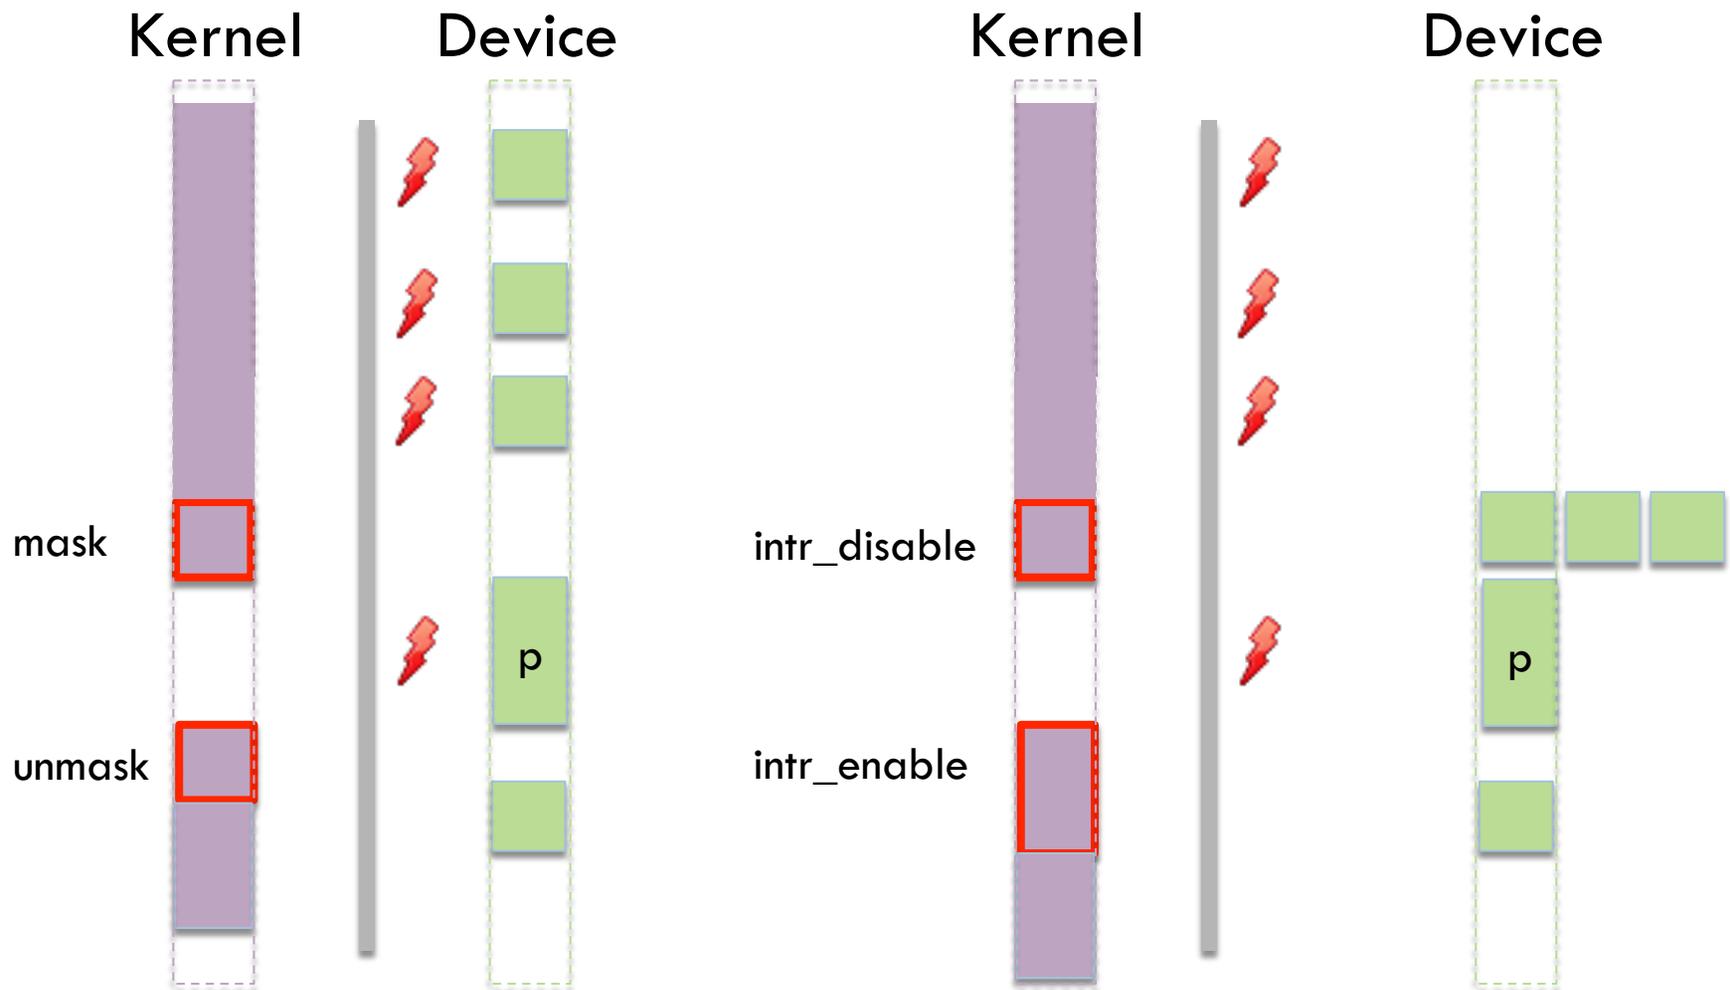


- ▣ New techniques for certifying abstraction layers with multiple *logical CPUs* and devices.
- ▣ New techniques for building formal *certified device hierarchies*.
- ▣ An abstraction-layer-based approach for reasoning about *interrupts*.
- ▣ **Case study:** interruptible mCertiKOS with device drivers.

Interrupt Models

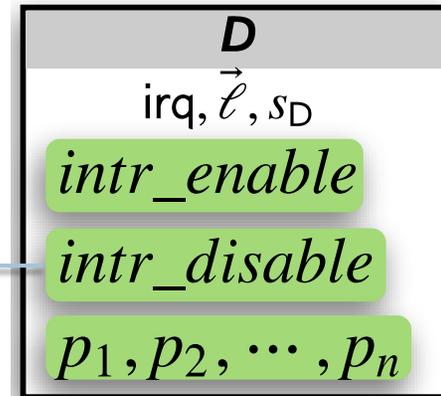


New HW Interrupt Model



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

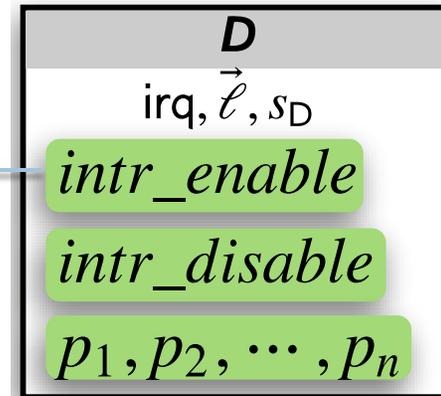
$$\frac{(e, l'_i) = \text{next}(l^{env}, l_i) \quad s_{tmp} = \delta^{env}(s, e) \quad s_{tmp}.irq = \text{false} \quad s' = s[\text{iFlag} \leftarrow 0]}{\text{intr_disable}(s, l_i, l^{env}) = (s', l_i)}$$

DISABLEINTR: Disable with unhandled interrupts

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

Semantics of *intr_enable*

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



details in the paper

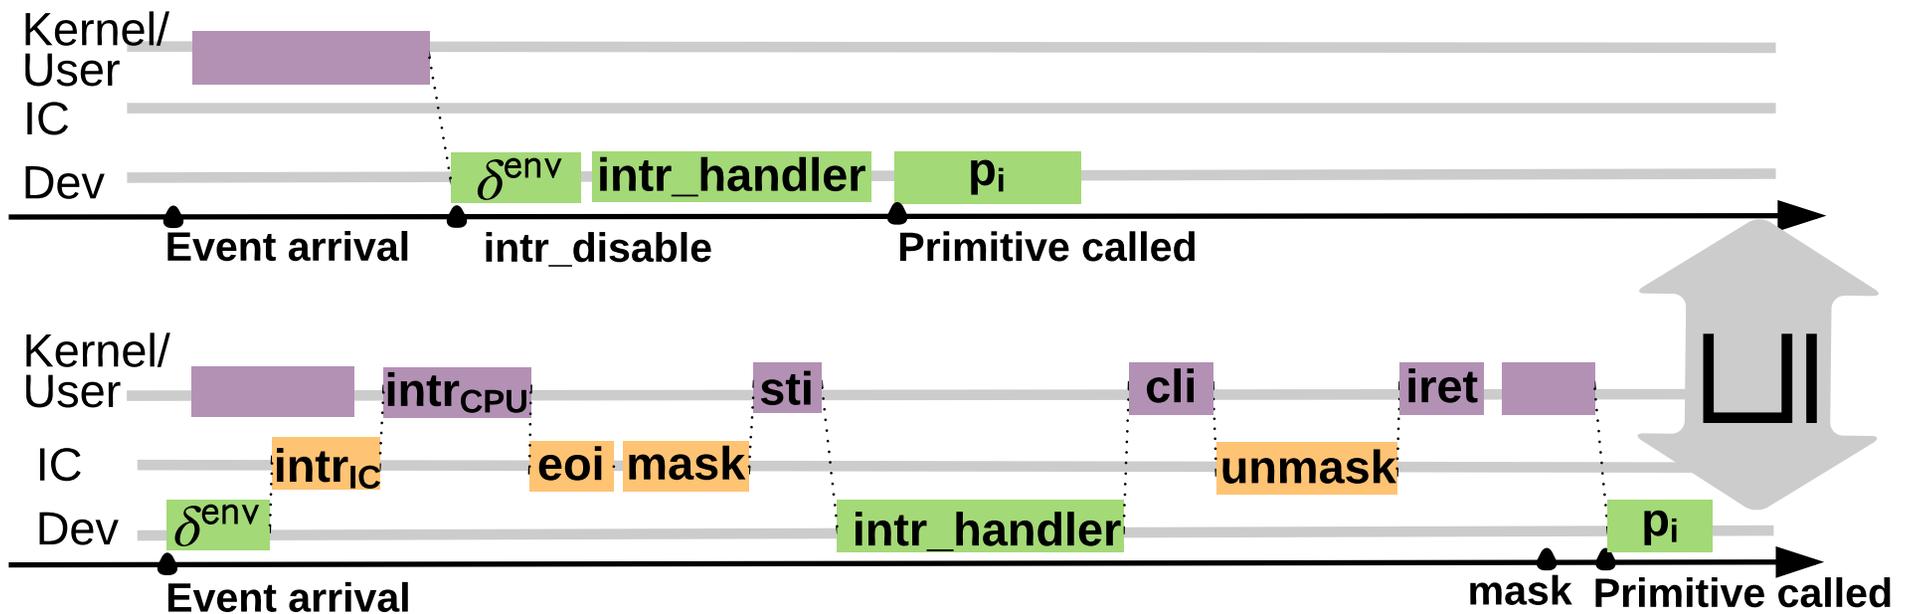
ENABLENOINTR: Enable with no pending interrupt

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

ENABLEINTR: Enable with pending interrupts

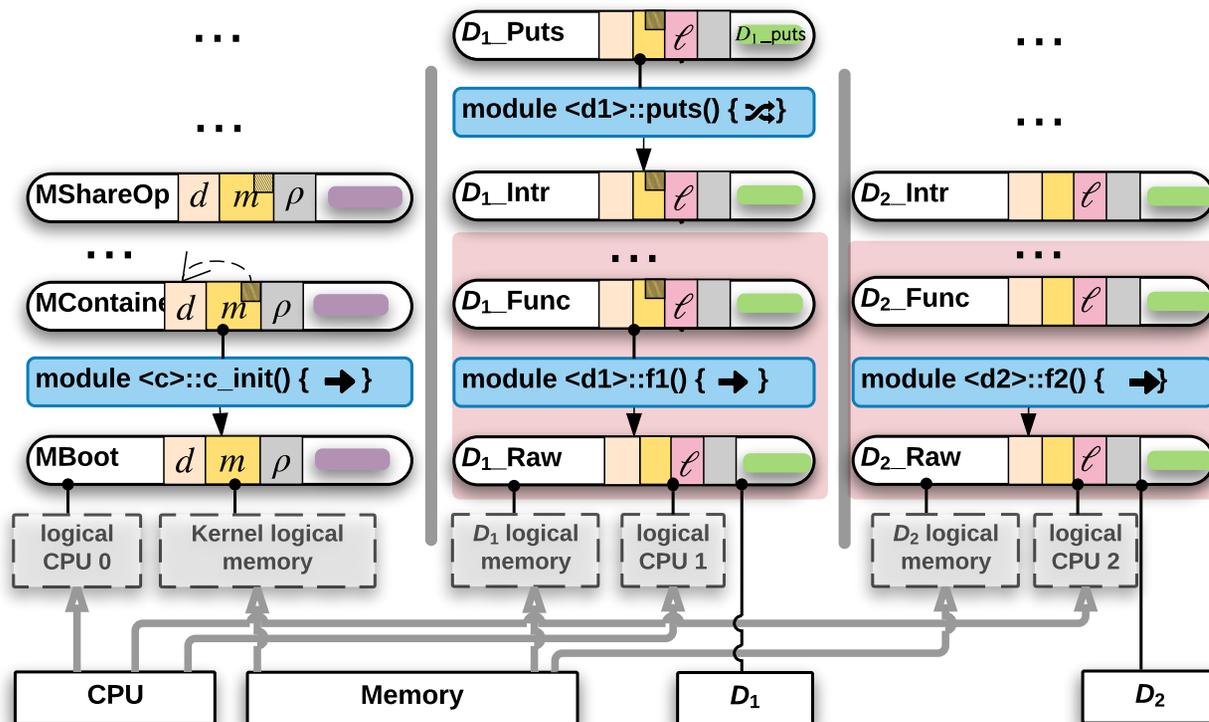
$$\frac{s.\text{irq} = \text{true} \quad (s', \ell'_i) = \text{intr_handler}(s, \ell_i, \ell^{env}) \quad (s'', \ell''_i) = \text{intr_enable}(s', \ell'_i, \ell^{env})}{\text{intr_enable}(s, \ell_i, \ell^{env}) = (s'', \ell''_i)}$$

Refinement btw. The HW & Abstract Interrupt Model



Our Approach

- The driver code of each device runs on its own “logical CPU”, operates its own internal states.
- Interruptible code can be naturally reasoned on top of the abstract interrupt model.

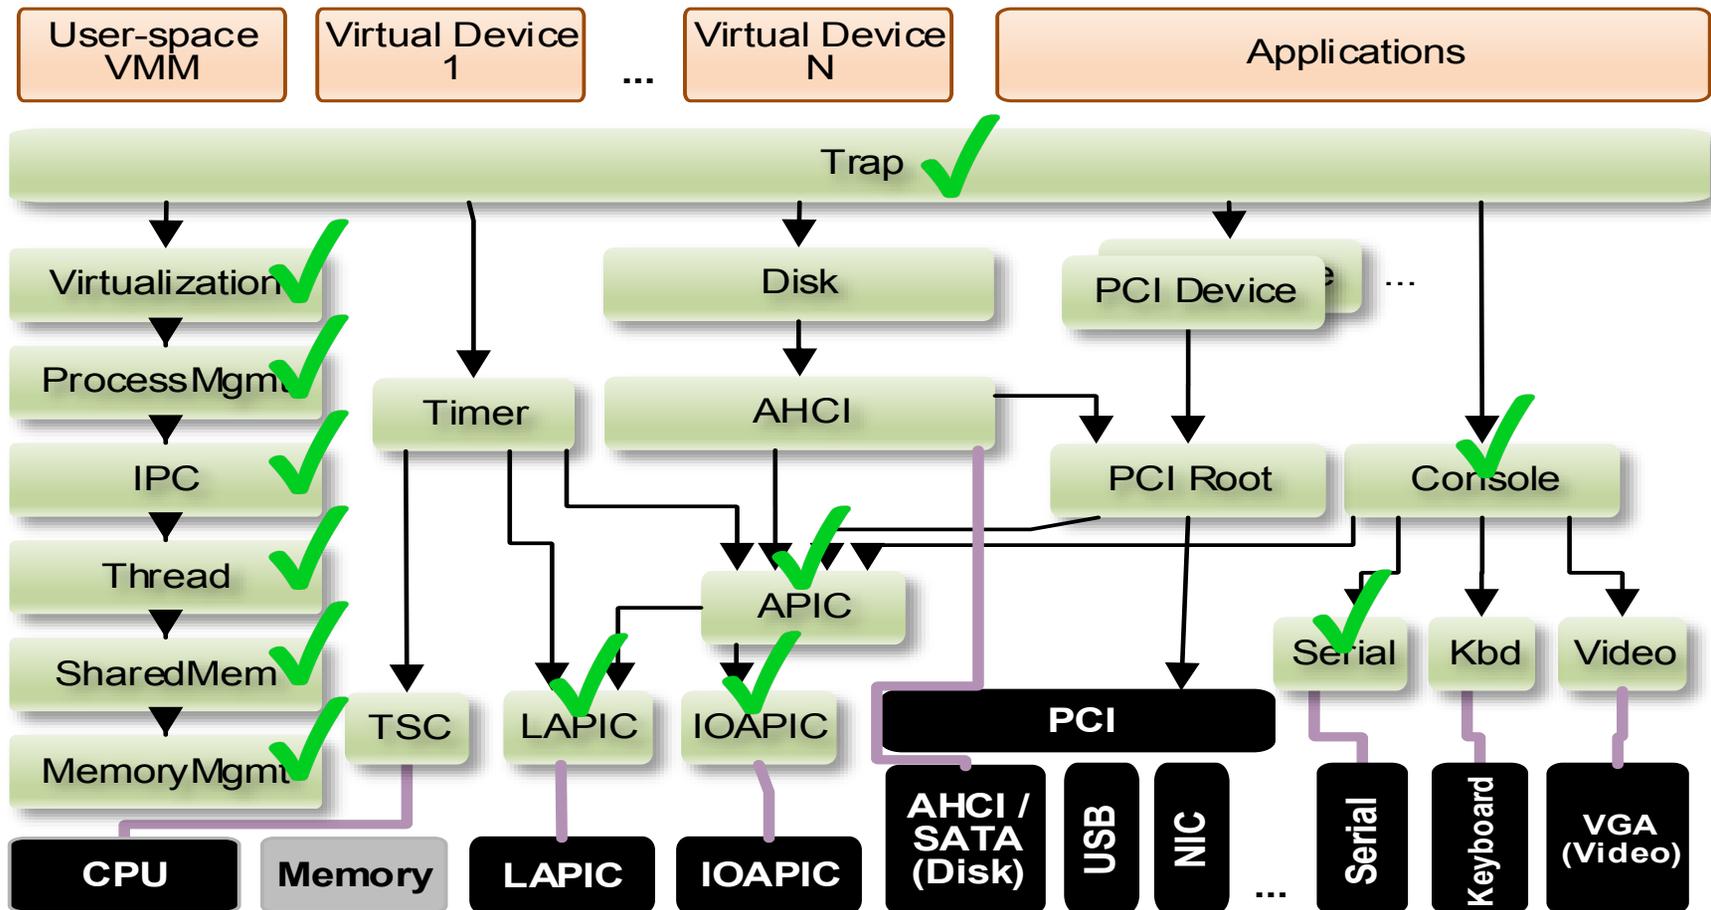


Our Contributions



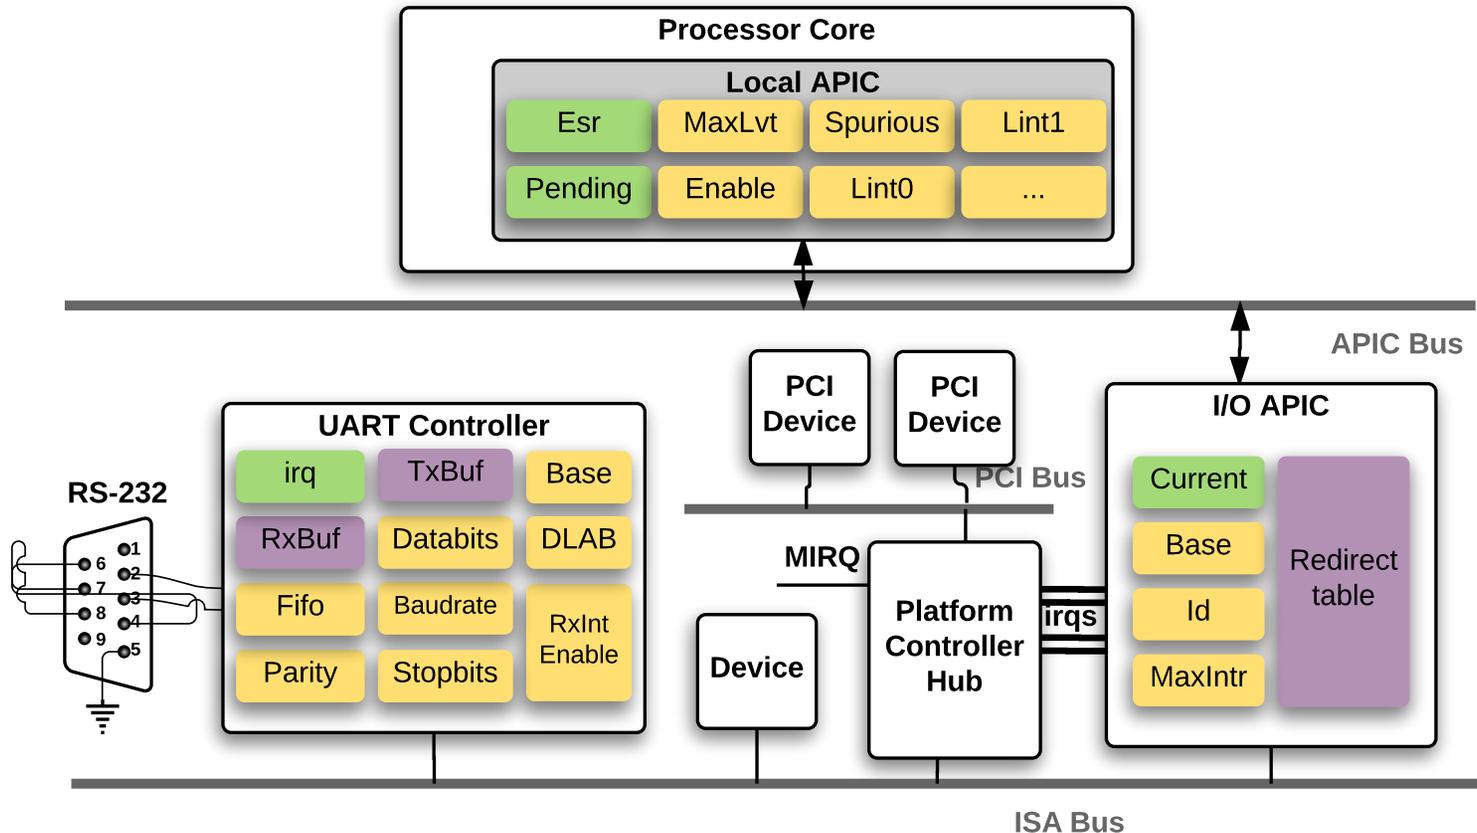
- New techniques for certifying abstraction layers with multiple *logical CPUs* and devices.
- New techniques for building formal *certified device hierarchies*.
- An abstraction-layer-based approach for reasoning about *interrupts*.
- **Case study:** interruptible mCertikOS with device drivers.

Interruptible mCertKOS with Drivers



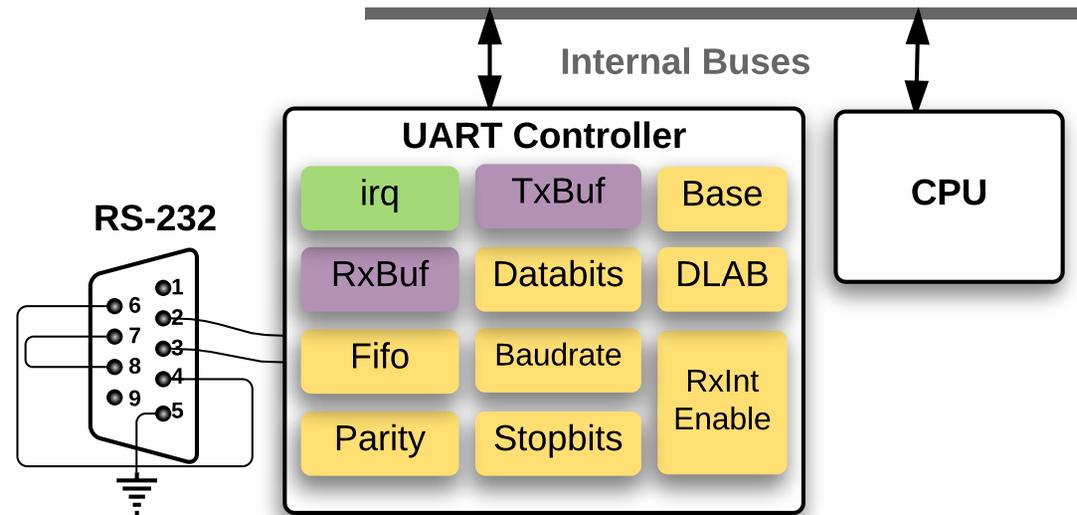
Case Study: Modeling HW Devices

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



Case Study: Serial Device

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



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 }
```

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         ...
}
```

What We Have Proved

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

$$\forall P, \llbracket K \bowtie P \rrbracket_{\text{x86}} \sqsubseteq \llbracket P \rrbracket_{\text{mCertiKOS}}$$

- 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.*

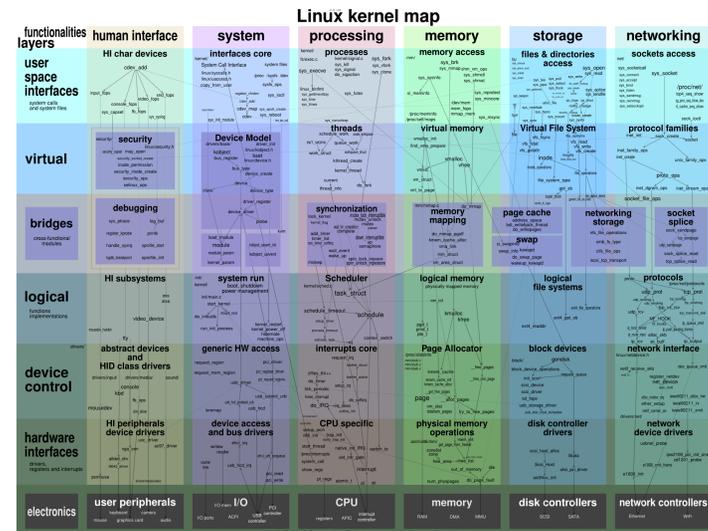
Size of TCB and Spec/Proof



- 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

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



Thank You

