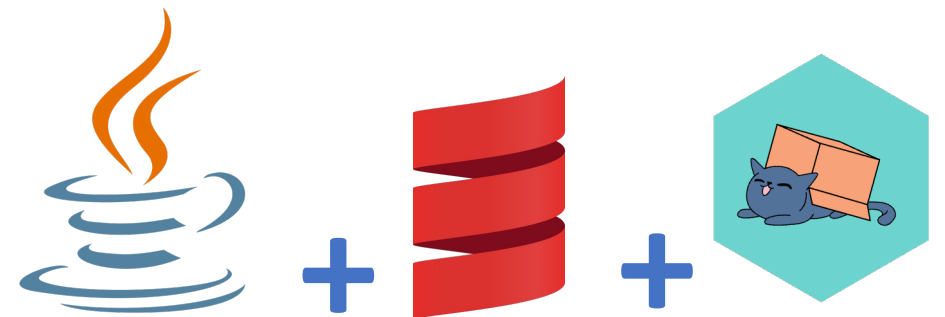


Elements of JVM Concurrency


In the context of Scala, and Cats-Effect

Alexandru Nedelcu
alexn.org







Plan

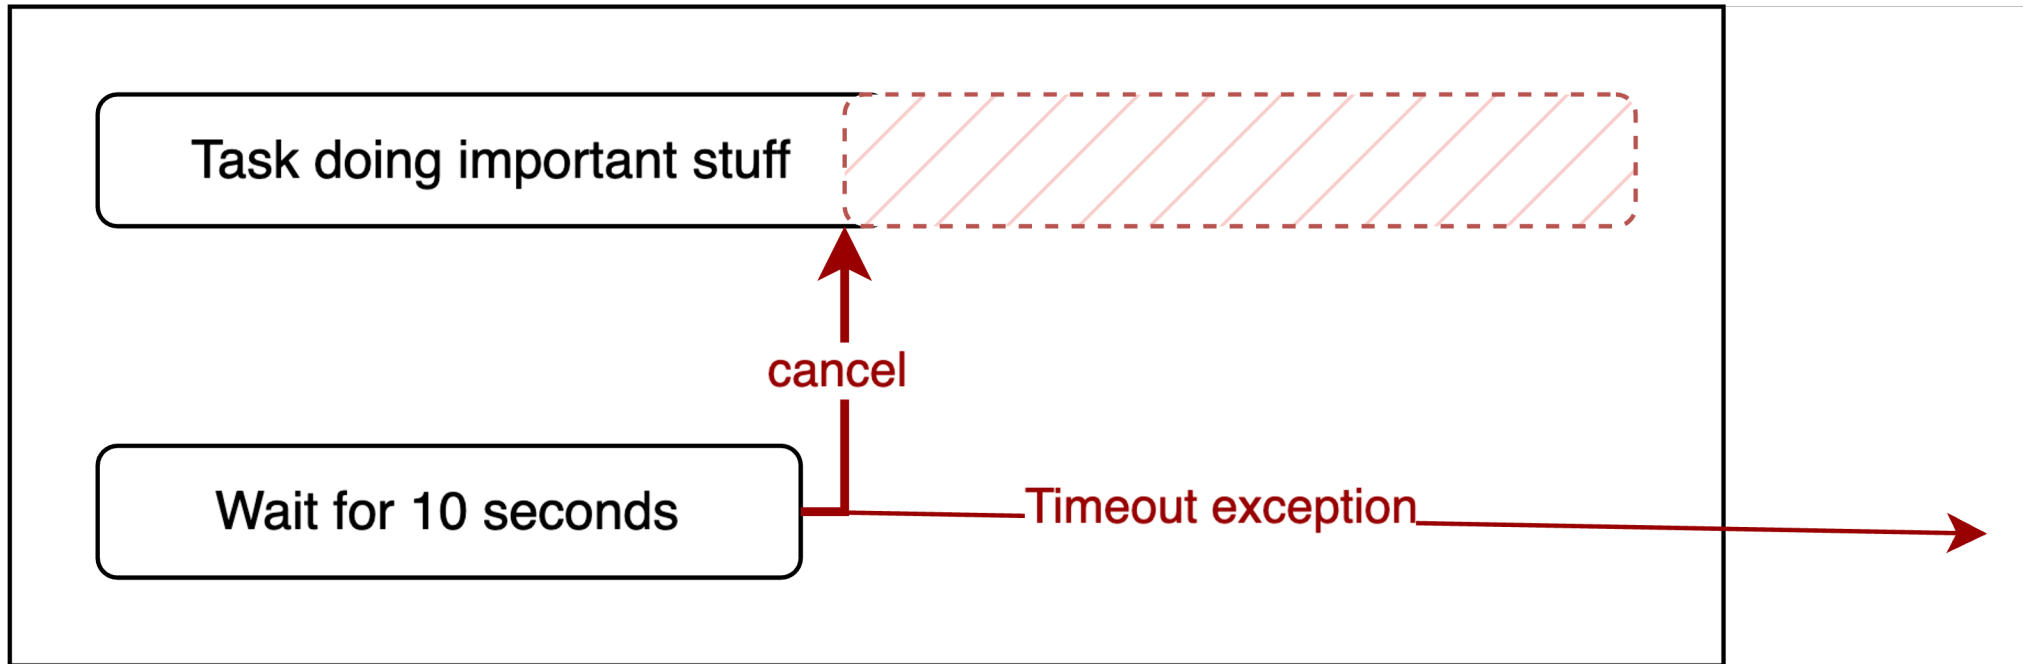
- Synchronization
 - Monitors, atomic references, Ref, volatiles
- Interruption
 - Java, Cats-Effect
- Goal: confidence boost



Concurrency vs. Parallelism



Interruption



```
val th = new Thread:  
    while !Thread.interrupted() do  
        process()  
  
th.start()  
//...  
th.interrupt()
```

```
val th = new Thread:  
    while true do  
        process()  
        Thread.sleep(1000)  
  
th.start()  
//...  
th.interrupt()
```

Java interruption is usable, but **error prone**

1. Must own the thread, or do expensive synchronization;
2. Code can catch and ignore interruption;
3. Interruption is too cooperative.



Just use `cats.effect.IO`

```
val task =  
  IO.interruptible:  
    while !Thread.interrupted() do  
      process()  
      Thread.sleep(1000)  
    .timeout(30.seconds)
```




```
trait ConcurrentQueue[A]:  
  def push(a: A): Unit  
  def pop(): A
```

```
ref.synchronized {  
  //...  
}
```

atomic references

```
import cats.effect.IO
```

```
class AtomicQueueVer1[A]:
```

```
  private val ref = new AtomicReference(immutable.Queue.empty[A])
```

```
  def push(event: A): IO[Unit] =
```

```
    IO.defer:
```

```
      val current = ref.get
```

```
      val update = current.enqueue(event)
```

```
      if ref.compareAndSet(current, update) then
```

```
        IO.unit
```

```
      else // RETRY!
```

```
        push(event)
```

```
import cats.effect.IO

class AtomicQueueVer1[A]:
  private val ref = new AtomicReference(immutable.Queue.empty[A])

  def push(event: A): IO[Unit] =
    IO.defer:
      val current = ref.get
      val update = current.enqueue(event)
      if ref.compareAndSet(current, update) then
        IO.unit
      else // RETRY!
        push(event)
```

```
import cats.effect.IO

class AtomicQueueVer1[A]:
  private val ref = new AtomicReference(immutable.Queue.empty[A])

  def push(event: A): IO[Unit] =
    IO.defer:
      val current = ref.get
      val update = current.enqueue(event)
      if ref.compareAndSet(current, update) then
        IO.unit
      else // RETRY!
        push(event)
```

```
import cats.effect.IO

class AtomicQueueVer1[A]:
  private val ref = new AtomicReference(immutable.Queue.empty[A])

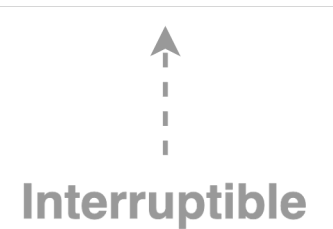
  def push(event: A): IO[Unit] =
    IO.defer:
      val current = ref.get
      val update = current.enqueue(event)
      if ref.compareAndSet(current, update) then
        IO.unit
      else // RETRY!
        push(event)
```

$f: \text{State} \Rightarrow (\text{State}, \text{Return})$


```
extension [State](self: AtomicReference[State])
  def modify[Return](f: State => (State, Return)): IO[Return] =
    IO.defer:
      val current = self.get
      val (updated, returnValue) = f(current)
      if (self.compareAndSet(current, updated))
        IO.pure(returnValue)
      else // RETRY!
        modify(f)
```

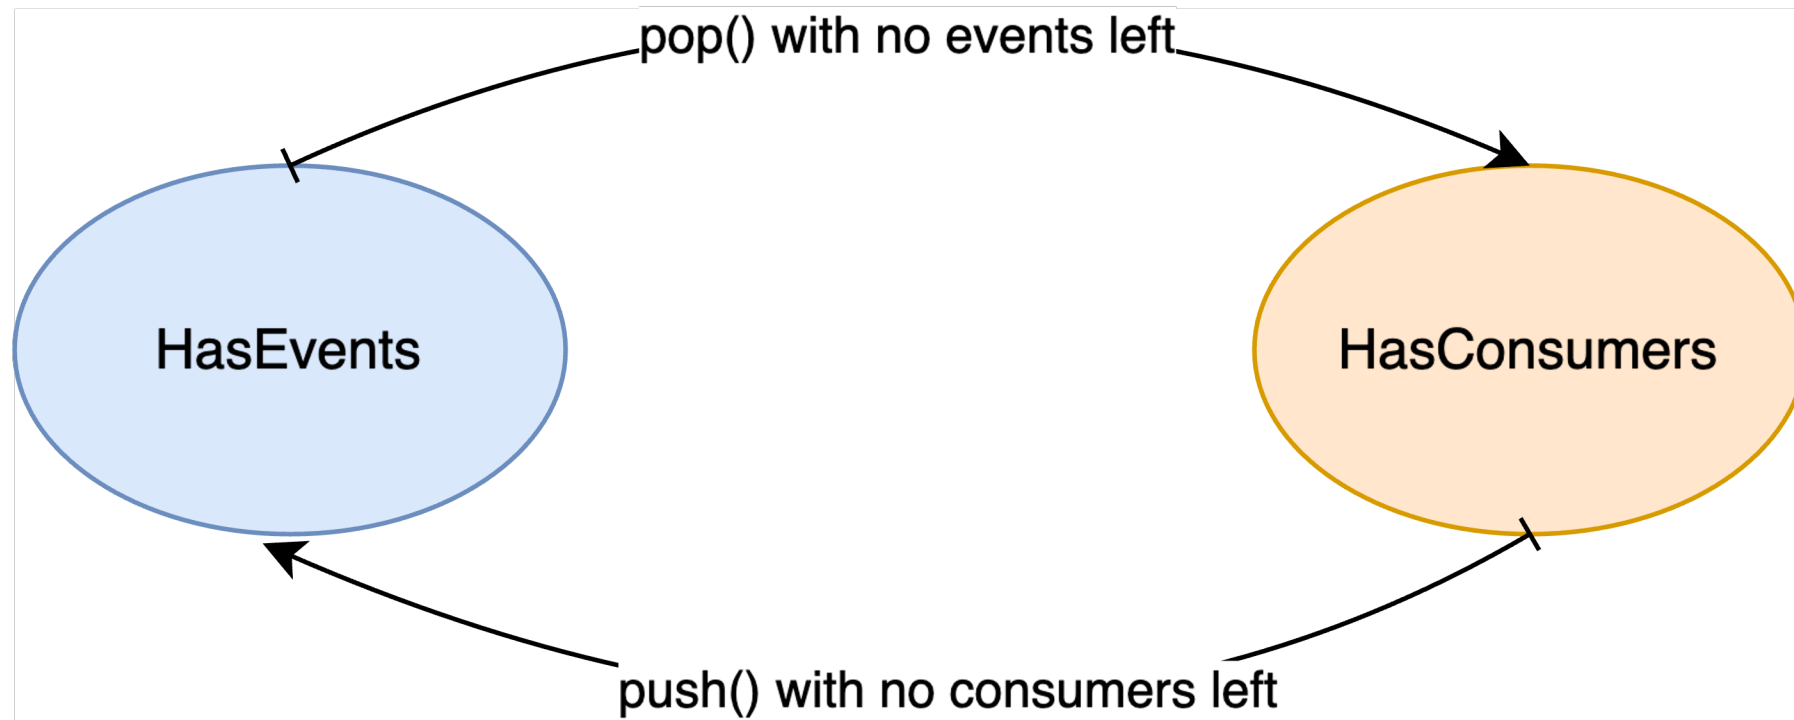
```
class AtomicQueueVer1[A]:  
  private val ref = new AtomicReference(immutable.Queue.empty[A])  
  
  def push(event: A): IO[Unit] =  
    ref.modify: queue =>  
      (queue.enqueue(event), ())
```

```
class AtomicQueueVer1[A]:  
  private val ref = new AtomicReference(immutable.Queue.empty[A])  
  
  def pop: IO[A] =  
    ref.modify: queue =>  
      queue.dequeueOption match  
      case Some((event, rest)) =>  
        (rest, IO.pure(event))  
      case None =>  
        val retry = IO.sleep(500.millis) *> pop  
        (queue, retry)  
    .flatten
```



Interruptible

State Machines, FTW!



```
import scala.collection.immutable.Queue as ScalaQueue

type Callback[-A] = Either[Throwable, A] => Unit

enum State[A]:
  case HasEvents(
    events: ScalaQueue[A]
  )
  case HasConsumers(
    callbacks: ScalaQueue[Callback[A]]
  )
```


`f: State => (State, IO[Return])`

```
enum State[A]:  
  //...  
  def pushEvent(event: A): (State[A], IO[Unit]) =  
    this match  
    case HasEvents(events) =>  
      (HasEvents(events.enqueue(event)), IO.unit)  
    case HasConsumers(consumers) =>  
      consumers.dequeueOption match  
      case None =>  
        (HasEvents(ScalaQueue(event)), IO.unit)  
      case Some((callback, rest)) =>  
        (HasConsumers(rest), IO(callback(Right(event))))
```



```
enum State[A]:  
  //...  
  def popEvent(cb: Callback[A]): (State[A], IO[Unit]) =  
    this match  
    case HasEvents(events) =>  
      events.dequeueOption match  
      case None =>  
        (HasConsumers(ScalaQueue(cb)), IO.unit)  
      case Some((event, rest)) =>  
        (HasEvents(rest), IO(cb(Right(event))))  
    case HasConsumers(callbacks) =>  
      (HasConsumers(callbacks.enqueue(cb)), IO.unit)
```

```
enum State[A]:  
  //...  
  def cancelPopAwait(cb: Callback[A]): State[A] =  
    this match  
    case ref: HasEvents[A] => ref  
    case HasConsumers(callbacks) =>  
      HasConsumers(callbacks.filterNot(_ == cb))
```



```
import cats.effect.*
```

```
class RefQueue[A](ref: Ref[IO, State[A]]):
```

```
  def push(event: A): IO[Unit] =  
    ref.flatModify(_.pushEvent(event))
```

```
  def pop: IO[A] =  
    IO.async: cb =>  
      for  
        _ <- ref.flatModify(_.popEvent(cb))  
      yield Some(  
        // Execute this in case of interruption  
        IO(ref.update(_.cancelPopAwait(cb)))  
      )
```

PROS:

- pure functions, immutable data structures
- cheap, often easy to implement
- async-friendly
- lock free

CONS:

- can have poor performance, sometimes

Plot Twist



Synchronization
is
Ordering!

```
class Queue[A]:  
  var hasEvent = false  
  var event: A | Null = null  
  
  def push(a: A): Unit = ???  
  
  def pop(): A = ???
```

```
class Queue[A]:  
  var hasEvent = false  
  var event: A | Null = null  
  
  def push(a: A): Unit =  
    while hasEvent do  
      Thread.onSpinWait()  
    event = a  
    hasEvent = true
```



```
class Queue[A]:  
  var hasEvent = false  
  var event: A | Null = null  
  
  def pop(): A =  
    while !hasEvent do  
      Thread.onSpinWait()  
    val ret = event.asInstanceOf[A]  
    hasEvent = false  
    ret
```

```
1 class Queue[A]:
2     var hasEvent = false
3     var event: A | Null = null
4
5     def pop(): A =
6         while !hasEvent do
7             Thread.onSpinWait()
8             hasEvent = false // RE-ORDERED!!!
9             val ret = event.asInstanceOf[A]
10            ret
```

```
1 class Queue[A]:
2     var hasEvent = false
3     var event: A | Null = null
4
5     def push(a: A): Unit =
6         while hasEvent do
7             Thread.onSpinWait()
8             hasEvent = true // RE-ORDERED!
9             event = a
```

```
// ...  
// producer thread  
queue.push("value")  
  
// consumer thread  
println(queue.hasEvent, queue.event)  
  
//=> (false, null)  
//=> (false, "value")  
//=> (true, "value")  
//=> (true, null)
```

@volatile

```
class SpScQueue[A]:  
  @volatile var hasEvent = false  
  var event: A | Null = null  
  
  def push(a: A): Unit =  
    while hasEvent do  
      Thread.onSpinWait()  
    event = a  
    hasEvent = true
```

```
1 class SpScQueue[A]:  
2     @volatile var hasEvent = false  
3     var event: A | Null = null  
4  
5     def pop(): A =  
6         while !hasEvent do  
7             Thread.onSpinWait()  
8         val ret = event.asInstanceOf[A]  
9         hasEvent = false  
10        ret
```

Reading from a volatile	Acquiring a lock
Writing to a volatile	Releasing a lock

- **acquisition and release needs to happen on the same monitor**
- **volatile reading and writing needs to happen on the same variable**

- **acquisition and release needs to happen in pair**
- **reading and writing from a volatile needs to happen in pair**

- **atomic references and monitor locks have the same visibility guarantees, due to volatile semantics**

Takeaways

- State machines, FTW
- **Synchronization is ordering**
- Meet concurrency issues with the confidence of Zuzi

