

Just because you can doesn't mean you should...

TypeStates for Increased Driver Correctness

Tyler Potyondy, Anthony Tarbinian, Leon Schuermann, Eric Mugnier, Adin Ackermann, Amit Levy, Pat Pannuto



Hello!

- PhD Student at UC San Diego
- Involved with Tock since 2023 (~2.5 years)
- My research centers around making systems secure-by-default



(Last weekend in the Sierra's!)

Rust will fix our problems!

Buffer overflows

Use-After-Free

Data Races

Uninitialized Accesses

Rust will fix our problems!

~~Buffer overflows~~

~~Use After Free~~

Device Protocol Bugs

FFI Bugs

~~Data Races~~

~~Uninitialized Accesses~~



What is a device protocol violation?

When software issues commands to hardware that violate the hardware specification

Low-level Driver Development

Specification /
reference manual



life.augmented

RM0461

Reference manual

STM32WLEx advanced Arm®-based 32-bit MCUs
with sub-GHz radio solution

Introduction

This document is addressed to application developers. It provides complete information on how to use the STM32WLEx microcontrollers memory and peripherals.

STM32WLEx MCUs with integrated sub-GHz radio operating in the 150 - 960 MHz ISM band, belong to a family of microcontrollers with different memory sizes, packages and peripherals.

For ordering information, mechanical and electrical device characteristics, refer to the corresponding datasheets.

For information on the Arm® Cortex®-M4 core, refer to the corresponding Arm® Technical Reference Manuals available on <http://infocenter.arm.com>.

STM32WLEx microcontrollers include ST state-of-the-art patented technology.

Related documents

- STM32WLE5xx STM32WLE4xx datasheet (DS13105)

For information on the device errata with respect to the datasheet and reference manual, refer to the STM32WLE5xx STM32WLE4xx errata sheet (ES0506).

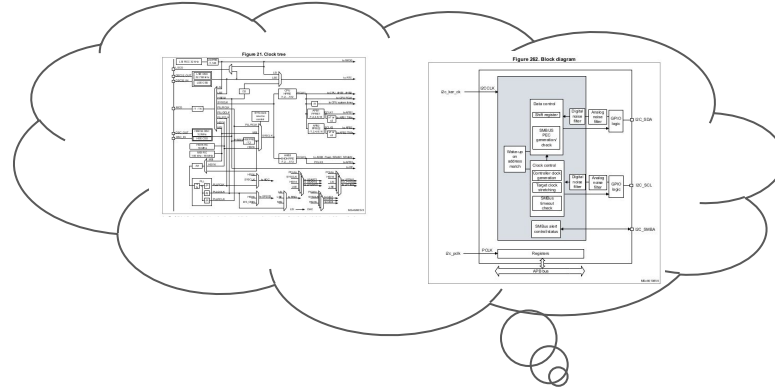
nRF52840

Product Specification

v1.0


Low-level Driver Development

Specification /
reference manual



Low-level Driver Development



Specification /
reference manual



```
27 /// An I2C master device.
28 ///
29 /// A 'TWI' instance wraps a 'registers::TWI' together with
30 /// additional data necessary to implement an asynchronous interface.
31 /// 4 implementations
32 pub struct TWI<'a> {
33     registers: StaticRef<TWIRegisters>,
34     client: OptionalCell<&a dyn hil::i2c::I2CMasterClient>,
35     slave_client: OptionalCell<&a dyn hil::i2c::I2CSlaveClient>,
36     buf: TakeCell<'static, [u8]>,
37     slave_read_buf: TakeCell<'static, [u8]>,
38 }
39
40 /// I2C bus speed.
41 #[repr(u32)]
42 pub enum Speed {
43     K100 = 0x01900000,
44     K250 = 0x04000000,
45     K400 = 0x06400000,
46 }
47
48 impl<'a> TWI<'a> {
49     const fn new(registers: StaticRef<TWIRegisters>) -> Self {
50         Self {
51             registers,
52             client: OptionalCell::empty(),
53             slave_client: OptionalCell::empty(),
54             buf: TakeCell::empty(),
55             slave_read_buf: TakeCell::empty(),
56         }
57     }
58
59     pub const fn new_tw10() -> Self {
60         TWI::new(registers::INSTANCES[0])
61     }
62 }
```


Low-level Driver Development

Specification /
reference manual




What can go wrong here? :)

```
27 // An I2C master device.
28 //
29 // A 'TWI' instance wraps a 'registers::TWI' together with
30 // additional data necessary to implement an asynchronous interface.
31 // 4 implementations
32 pub struct TWI<'a> {
33     registers: StaticRef<TWIRegisters>,
34     client: OptionalCell<&a dyn hil::i2c::i2c::MasterClient>,
35     slave_client: OptionalCell<&a dyn hil::i2c::i2c::SlaveClient>,
36     buf: TakeCell<'static, [u8]>,
37     slave_read_buf: TakeCell<'static, [u8]>,
38 }
39
40 // I2C bus speed.
41 #[repr(u32)]
42 pub enum Speed {
43     K100 = 0x01900000,
44     K250 = 0x04000000,
45     K400 = 0x06400000,
46 }
47
48 impl<'a> TWI<'a> {
49     const fn new(registers: StaticRef<TWIRegisters>) -> Self {
50         Self {
51             registers,
52             client: OptionalCell::empty(),
53             slave_client: OptionalCell::empty(),
54             buf: TakeCell::empty(),
55             slave_read_buf: TakeCell::empty(),
56         }
57     }
58
59     pub const fn new_twio() -> Self {
60         TWI::new(registers::INSTANCES[0])
61     }
62 }
```

Low-level Driver Development

Specification /
reference manual



What can go wrong here? :)

```
27 // An I2C master device.
28 //
29 // A 'TWI' instance wraps a 'registers::TWI' together with
30 // additional data necessary to implement an asynchronous interface.
31 // 4 implementations
32 pub struct TWI<'a> {
33     registers: StaticRef<TWIRegisters>,
34     client: OptionalCell<&a dyn hil::I2C::I2CMasterClient>,
35     slave_client: OptionalCell<&a dyn hil::I2C::I2CSlaveClient>,
36     buf: TakeCell<'static, [u8]>,
37     slave_read_buf: TakeCell<'static, [u8]>,
38 }
39
40 // I2C bus speed.
41 #[repr(u32)]
42 pub enum Speed {
43     K100 = 0x01900000,
44     K250 = 0x04000000,
45     K400 = 0x06400000,
46 }
47
48 impl<'a> TWI<'a> {
49     const fn new(registers: StaticRef<TWIRegisters>) -> Self {
50         Self {
51             registers,
52             client: OptionalCell::empty(),
53             slave_client: OptionalCell::empty(),
54             buf: TakeCell::empty(),
55             slave_read_buf: TakeCell::empty(),
56         }
57     }
58
59     pub const fn new_twio() -> Self {
60         TWI::new(registers::INSTANCES[0])
61     }
62 }
```

The diagram shows a neural network architecture with three main layers: Input, Hidden, and Output. The Input layer is connected to the Hidden layer, which is connected to the Output layer. The flow of information is indicated by arrows. The diagram is enclosed in a cloud-like shape.




```

27 // // An I2C interface device.
28
29 // // The "I2C" instance wraps a "registers" TML together with
30 // // "registers" and is necessary to implement an asynchronous interface.
31 // //
32 // // registers: StaticRef<TwoRegisters>,
33 // // client: Optional<Client> & dyn I2C::I2CMasterClient,
34 // // slave_client: Optional<Client> & dyn I2C::I2CSlaveClient,
35 // // buf: TakeRef<Static, [u8]>,
36 // // slave_read_buf: TakeRef<Static, [u8]>,
37
38 // // I2C bus speed.
39 #pragma clang optimize off
40 #pragma clang inlinehint off
41 pub struct I2C {
42     K18M = 0x01000000,
43     K250 = 0x00000000,
44     K400 = 0x00000000
45 }
46
47 impl as TML as {
48     const fn new(registers: StaticRef<TwoRegisters>) -> Self {
49         Self {
50             registers,
51             client: Optional<Client>::empty(),
52             slave_client: Optional<Client>::empty(),
53             buf: TakeRef<Static, [u8]>::empty(),
54             slave_read_buf: TakeRef<Static, [u8]>::empty(),
55         }
56     }
57
58     pub const fn new_twai() -> Self {
59         TML::new(registers).into(Self)
60     }

```

3

Why is this challenging?



Challenging!


```
27 // An I2C master device.
28 //
29 // A TWI instance wraps a registers::TWI together with
30 // additional data necessary to implement an asynchronous interface.
31 //
32 // Implementation.
33 pub struct TWI<as {
34     registers: StaticRef<TwoRegisters>,
35     client: OptionalCell<dyn hil::i2c::I2CMasterClient>,
36     slave_client: OptionalCell<dyn hil::i2c::I2CSlaveClient>,
37     buf: TakeCell<static, [u8]>,
38     slave_read_buf: TakeCell<static, [u8]>,
39 }
40
41 // I2C bus speed.
42 // [repr(u32)]
43 pub enum Speed {
44     K100 = 0x00000000,
45     K200 = 0x00000000,
46     K400 = 0x00000000,
47 }
48
49 impl<as TWI<as {
50     const fn new(registers: StaticRef<TwoRegisters>) -> Self {
51         Self {
52             registers,
53             client: OptionalCell::empty(),
54             slave_client: OptionalCell::empty(),
55             buf: TakeCell::empty(),
56             slave_read_buf: TakeCell::empty(),
57         }
58     }
59
60     pub const fn new_tw181 -> Self {
61         TWI::new(registers, INSTANCES[0])
62     }
63 }
```

Validity of a given MMIO operation depends on the current hardware state.



Modern hardware may transition the hardware state without input from the driver.

Why is this challenging?



Challenging!



```
27 // An I2C master device.
28 //
29 // A TWI instance wraps a registers::TWI together with
30 // additional data necessary to implement an asynchronous interface.
31 //
32 // Implementation.
33 pub struct TWI as {
34     registers: StaticRef<TwoRegisters>,
35     client: OptionalCell<dyn hil::i2c::I2CMasterClient>,
36     slave_client: OptionalCell<dyn hil::i2c::I2CSlaveClient>,
37     buf: TakeCell<static, [u8]>,
38     slave_read_buf: TakeCell<static, [u8]>,
39 }
40
41 // I2C bus speed.
42 #[repr(u32)]
43 pub enum Speed {
44     K100 = 0x00000000,
45     K200 = 0x00000000,
46     K400 = 0x00000000,
47 }
48
49 impl<'a> TWI as {
50     const fn new(registers: StaticRef<TwoRegisters>) -> Self {
51         Self {
52             registers,
53             client: OptionalCell::empty(),
54             slave_client: OptionalCell::empty(),
55             buf: TakeCell::empty(),
56             slave_read_buf: TakeCell::empty(),
57         }
58     }
59
60     pub const fn new_tw181() -> Self {
61         TWI::new(registers::INSTANCES[0])
62     }
63 }
```

Validity of a given MMIO operation depends on the current hardware state.



Modern hardware may transition the hardware state without input from the driver.



- ! may result in a buggy driver.
- ! at worst, may cause systematic failures (e.g. hanging the system's bus).



Challenging!

```
27 /// An I2C master device.
28 ///
29 /// A 'TWI' instance wraps a 'registers::TWI' together with
30 /// additional data necessary to implement an asynchronous interface.
31 pub struct TWI<'a> {
32     registers: StaticRef<TWIRegisters>,
33     client: OptionalCell<'a dyn hil::i2c::I2CMasterClient>,
34     slave_client: OptionalCell<'a dyn hil::i2c::I2CSlaveClient>,
35     buf: TakeCell<'static, [u8]>,
36     slave_read_buf: TakeCell<'static, [u8]>,
37 }
38
39 /// I2C bus speed.
40 #[repr(u32)]
41 pub enum Speed {
42     K100 = 0x01980000,
43     K250 = 0x04080000,
44     K400 = 0x06400000,
45 }
46
47 impl<'a> TWI<'a> {
48     const fn new(registers: StaticRef<TWIRegisters>) -> Self {
49         Self {
50             registers,
51             client: OptionalCell::empty(),
52             slave_client: OptionalCell::empty(),
53             buf: TakeCell::empty(),
54             slave_read_buf: TakeCell::empty(),
55         }
56     }
57
58     pub const fn new_twio() -> Self {
59         TWI::new(registers; INSTANCES[0])
60     }
61 }
```

Q: Can we enforce, at compile time, that the implemented driver will always comply with the developer's hw mental model?



Verified!

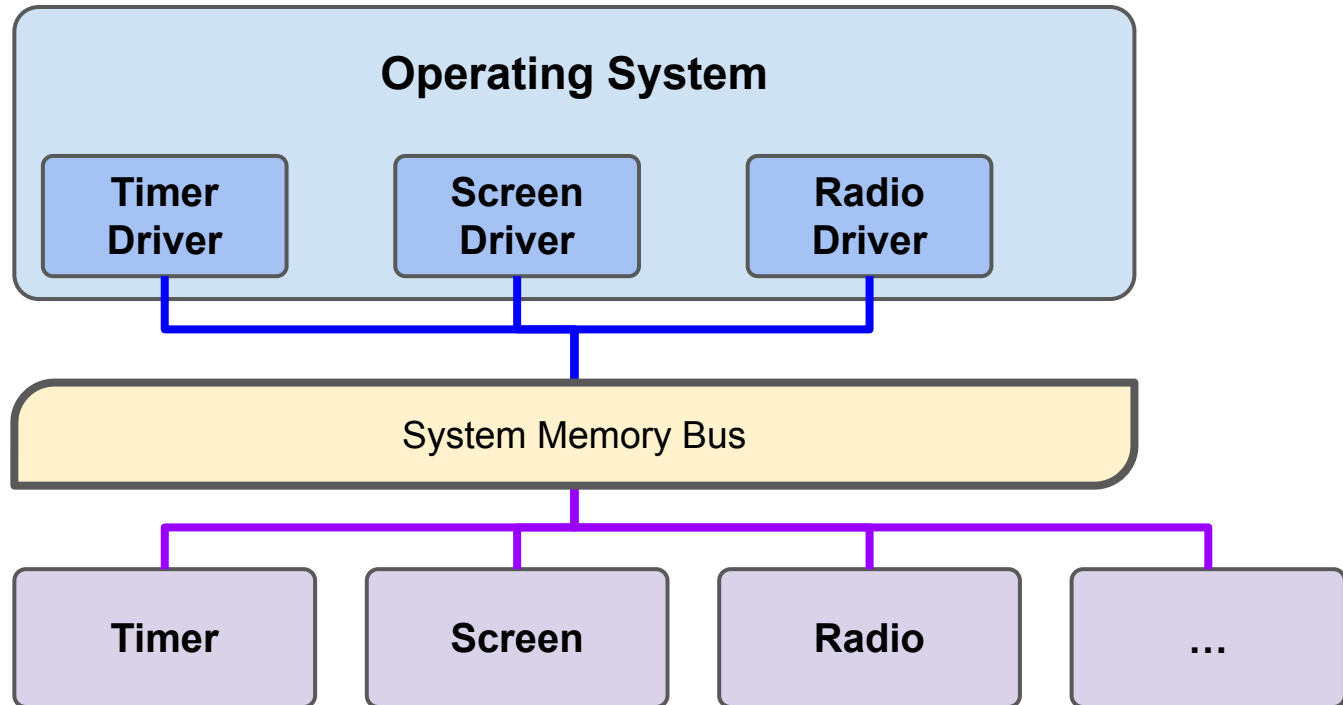
```
27 /// An I2C master device.
28
29 /// A 'TWI' instance wraps a 'registers::TWI' together with
30 /// additional data necessary to implement an asynchronous interface.
31 pub struct TWI<'a> {
32     registers: StaticRef<TwoRegisters>,
33     client: OptionalCell<'a dyn hil::I2C::I2CMasterClient>,
34     slave_client: OptionalCell<'a dyn hil::I2C::I2CSlaveClient>,
35     buf: TakeCell<'static, [u8]>,
36     slave_read_buf: TakeCell<'static, [u8]>,
37 }
38
39 /// I2C bus speed.
40 #[repr(u32)]
41 pub enum Speed {
42     K100 = 0x01980000,
43     K250 = 0x04000000,
44     K400 = 0x06400000,
45 }
46
47 impl<'a> TWI<'a> {
48     const fn new(registers: StaticRef<TwoRegisters>) -> Self {
49         Self {
50             registers,
51             client: OptionalCell::empty(),
52             slave_client: OptionalCell::empty(),
53             buf: TakeCell::empty(),
54             slave_read_buf: TakeCell::empty(),
55         }
56     }
57
58     pub const fn new_twio() -> Self {
59         TWI::new(registers; INSTANCES[0])
60     }
61 }
```

Q: Can we enforce, at compile time, that the implemented driver will always comply with the developer's hw mental model?



software driver adheres to hardware's specification

Key Insight: **Software** talks to **hardware** through a “narrow waist” — memory-mapped I/O




We present a framework that statically
(compile-time) prevents device protocol violations



- software driver adheres to hardware's specification
- i.e., only performs MMIO operations valid for the given hw state

Statically eliminate device protocol violations with minimal-to-no overheads in runtime and code size.

TypeStates



DSL

Outline

- Introducing device protocol violations
- **How do we build drivers today?**
- TypeState programming
- Our System
- Evaluation & Closing Thoughts

Let's build a UART driver...

DATA

WriteOnly

7 6 5 4 3 2 1 0



Write a byte to transmit. Bytes are placed in an internal FIFO queue. The UART transmits whenever queue is non-empty and pops entries once sent.

STATUS

ReadOnly

7 6 5 4 3 2 1 0



Read hardware status. BUSY indicates when a transmission is active. FULL indicates when the FIFO transmit queue is full; **DATA** must not be written when FULL is asserted.

(Hypothetical UART Hardware Specification)

Let's build a UART driver...

DATA

WriteOnly

7 6 5 4 3 2 1 0



Write a byte to transmit. Bytes are placed in an internal FIFO queue. The UART transmits whenever queue is non-empty and pops entries once sent.

STATUS

ReadOnly

7 6 5 4 3 2 1 0



Read hardware status. BUSY indicates when a transmission is active. FULL indicates when the FIFO transmit queue is full; **DATA** must not be written when FULL is asserted.

(Hypothetical UART Hardware Specification)

Let's build a UART driver...

DATA

WriteOnly

7 6 5 4 3 2 1 0

Byte

STATUS

ReadOnly

7 6 5 4 3 2 1 0

reserved

FULL
BUSY

Write a byte to transmit. Bytes are placed in an internal FIFO queue. The UART transmits whenever queue is non-empty and pops entries once sent.

Read hardware status. BUSY indicates when a transmission is active. FULL indicates when the FIFO transmit queue is full; **DATA** must not be written when FULL is asserted.

QueueNotFull

QueueFull

(Hypothetical UART Hardware Specification)

(Developer Mental Model of HW Specification)

Let's build a UART driver...

DATA

WriteOnly

7 6 5 4 3 2 1 0

Byte

STATUS

ReadOnly

7 6 5 4 3 2 1 0


reserved

FULL
BUSY

Write a byte to transmit. Bytes are placed in an internal FIFO queue. The UART transmits whenever queue is non-empty and pops entries once sent.

Read hardware status. BUSY indicates when a transmission is active. FULL indicates when the FIFO transmit queue is full; **DATA** must not be written when FULL is asserted.

(Hypothetical UART Hardware Specification)



(Developer Mental Model of HW Specification)

Let's build a UART driver...

DATA

WriteOnly

7 6 5 4 3 2 1 0

Byte

STATUS

ReadOnly

7 6 5 4 3 2 1 0

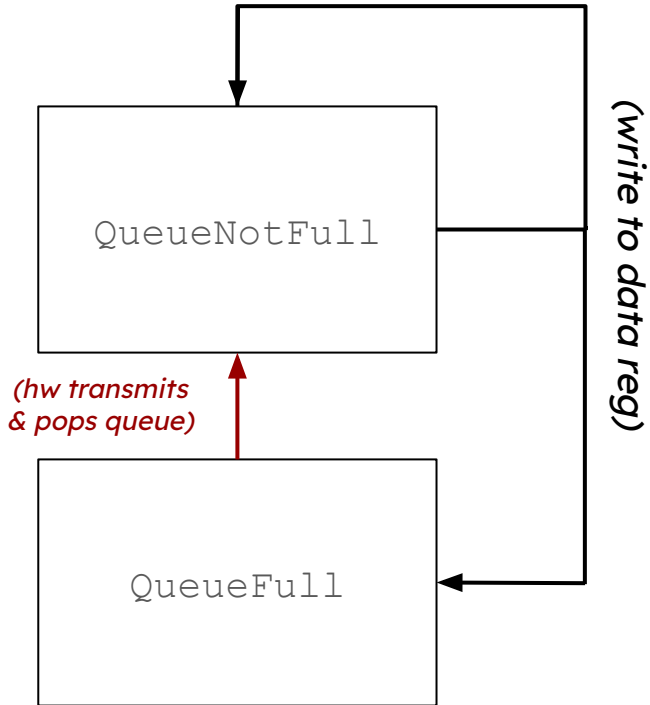
reserved

FULL
BUSY

Write a byte to transmit. Bytes are placed in an internal FIFO queue. The UART transmits whenever queue is non-empty and pops entries once sent.

Read hardware status. BUSY indicates when a transmission is active. FULL indicates when the FIFO transmit queue is full; **DATA** must not be written when FULL is asserted.

(Hypothetical UART Hardware Specification)



(Developer Mental Model of HW Specification)

Let's build a UART driver...

DATA

WriteOnly

7 6 5 4 3 2 1 0

Byte

STATUS

ReadOnly

7 6 5 4 3 2

reserved

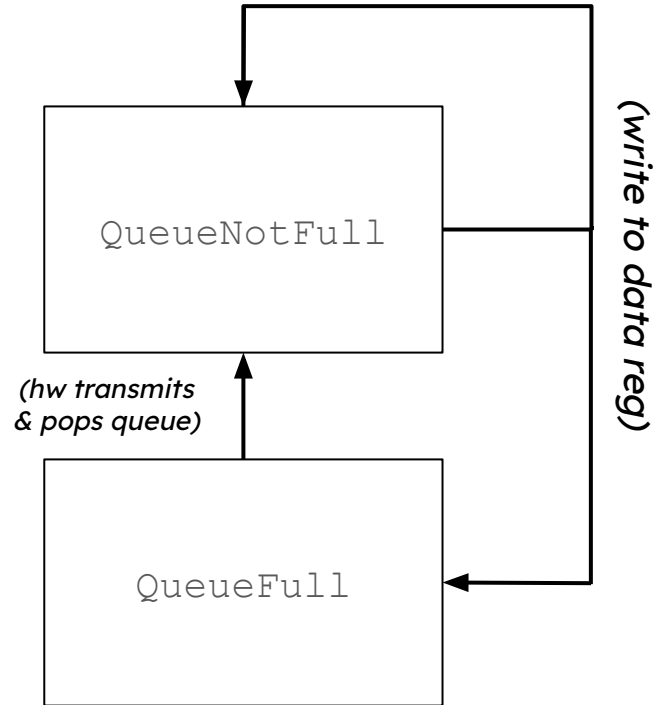


FULL
BUSY

Write a byte to transmit. Bytes are placed in an internal FIFO queue. The UART transmits whenever queue is non-empty and pops entries once sent.

Read hardware status. BUSY indicates when a transmission is active. FULL indicates when the FIFO transmit queue is full; **DATA** must not be written when FULL is asserted.

(Hypothetical UART Hardware Specification)



(Developer Mental Model of HW Specification)

Let's build a UART driver...

DATA

WriteOnly

7 6 5 4 3 2 1 0

Byte

STATUS

ReadOnly

7 6 5 4 3 2

reserved

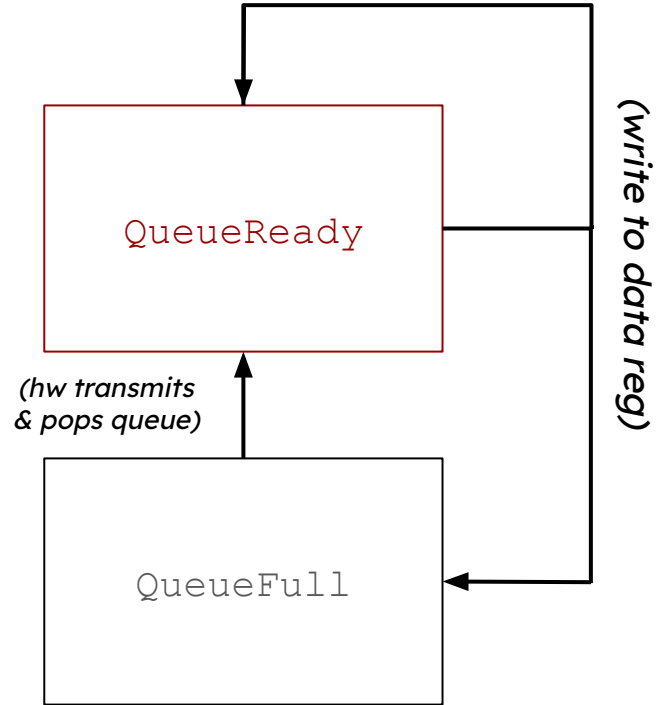


FULL
BUSY

Write a byte to transmit. Bytes are placed in an internal FIFO queue. The UART transmits whenever queue is non-empty and pops entries once sent.



Read hardware status. BUSY indicates when a transmission is active. FULL indicates when the FIFO transmit queue is full; **DATA** must not be written when FULL is asserted.

(Hypothetical UART Hardware Specification)



(Developer Mental Model of HW Specification)

Let's build a UART driver...





```
1 struct UartRegisters {
2     data: RegisterWO<u8>,
3     status: RegisterRO<StatusReg>
4 } // ^^^^^^^^^^^ helper for bitfields
5
6 fn transmit(reg: &UartRegisters, buf: &[u8]) {
7     for index in len(buf):
8         reg.data.write(buf[index])
9         // busy wait until queue has space
10        while (reg.status.read().is_set(StatusReg::FULL)) {}
11 }
```

(Implemented UART driver - based on our mental model)

(Developer Mental Model of HW Specification)

Let's build a UART driver...



```
1 struct UartRegisters {
2     data: RegisterWO<u8>,
3     status: RegisterRO<StatusReg>
4 } // ^^^^^^^^^^^ helper for bitfields
5
6 fn transmit(reg: &UartRegisters, buf: &[u8]) {
7     for index in len(buf):
8         reg.data.write(buf[index])
9         // busy wait until queue has space
10        while (reg.status.read().is_set(StatusReg::FULL)) {}
11 }
```

(Implemented UART driver - based on our mental model)

Do you see the bug?

(Developer Mental Model of HW Specification)

Let's build a UART driver...

```
1 struct UartRegisters {
2     data: RegisterWO<u8>,
3     status: RegisterRO<StatusReg>
4 } // ^^^^^^^^^^^ helper for bitfields
5
6 fn transmit(reg: &UartRegisters, buf: &[u8]) {
7     for index in len(buf):
8         reg.data.write(buf[index])
9         // busy wait until queue has space
10        while (reg.status.read().is_set(StatusReg::FULL) {}
11 }
```

(Implemented UART driver – based on our mental model)

Recall...



DATA must not be written
when FULL is asserted.

Do you see the bug?

Let's build a UART driver...

Violate device protocol!

We assume that the hw transmit queue is NOT full when calling this function



```
1 struct UartRegisters {
2     data: RegisterWO<u8>,
3     status: RegisterRO<StatusReg>
4 } // ^^^^^^^^^^^ helper for bitfields
5
6 fn transmit(reg: &UartRegisters, buf: &[u8]) {
7     for index in len(buf):
8         reg.data.write(buf[index])
9         // busy wait until queue has space
10        while (reg.status.read().is_set(StatusReg::FULL) {}
11 }
```

(Implemented UART driver – based on our mental model)

Recall...



DATA must not be written
when FULL is asserted.

Do you see the bug?

How might we prevent this bug?

Standard approaches for enforcing system properties (generally)...

Testing

(only proves the absence of tested bugs).

Formal Verification

(challenging; requires domain specific expertise).

How might we prevent this bug?

Standard approaches for enforcing system properties (generally)...

Testing

(only proves the absence of tested bugs).

TypeState Programming

Formal Verification

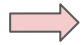
(challenging; requires domain specific expertise).

Outline

- Introducing device protocol violations
- How do we build drivers today?
- **TypeState programming**
- Our System
- Evaluation & Closing Thoughts

A TypeStated Queue

- Encode system properties into the type-system.




```
1 struct Full {} // 3 items in queue
2 struct Two {} // 2 items in queue
3 struct One {} // 1 item in queue
4 struct Empty {} // 0 items in queue
5
6 struct Queue<S: State> {
7     queue: [u8; 3]
8 }
```

(Using typestates to *statically enforce a correct implementation for a queue of size 3*)

A TypeStated Queue

- Encode system properties into the type-system.
- Define valid operations as functions on respective type.

```
1 struct Full {} // 3 items in queue
2 struct Two {} // 2 items in queue
3 struct One {} // 1 item in queue
4 struct Empty {} // 0 items in queue
5
6 struct Queue<S: State> {
7     queue: [u8; 3]
8 }
9
10 impl Queue<Empty> {
11     fn push(self) -> Queue<One>
12 }
13
14 impl Queue<One> {
15     fn push(self) -> Queue<Two>
16     fn pop(self) -> Queue<Empty>
17 }
18
19 // similar form to Queue<One>
20 impl Queue<Two> { ... }
21
22 impl Queue<Full> {
23     fn pop(self) -> Queue<Two>
24 }
```




(Using typestates to *statically enforce a correct implementation for a queue of size 3*)

A TypeStated Queue

- Encode system properties into the type-system.
- Define valid operations as functions on respective type.
- Incorrect usages result in a compilation error!

```
1 struct Full {} // 3 items in queue
2 struct Two {} // 2 items in queue
3 struct One {} // 1 item in queue
4 struct Empty {} // 0 items in queue
5
6 struct Queue<S: State> {
7     queue: [u8; 3]
8 }
9
10 impl Queue<Empty> {
11     fn push(self) -> Queue<One>
12 }
13
14 impl Queue<One> {
15     fn push(self) -> Queue<Two>
16     fn pop(self) -> Queue<Empty>
17 }
18
19 // similar form to Queue<One>
20 impl Queue<Two> { ... }
21
22 impl Queue<Full> {
23     fn pop(self) -> Queue<Two>
24 }
```



(Using typestates to *statically enforce a correct implementation for a queue of size 3*)

A TypeStated Queue

Recall from hw spec...

! The UART transmits whenever queue is non-empty and pops entries once sent.

Out-of-the-box typestates cannot model this state transition!

(Using typestates to *statically enforce a correct implementation for a queue of size 3*)

```
1 struct Full {} // 3 items in queue
2 struct Two {} // 2 items in queue
3 struct One {} // 1 item in queue
4 struct Empty {} // 0 items in queue
5
6 struct Queue<S: State> {
7     queue: [u8; 3]
8 }
9
10 impl Queue<Empty> {
11     fn push(self) -> Queue<One>
12 }
13
14 impl Queue<One> {
15     fn push(self) -> Queue<Two>
16     fn pop(self) -> Queue<Empty>
17 }
18
19 // similar form to Queue<One>
20 impl Queue<Two> { ... }
21
22 impl Queue<Full> {
23     fn pop(self) -> Queue<Two>
24 }
```


Outline

- Introducing device protocol violations
- How do we build drivers today?
- TypeState programming
- **Our System**
- Evaluation & Closing Thoughts

We present a framework that statically (at compile time) prevents device protocol violations

- Achieve device protocol enforcement with minimal to no overheads in runtime and code size.

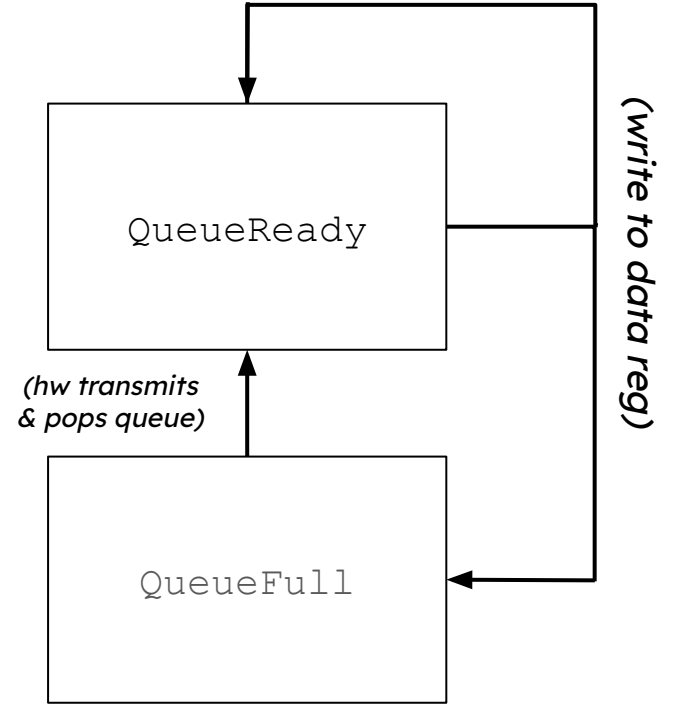
TypeStates



DSL

- **Primary contribution:** Introduce a refinement to type-states and principled approach to model hardware-software concurrency using type-states.

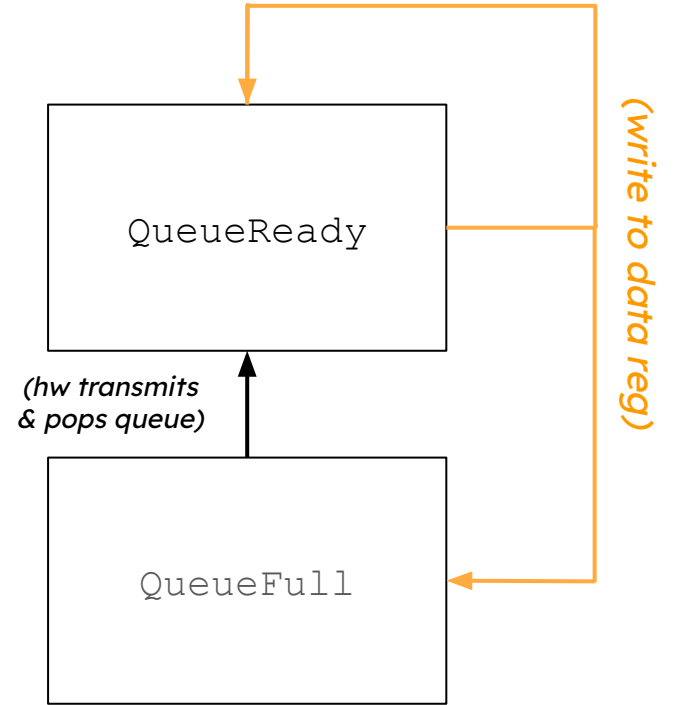
We observe... there are two classes of hw state transitions



(Developer Mental Model of HW Specification)

We observe... there are two classes of hw state transitions

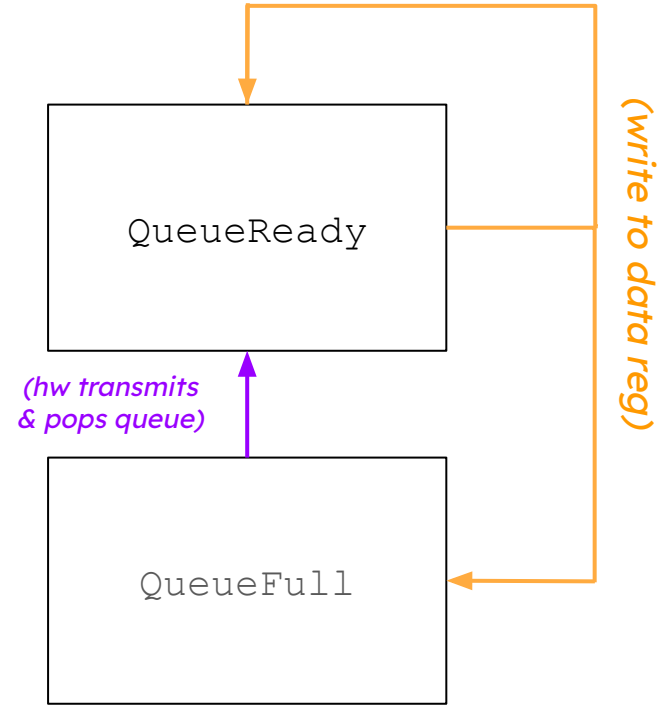
- Software-initiated



(Developer Mental Model of HW Specification)

We observe... there are two classes of hw state transitions

- Software-initiated
- Hardware-initiated



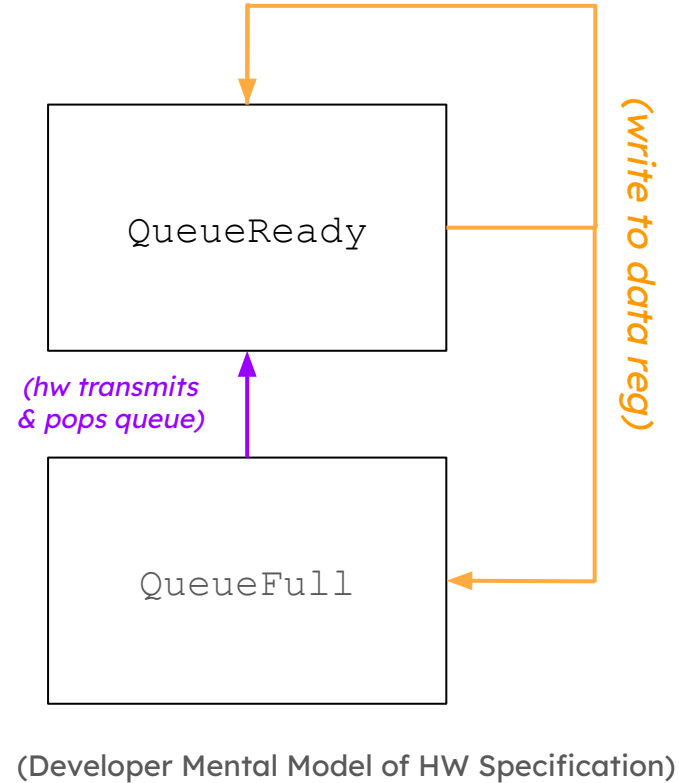
(Developer Mental Model of HW Specification)

We observe... there are two classes of hw state transitions

- Software-initiated
- Hardware-initiated

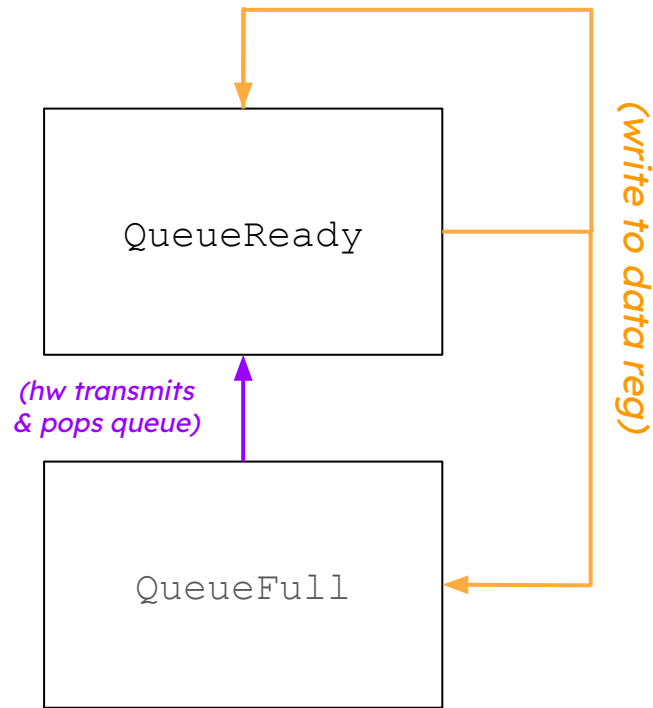
Categorize hardware states into two mutually exclusive families

- transient state
- stable state



Stable State

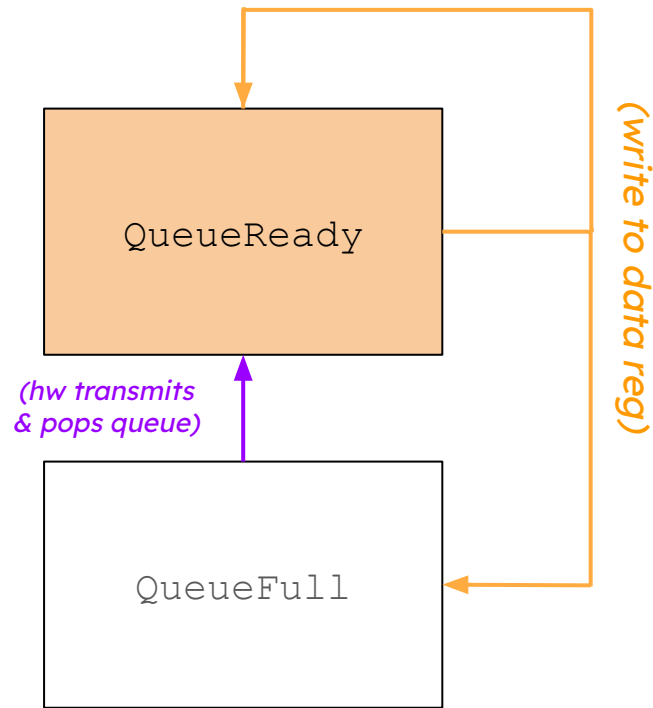
- Hw state that can only be exited with a software-initiated state transition.



(Developer Mental Model of HW Specification)

Stable State

- Hw state that can only be exited with a software-initiated state transition.



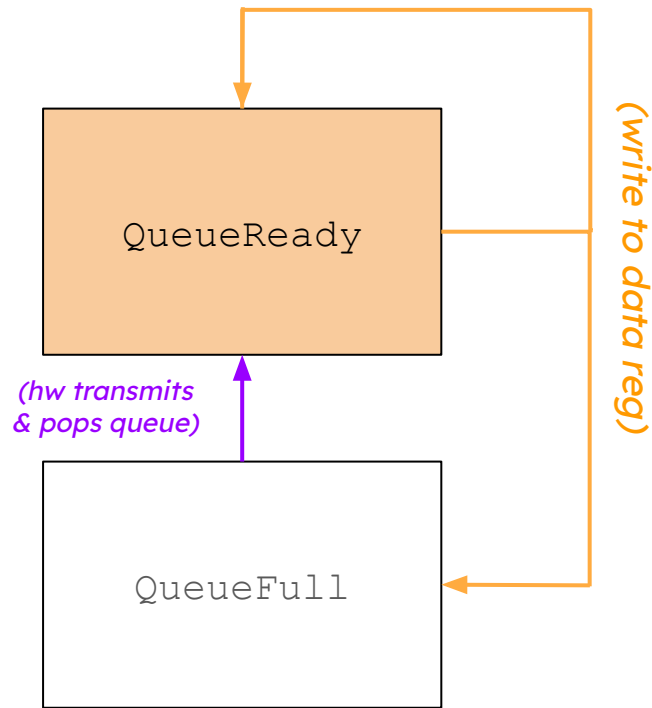
(Developer Mental Model of HW Specification)

Stable State

- Hw state that can only be exited with a software-initiated state transition.

Transient State

- Hw state with at least one hw-initiated state transition.
- Transition from transient state without explicit software involvement.




(Developer Mental Model of HW Specification)

Stable State

- Hw state that can only be exited with a software-initiated state transition.

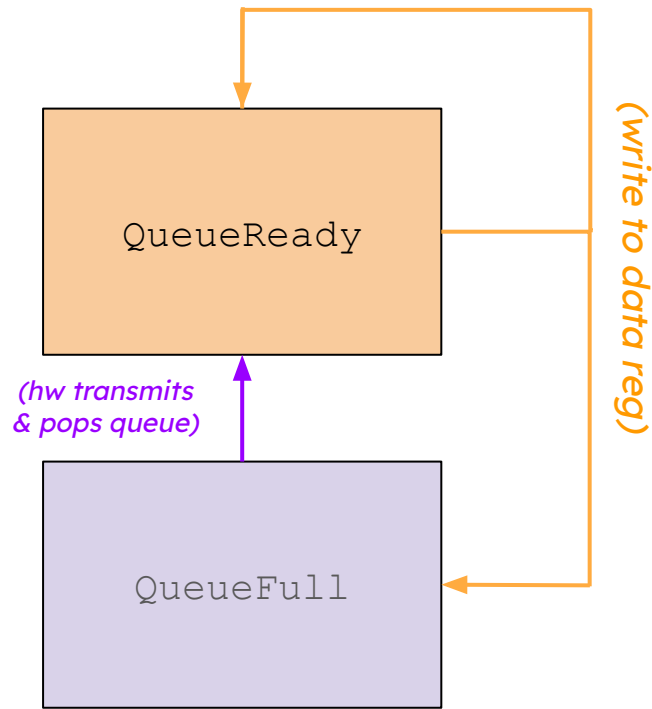
Transient State

- Hw state with at least one hw-initiated state transition.
- Transition from transient state without explicit software involvement.



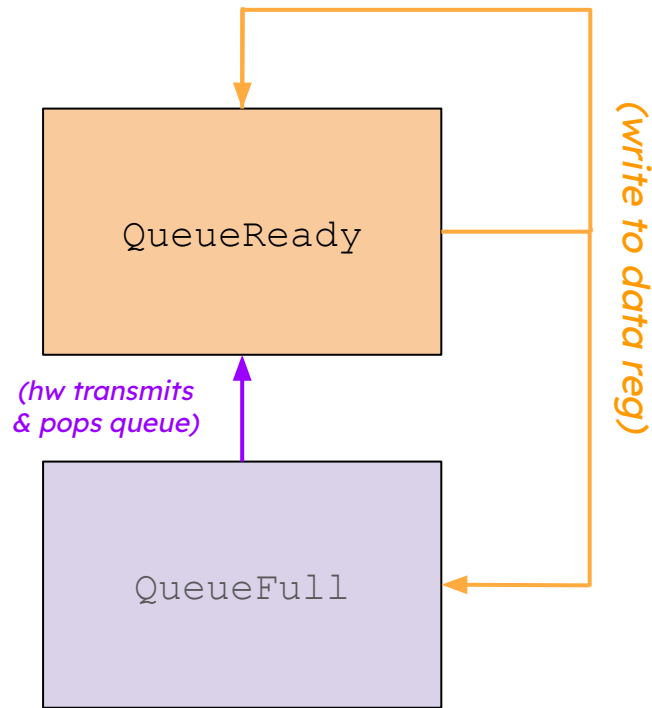
(Developer Mental Model of HW Specification)

- Stable states can be modeled with out-of-the-box typestates.



(Developer Mental Model of HW Specification)


- Stable states can be modeled with out-of-the-box typestates.
- Transient states cause typestates to no longer accurately model hw (violate static invariance).



(Developer Mental Model of HW Specification)

- Stable states can be modeled with out-of-the-box typestates.
- Transient states cause typestates to no longer accurately model hw (violate static invariance).

Typestates + restrict transient state operations & re-synchronization mechanism




(Developer Mental Model of HW Specification)

- Stable states can be modeled with out-of-the-box typestates.
- Transient states cause typestates to no longer accurately model hw (violate static invariance).

Typestates + restrict transient state operations & re-synchronization mechanism

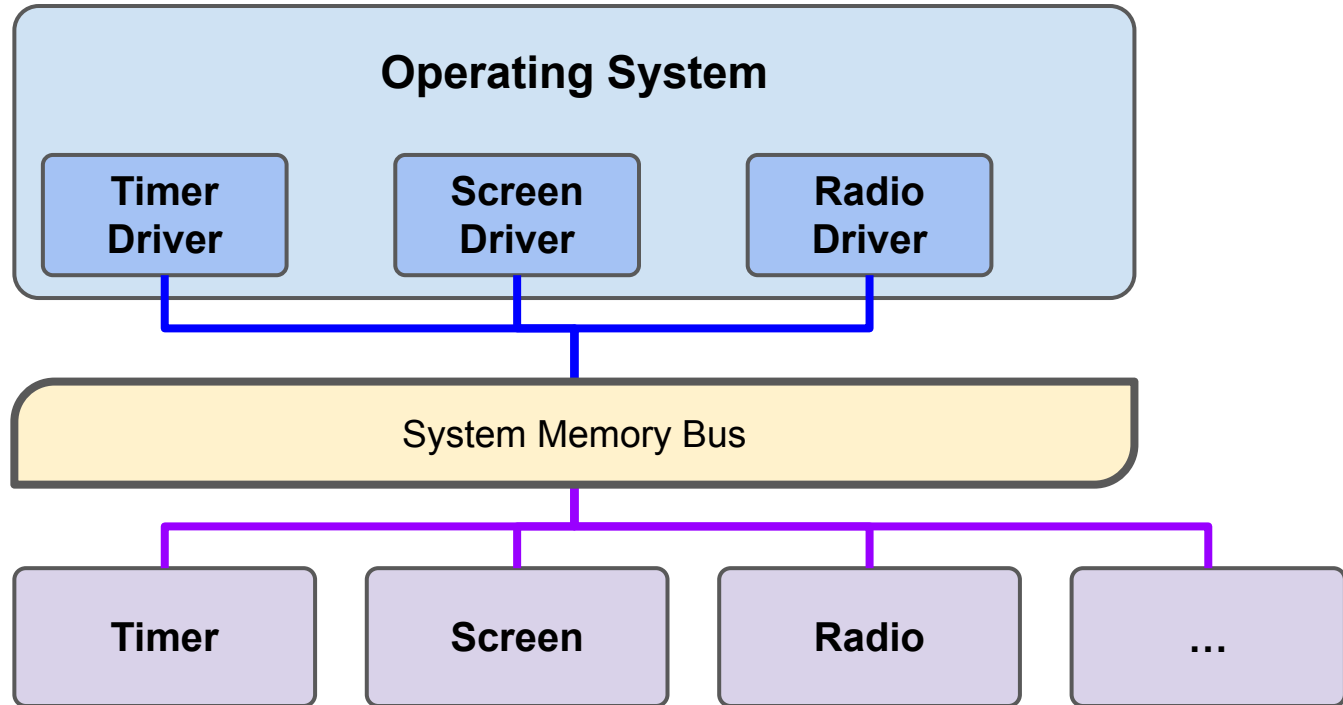


Careful: Transient states have potential for **TOCTOU bugs**!



(Developer Mental Model of HW Specification)

(Recall) Key Insight: **Software** talks to **hardware** through a “narrow waist” — memory-mapped I/O



```

1  +#[(states=[ QueueReady<Idle>,
2  +             QueueReady<Busy>(*T*),
3  +             QueueMaybeFull(*T*)      ])]
4  struct UartRegisters {
5  +  #[attribute(SC(QueueReady<Any>, QueueMaybeFull))]
6      data: WriteOnly<u8, Data::Register>,
7      // No attributes are required for `Status`
8      status: ReadOnly<u8, Status::Register>,
9  +  #[attribute(SC(Any, QueueReady<Idle>))]
10     flush: WriteOnly<u8, Flush::Register>,
11 +  #[attribute(QueueReady<Idle>)]
12     config: ReadWrite<u8m Config::Register>,
13 }

```

Annotations for updated UART driver

Enforce device protocols by constraining MMIO using type-states.

1. Label states and mark transient states

```
1  +#[states=[ QueueReady<Idle>,  
2  +           QueueReady<Busy>(*T*),  
3  +           QueueMaybeFull(*T*)  ])]  
4  struct UartRegisters {  
5  +  #[attribute(SC(QueueReady<Any>, QueueMaybeFull))]  
6    data: WriteOnly<u8, Data::Register>,  
7    // No attributes are required for `Status`  
8    status: ReadOnly<u8, Status::Register>,  
9  +  #[attribute(SC(Any, QueueReady<Idle>))]  
10   flush: WriteOnly<u8, Flush::Register>,  
11  +  #[attribute(QueueReady<Idle>)]  
12   config: ReadWrite<u8m Config::Register>,  
13 }
```

Annotations for updated UART driver.

Enforce device protocols by constraining MMIO using type-states.

1. Label states and mark transient states
2. Add constraints to registers

```
1  +#[states=[ QueueReady<Idle>,  
2  +           QueueReady<Busy>(*T*),  
3  +           QueueMaybeFull(*T*)      ])]  
4  struct UartRegisters {  
5  +  #[attribute(SC(QueueReady<Any>, QueueMaybeFull))]  
6    data: WriteOnly<u8, Data::Register>,  
7    // No attributes are required for `Status`  
8    status: ReadOnly<u8, Status::Register>,  
9  +  #[attribute(SC(Any, QueueReady<Idle>))]  
10   flush: WriteOnly<u8, Flush::Register>,  
11 +  #[attribute(QueueReady<Idle>)]  
12   config: ReadWrite<u8m Config::Register>,  
13 }
```

Annotations for updated UART driver.

Enforce device protocols by constraining MMIO using type-states.

1. Label states and mark transient states
2. Add constraints to registers

```
1  +#[states=[ QueueReady<Idle>,  
2  +           QueueReady<Busy>(*T*),  
3  +           QueueMaybeFull(*T*)      ])]  
4  struct UartRegisters {  
5  +  #[attribute(SC(QueueReady<Any>, QueueMaybeFull))]  
6    data: WriteOnly<u8, Data::Register>,  
7    // No attributes are required for `Status`  
8    status: ReadOnly<u8, Status::Register>,  
9  +  #[attribute(SC(Any, QueueReady<Idle>))]  
10   flush: WriteOnly<u8, Flush::Register>,  
11  +  #[attribute(QueueReady<Idle>)]  
12   config: ReadWrite<u8m Config::Register>,  
13 }
```

Annotations for updated UART driver.

Enforce device protocols by constraining MMIO using type-states.

DSL / proc-macro autogenerates type-states

```
1 // Hardware object made generic over device state.
2 struct UartRegisters<S: State> {
3     data: SCRegisterWO<S, u8>,
4     status: ReadOnly<u8, Status>,
5     reset: SCRegisterWO<S, u8>,
6     config: RWRegister<N, S, u8>,
7 }
8
9 // Wrapper around state changing registers.
10 struct SCRegisterWO<S: State, T> {
11     reg: WriteOnly<T>,
12     associated_state: PhantomData<S>,
13 }
14
15 // Wrapper around constrained MMIO.
16 struct RWRegister<N, S: State, T> {
17     reg: ReadWrite<T>,
18     associated_name: PhantomData<N>,
19     associated_state: PhantomData<S>,
20 }
21
22 // This impl is only generated for S==QueueReady<Idle>,
23 // which enforces config's device protocol invariants.
24 impl <T> RWRegister<Config, QueueReady<Idle>, T> {
25     fn read(&self) -> T {..}
26     fn write(&self, T) {..}
27 }
```

DSL / proc-macro autogenerates type-states

1. Modified MMIO register struct

```
1 // Hardware object made generic over device state.
2 struct UartRegisters<S: State> {
3     data: SCRegisterWO<S, u8>,
4     status: ReadOnly<u8, Status>,
5     reset: SCRegisterWO<S, u8>,
6     config: RWRegister<N, S, u8>,
7 }
8
9 // Wrapper around state changing registers.
10 struct SCRegisterWO<S: State, T> {
11     reg: WriteOnly<T>,
12     associated_state: PhantomData<S>,
13 }
14
15 // Wrapper around constrained MMIO.
16 struct RWRegister<N, S: State, T> {
17     reg: ReadWrite<T>,
18     associated_name: PhantomData<N>,
19     associated_state: PhantomData<S>,
20 }
21
22 // This impl is only generated for S==QueueReady<Idle>,
23 // which enforces config's device protocol invariants.
24 impl <T> RWRegister<Config, QueueReady<Idle>, T> {
25     fn read(&self) -> T {..}
26     fn write(&self, T) {..}
27 }
```

DSL / proc-macro autogenerates type-states

1. Modified MMIO register struct

```
1 // Hardware object made generic over device state.
2 struct UartRegisters<S: State> {
3     data: SCRegisterWO<S, u8>,
4     status: ReadOnly<u8, Status>,
5     reset: SCRegisterWO<S, u8>,
6     config: RWRegister<N, S, u8>,
7 }
8
9 // Wrapper around state changing registers.
10 struct SCRegisterWO<S: State, T> {
11     reg: WriteOnly<T>,
12     associated_state: PhantomData<S>,
13 }
14
15 // Wrapper around constrained MMIO.
16 struct RWRegister<N, S: State, T> {
17     reg: ReadWrite<T>,
18     associated_name: PhantomData<N>,
19     associated_state: PhantomData<S>,
20 }
21
22 // This impl is only generated for S==QueueReady<Idle>,
23 // which enforces config's device protocol invariants.
24 impl <T> RWRegister<Config, QueueReady<Idle>, T> {
25     fn read(&self) -> T {...}
26     fn write(&self, T) {...}
27 }
```

DSL / proc-macro autogenerates type-states

1. Modified MMIO register `struct`
2. Wrap tock registers in type-state

```
1 // Hardware object made generic over device state.
2 struct UartRegisters<S: State> {
3     data: SCRegisterWO<S, u8>,
4     status: ReadOnly<u8, Status>,
5     reset: SCRegisterWO<S, u8>,
6     config: RWRegister<N, S, u8>,
7 }
8
9 // Wrapper around state changing registers.
10 struct SCRegisterWO<S: State, T> {
11     reg: WriteOnly<T>,
12     associated_state: PhantomData<S>,
13 }
14
15 // Wrapper around constrained MMIO.
16 struct RWRegister<N, S: State, T> {
17     reg: ReadWrite<T>,
18     associated_name: PhantomData<N>,
19     associated_state: PhantomData<S>,
20 }
21
22 // This impl is only generated for S==QueueReady<Idle>,
23 // which enforces config's device protocol invariants.
24 impl <T> RWRegister<Config, QueueReady<Idle>, T> {
25     fn read(&self) -> T {...}
26     fn write(&self, T) {...}
27 }
```

DSL / proc-macro autogenerates type-states

1. Modified MMIO register `struct`
2. Wrap tock registers in type-state
3. Only define valid transitions

```
1 // Hardware object made generic over device state.
2 struct UartRegisters<S: State> {
3     data: SCRegisterWO<S, u8>,
4     status: ReadOnly<u8, Status>,
5     reset: SCRegisterWO<S, u8>,
6     config: RWRegister<N, S, u8>,
7 }
8
9 // Wrapper around state changing registers.
10 struct SCRegisterWO<S: State, T> {
11     reg: WriteOnly<T>,
12     associated_state: PhantomData<S>,
13 }
14
15 // Wrapper around constrained MMIO.
16 struct RWRegister<N, S: State, T> {
17     reg: ReadWrite<T>,
18     associated_name: PhantomData<N>,
19     associated_state: PhantomData<S>,
20 }
21
22 // This impl is only generated for S==QueueReady<Idle>,
23 // which enforces config's device protocol invariants.
24 impl <T> RWRegister<Config, QueueReady<Idle>, T> {
25     fn read(&self) -> T {..}
26     fn write(&self, T) {..}
27 }
```


Updated transmit

```
1 // Driver object holds hardware reference & driver-specific state
2 struct UartDriver {
3 -   registers: &UartRegisters,
4 +   registers: MMIOCell<UartStates>,
5 }
6 impl UartDriver {
7     pub fn transmit(&self, buf: &[u8]) {
8         for data in buf.iter() {
9 -             self.registers.data.write(data);
10 -             while self.registers.status.is_set(Status::FULL) {};
11 +             self.registers.map(|state| {
12 +                 match state {
13 +                     UartStates::QueueReadyIdle(regs) => {
14 +                         regs.data.write(data).sync_state()
15 +                     }
16 +                     UartStates::QueueReadyBusy(regs) => {
17 +                         regs.data.write(data).sync_state()
18 +                     }
19 +                     UartStates::QueueMaybeFull(regs) => {
20 +                         regs.sync_state() /* no regs.data.write() exists */
21 +                     }
22 }}}}
```

No longer possible for the transmit method to attempt to write DATA when the queue is possibly full!

Updated transmit

```
1 // Driver object holds hardware reference & driver-specific state
2 struct UartDriver {
3     - registers: &UartRegisters,
4     + registers: MMIOCell<UartStates>,
5 }
6 impl UartDriver {
7     pub fn transmit(&self, buf: &[u8]) {
8         for data in buf.iter() {
9             - self.registers.data.write(data);
10            - while self.registers.status.is_set(Status::FULL) {};
11            + self.registers.map(|state| {
12                + match state {
13                +     UartStates::QueueReadyIdle(regs) => {
14                +         regs.data.write(data).sync_state()
15                +     }
16                +     UartStates::QueueReadyBusy(regs) => {
17                +         regs.data.write(data).sync_state()
18                +     }
19                +     UartStates::QueueMaybeFull(regs) => {
20                +         regs.sync_state() /* no regs.data.write() exists */
21                +     }
22            }
23        }
24    }
25 }
```

No longer possible for the transmit method to attempt to write DATA when the queue is possibly full!

Updated transmit

```
1 // Driver object holds hardware reference & driver-specific state
2 struct UartDriver {
3 -   registers: &UartRegisters,
4 +   registers: MMIOCell<UartStates>,
5 }
6 impl UartDriver {
7     pub fn transmit(&self, buf: &[u8]) {
8         for data in buf.iter() {
9 -         self.registers.data.write(data);
10 -         while self.registers.status.is_set(Status::FULL) {};
11 +         self.registers.map(|state| {
12 +             match state {
13 +                 UartStates::QueueReadyIdle(regs) => {
14 +                     regs.data.write(data).sync_state()
15 +                 }
16 +                 UartStates::QueueReadyBusy(regs) => {
17 +                     regs.data.write(data).sync_state()
18 +                 }
19 +                 UartStates::QueueMaybeFull(regs) => {
20 +                     regs.sync_state() /* no regs.data.write() exists */
21 +                 }
22             }
23         });
24     }
25 }
```

No longer possible for the transmit method to attempt to write DATA when the queue is possibly full!

Updated transmit



```
1 // Driver object holds hardware reference & driver-specific state
2 struct UartDriver {
3 -   registers: &UartRegisters,
4 +   registers: MMIOCell<UartStates>,
5 }
6 impl UartDriver {
7     pub fn transmit(&self, buf: &[u8]) {
8         for data in buf.iter() {
9 -             self.registers.data.write(data);
10 -             while self.registers.status.is_set(Status::FULL) {};
11 +             self.registers.map(|state| {
12 +                 match state {
13 +                     UartStates::QueueReadyIdle(regs) => {
14 +                         regs.data.write(data).sync_state()
15 +                     }
16 +                     UartStates::QueueReadyBusy(regs) => {
17 +                         regs.data.write(data).sync_state()
18 +                     }
19 +                     UartStates::QueueMaybeFull(regs) => {
20 +                         regs.sync_state() /* no regs.data.write() exists */
21 +                     }
22 +                 }
23             });
24     }
25 }
```

No longer possible for the transmit method to attempt to write DATA when the queue is possibly full!

Outline

- Introducing device protocol violations
- How do we build drivers today?
- TypeState programming
- Our System
- **Evaluation & Closing Thoughts**

Implementation with TockOS



What's the catch...


- Code size?
- Developer effort?
- Runtime performance?

Driver	Platform	Binary Size (B)	Diff (B)	Percent Diff
Baseline	Nrf52840	218594	–	–
UART	Nrf52840	218594	+0	0.00%
Temperature Sensor	Nrf52840	218594	+0	0.00%
IEEE 802.15.4 Radio	Nrf52840	218602	+8	0.00%
Baseline	STM	107482	–	–
TRNG	STM	107490	+8	0.00%
UART	STM	107490	+8	0.00%

Code size of total kernel binary image for a baseline kernel image and kernel integrating our system into drivers.

Our system adds no code size overhead!

Our system adds negligible runtime overheads.



MacroBenchmark Performance (in CPU cycles).

Driver	States	Original LoC	Annotations	Integration
nRF52 UARTE	5	526	43 (+)	492 (+) 110 (-)
nRF5x Temperature	2	151	4 (+)	53 (+) 17 (-)
nRF52 15.4 Radio	8	1352	33 (+)	518 (+) 157 (-)
STM USART	5	743	45 (+)	351 (+) 79 (-)
STM TRNG	2	159	13 (+)	69 (+) 25 (-)
xHCI PortSC	5	6748	14 (+)	330 (+) 194 (-)

Our system adds some developer overheads

(improving the usability is ongoing!)

Case Study – NRF52 IEEE802.15.4 Driver



OPENTHREAD
released by Google

July 2023

Case Study – NRF52 IEEE802.15.4 Driver



July 2023

*Add SW ACKs to
radio driver*

August 2023

3 possible HW “shortcuts” to enable faster radio TX

Case Study – NRF52 IEEE802.15.4 Driver



*Add SW ACKs to
radio driver*

July 2023

August 2023

3 possible HW “shortcuts” to enable faster radio TX

Unable to get all 3 working (~2 weeks of development)

Case Study – NRF52 IEEE802.15.4 Driver



July 2023

*Add SW ACKs to
radio driver*

August 2023



*Integrate our
system into 15.4
driver*

March 2025

Unable to get all 3 shortcuts working (~2 weeks of development)

Case Study – NRF52 IEEE802.15.4 Driver



July 2023

*Add SW ACKs to
radio driver*

August 2023



*Integrate our
system into 15.4
driver*

March 2025

Unable to get all 3 shortcuts working (~2 weeks of development)

- Updated our state machine in DSL to use all 3 TX HW shortcuts

Case Study – NRF52 IEEE802.15.4 Driver



July 2023

*Add SW ACKs to
radio driver*

August 2023



*Integrate our
system into 15.4
driver*

March 2025

Unable to get all 3 shortcuts working (~2 weeks of development)

- Updated our state machine in DSL to use all 3 TX HW shortcuts
- Compiler identifies sections of driver that must be updated (errors)

Case Study – NRF52 IEEE802.15.4 Driver



July 2023

*Add SW ACKs to
radio driver*

August 2023



*Integrate our
system into 15.4
driver*

March 2025

Unable to get all 3 shortcuts working (~2 weeks of development)

- Updated our state machine in DSL to use all 3 TX HW shortcuts
- Compiler identifies sections of driver that must be updated (errors)

Updated and working driver in ~2 hours!

Case Study – NRF52 IEEE802.15.4 Driver



July 2023

*Add SW ACKs to
radio driver*

August 2023



*Integrate our
system into 15.4
driver*

March 2025

Unable to get all 3 shortcuts working (~2 weeks of development)

- Updated our state machine in DSL to use all 3 TX HW shortcuts
- Compiler identifies sections of driver that must be updated (errors)

Updated and working driver in ~2 hours!

50% decrease in driver interrupts; 8% runtime improvement

Our system statically prevents device protocol violations using tpestates.

Imposes minimal to no code size and runtime overheads.

Crate coming soon!

(will be a counterpart to tock-registers)

