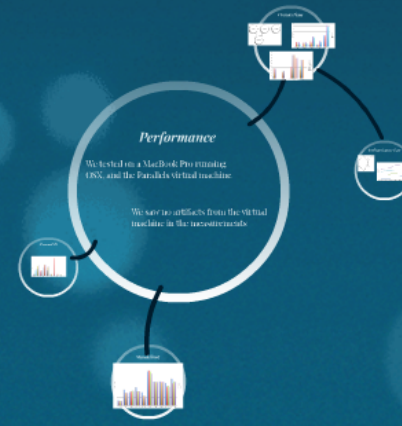
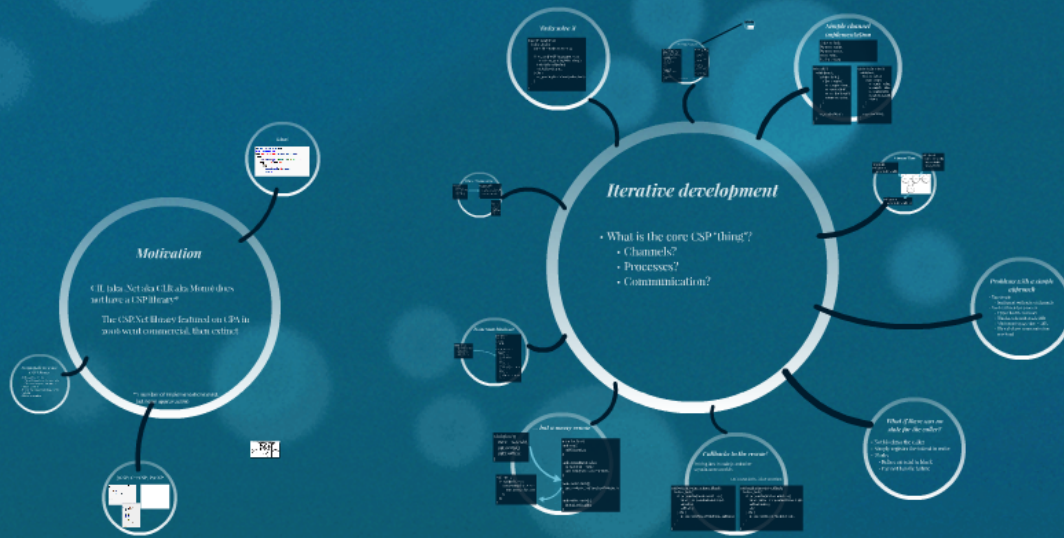




# CoCoL

## Concurrent Communication Library



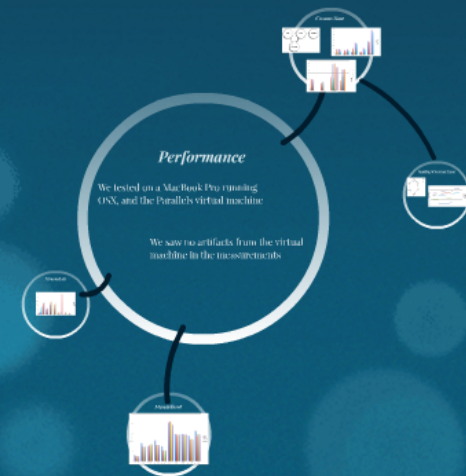
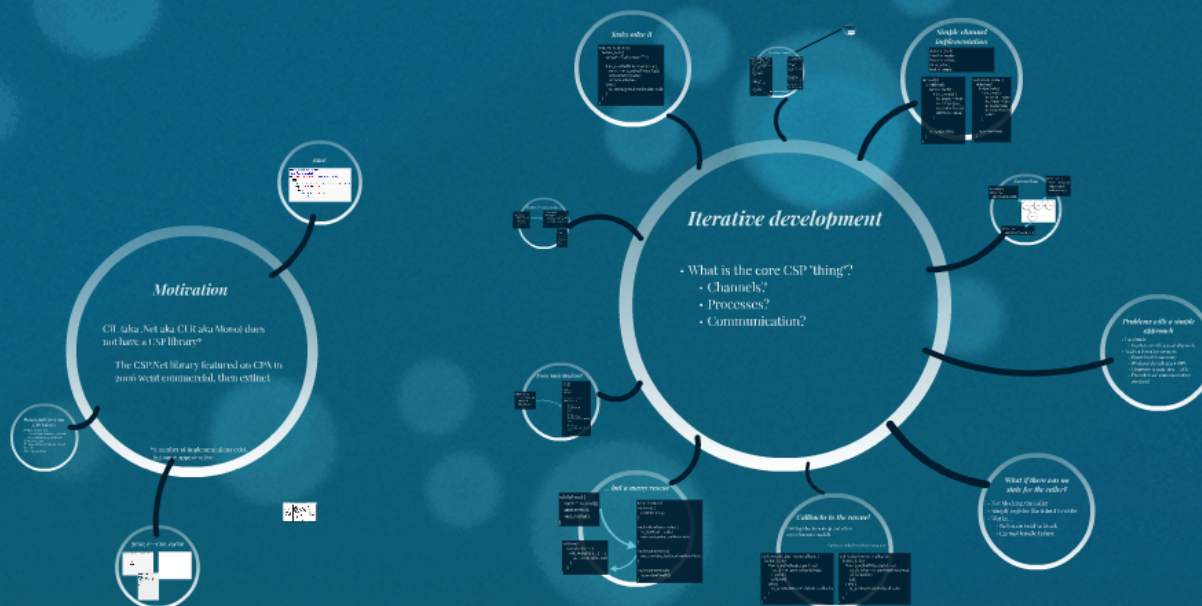
CPA 2015-08-25  
 Kenneth Skovhede  
 Niels Bohr Institute  
 University of Copenhagen





# CoCoL

## Concurrent Communication Library



CPA 2015-08-25  
Kenneth Skovhede  
Niels Bohr Institute  
University of Copenhagen







# *Motivation*

CIL (aka .Net aka CLR aka Mono) does not have a CSP library\*

The CSP.Net library featured on CPA in 2006 went commercial, then extinct

\*A number of implementations exist, but none appear active

## *Design goals for a new CSP Library*

- Fully contained in C#
  - No additional compilers or tools
  - No extra language constructs
- Simple codebase
- Fit into existing terminology & code
- Scalable
- Efficient execution



# *KRoC*

```
<<hello_world-kroc.occ>>=  
#USE "course.lib"  
PROC hello.world (CHAN BYTE scr!)  
  SEQ  
    out.string("Hello World!*n", 0, scr)  
    SEQ i = 1 FOR 10  
      SEQ  
        out.int(i, 0, scr)  
        scr ! '*n'  
  :
```



# JCSP, C++CSP, PyCSP

```
class IncreasingNumbers : public CSProcess
{
private:
    Chanout<int> out;
protected:
    void run()
    {
        for (int i = 1; ;i++)
        {
            out << i;
        }
    }
public:
    IncreasingNumbers(const Chanout<int>& _out)
        : out(_out)
    {
    }
};
```

```
from pycsp.parallel import *
```

```
@process
def counter(cout, limit):
    for i in xrange(limit):
        cout(i)
    poison(cout)
```

```
@process
def printer(cin):
    while True:
        print cin(),
```

```
A = Channel('A')
Parallel(
    counter(A.writer(), limit=10),
    printer(A.reader())
)
```

```
shutdown()
```

```
import org.jcsp.lang.*;
import org.jcsp.pluginPlay.*;

class ParalelexIntExample {

    public static void main (String[] args) {

        final One2OneChannelInt[] a = Channel.one2OneIntArray (3);
        final One2OneChannel b = Channel.one2One ();

        new Parallel (
            new CSProcess[] {
                new NumbersInt (a[0].out ()),
                new SquaresInt (a[1].out ()),
                new FibonacciInt (a[2].out ()),
                new ParalelexInt (Channel.getInputArray (a), b.out ()),
                new CSProcess () {
                    public void run () {
                        System.out.println ("\n\t\tNumbers\t\tSquares\t\tFibonacci\n");
                        while (true) {
                            int[] data = (int[]) b.in ().read ();
                            for (int i = 0; i < data.length; i++) {
                                System.out.print ("\t\t" + data[i]);
                            }
                            System.out.println ();
                        }
                    }
                }
            )
        }.run ();
    }
}
```



## *Design goals for a new CSP Library*

- Fully contained in C#
  - No additional compilers or tools
  - No extra language constructs
- Simple codebase
- Fit into existing terminology & code
- Scalable
- Efficient execution



# HOW STANDARDS PROLIFERATE:

(SEE: A/C CHARGERS, CHARACTER ENCODINGS, INSTANT MESSAGING, ETC.)

SITUATION:  
THERE ARE  
14 COMPETING  
STANDARDS.

14?! RIDICULOUS!  
WE NEED TO DEVELOP  
ONE UNIVERSAL STANDARD  
THAT COVERS EVERYONE'S  
USE CASES.



YEAH!

SOON:

SITUATION:  
THERE ARE  
15 COMPETING  
STANDARDS.



# Iterative development

- What is the core CSP "thing"?
- Channels?
- Processes?
- Communication?

## Tasks solve it

```
Task T = read(T) {
    lock(m_lock)
    var task = Task.Create(T) {
        if (m_pendingWrites.Count > 0) {
            var w = m_pendingWrites.Pop();
            task.SetResult(w);
            w.SetResult(true);
        } else {
            m_pendingReads.Push(task);
        }
    }
}
```

### Two processes

```
using System;
using System.Threading;

class Program {
    static void Main() {
        var lock = new Lock();
        var t1 = Task.Run(() => {
            lock.Acquire();
            Console.WriteLine("Process 1");
            lock.Release();
        });
        var t2 = Task.Run(() => {
            lock.Acquire();
            Console.WriteLine("Process 2");
            lock.Release();
        });
        Task.WaitAll(t1, t2);
    }
}
```

## Simple channel implementation

```
object m_lock;
Event m_reader;
Event m_writer;
int m_value;
bool m_empty;

int read() {
    while(true) {
        lock(m_lock);
        if (m_empty) {
            m_empty = true;
            m_reader.Set();
            m_reader.Reset();
            return m_value;
        }
        m_reader.Wait();
    }
}

void write(int value) {
    while(true) {
        lock(m_lock);
        if (m_empty) {
            m_value = value;
            m_empty = false;
            m_writer.Set();
            m_writer.Reset();
            return;
        }
        m_writer.Wait();
    }
}
```

## Commas Time



## Problems with a simple approach

- Too simple
- Inefficient with sets of channels
- Needs a thread per process
- Upper limit is memory
- Windows default stack = 1MB
- Minimum is page size = 4KB
- Threads have communication overhead

## What if there was no state for the caller?

- Not blocking the caller
- Simply register the intent to write
- Works....
- Relies on read to block
- Cannot handle failure

## Callbacks to the rescue!

Driving idea in node.js and other asynchronous models

Can be extended to deliver exceptions

```
void write(int value, Action callback) {
    lock(m_lock);
    if (m_pendingWrites.Count > 0) {
        var cb = m_pendingWrites.Pop();
        cb(value);
        cb();
    } else {
        m_pendingWrites.Push(value, callback);
    }
}

void read(Action cb = null) {
    lock(m_lock);
    if (m_pendingReads.Count > 0) {
        var cb = m_pendingReads.Pop();
        cb();
    } else {
        m_pendingReads.Push(cb);
    }
}
```

## ... but a messy rescue

```
while(true) {
    var v = in.read();
    out1.write(v);
    out2.write(v);
}

int m_inRead;
void read() {
    lock(m_lock);
    m_inRead = value;
    out1.write(value, mWriteOut1);
}

void mWriteOut1() {
    out1.write(m_inRead, mWriteOut1);
}

void mWriteOut2() {
    out2.write(m_inRead);
}

void read() {
    in.read(value => {
        out1.write(value, () => {
            out2.write(value, read);
        });
    });
}
```

## Finite State Machine?

```
enum State {
    Idle,
    Reading,
    Writing,
    Error
};

class Channel {
    State state;
    int value;
    Action callback;
    void read() {
        lock(m_lock);
        if (state == Reading) {
            return;
        }
        state = Reading;
        value = in.read();
        callback();
        state = Idle;
    }
    void write(int value) {
        lock(m_lock);
        if (state == Writing) {
            return;
        }
        state = Writing;
        this.value = value;
        callback();
        state = Idle;
    }
}
```

## Future, Promise, Async

```
void read() {
    lock(m_lock);
    if (m_empty) {
        m_empty = true;
        m_reader.Set();
        m_reader.Reset();
        return m_value;
    }
    m_reader.Wait();
}

void write(int value) {
    lock(m_lock);
    if (m_empty) {
        m_value = value;
        m_empty = false;
        m_writer.Set();
        m_writer.Reset();
        return;
    }
    m_writer.Wait();
}
```



# *Iterative development*

- What is the core CSP "thing"?
  - Channels?
  - Processes?
  - Communication?

te Machine?

```
let in, state, out
  into value
  while read!
    call back!
  }
  out! call back! end!
  switch on, state!
  case 0:
    in, state = 1;
    in, read! back!
    break;
  case 1:
    in, state = 2;
    in, state = value;
    out, write! in, value, call back!
    break;
  case 2:
    in, state = 0;
    out, call back! in, value, call back!
    break;
  }
```

```
in, write! 0;
while (true)
  out, write! in, read!
```

```
while (true)
  out, write!
```



# *Simple channel implementation*

```
object m_lock;  
Event m_reader;  
Event m_writer;  
int m_value;  
bool m_empty
```

```
int read() {  
    while(true) {  
        lock(m_lock) {  
            if (!m_empty) {  
                m_empty = true;  
                m_writer.Set();  
                m_reader.Reset();  
                return m_value;  
            }  
        }  
        m_reader.Wait();  
    }  
}
```

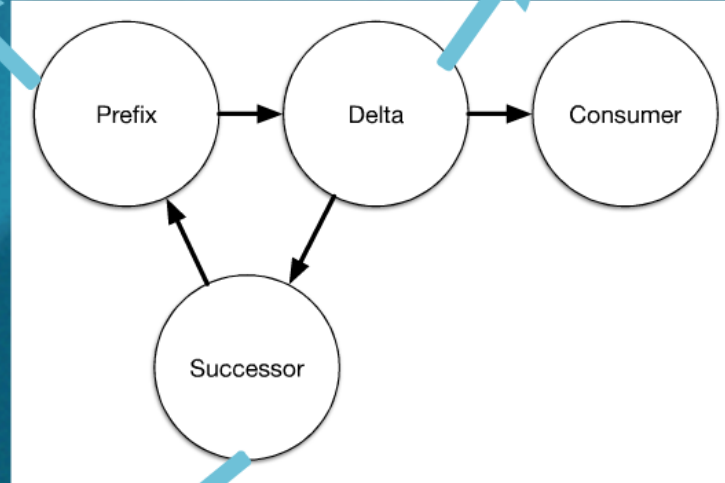
```
void write(int value) {  
    while(true) {  
        lock(m_lock) {  
            if (m_empty) {  
                m_result = value;  
                m_empty = false;  
                m_reader.Set();  
                m_writer.Reset();  
                return;  
            }  
        }  
        m_writer.Wait();  
    }  
}
```



# CommsTime

```
in.write(o);  
while(true)  
  out.write(in.read());
```

```
while(true) {  
  var v = in.read();  
  out1.write(v);  
  out2.write(v);  
}
```



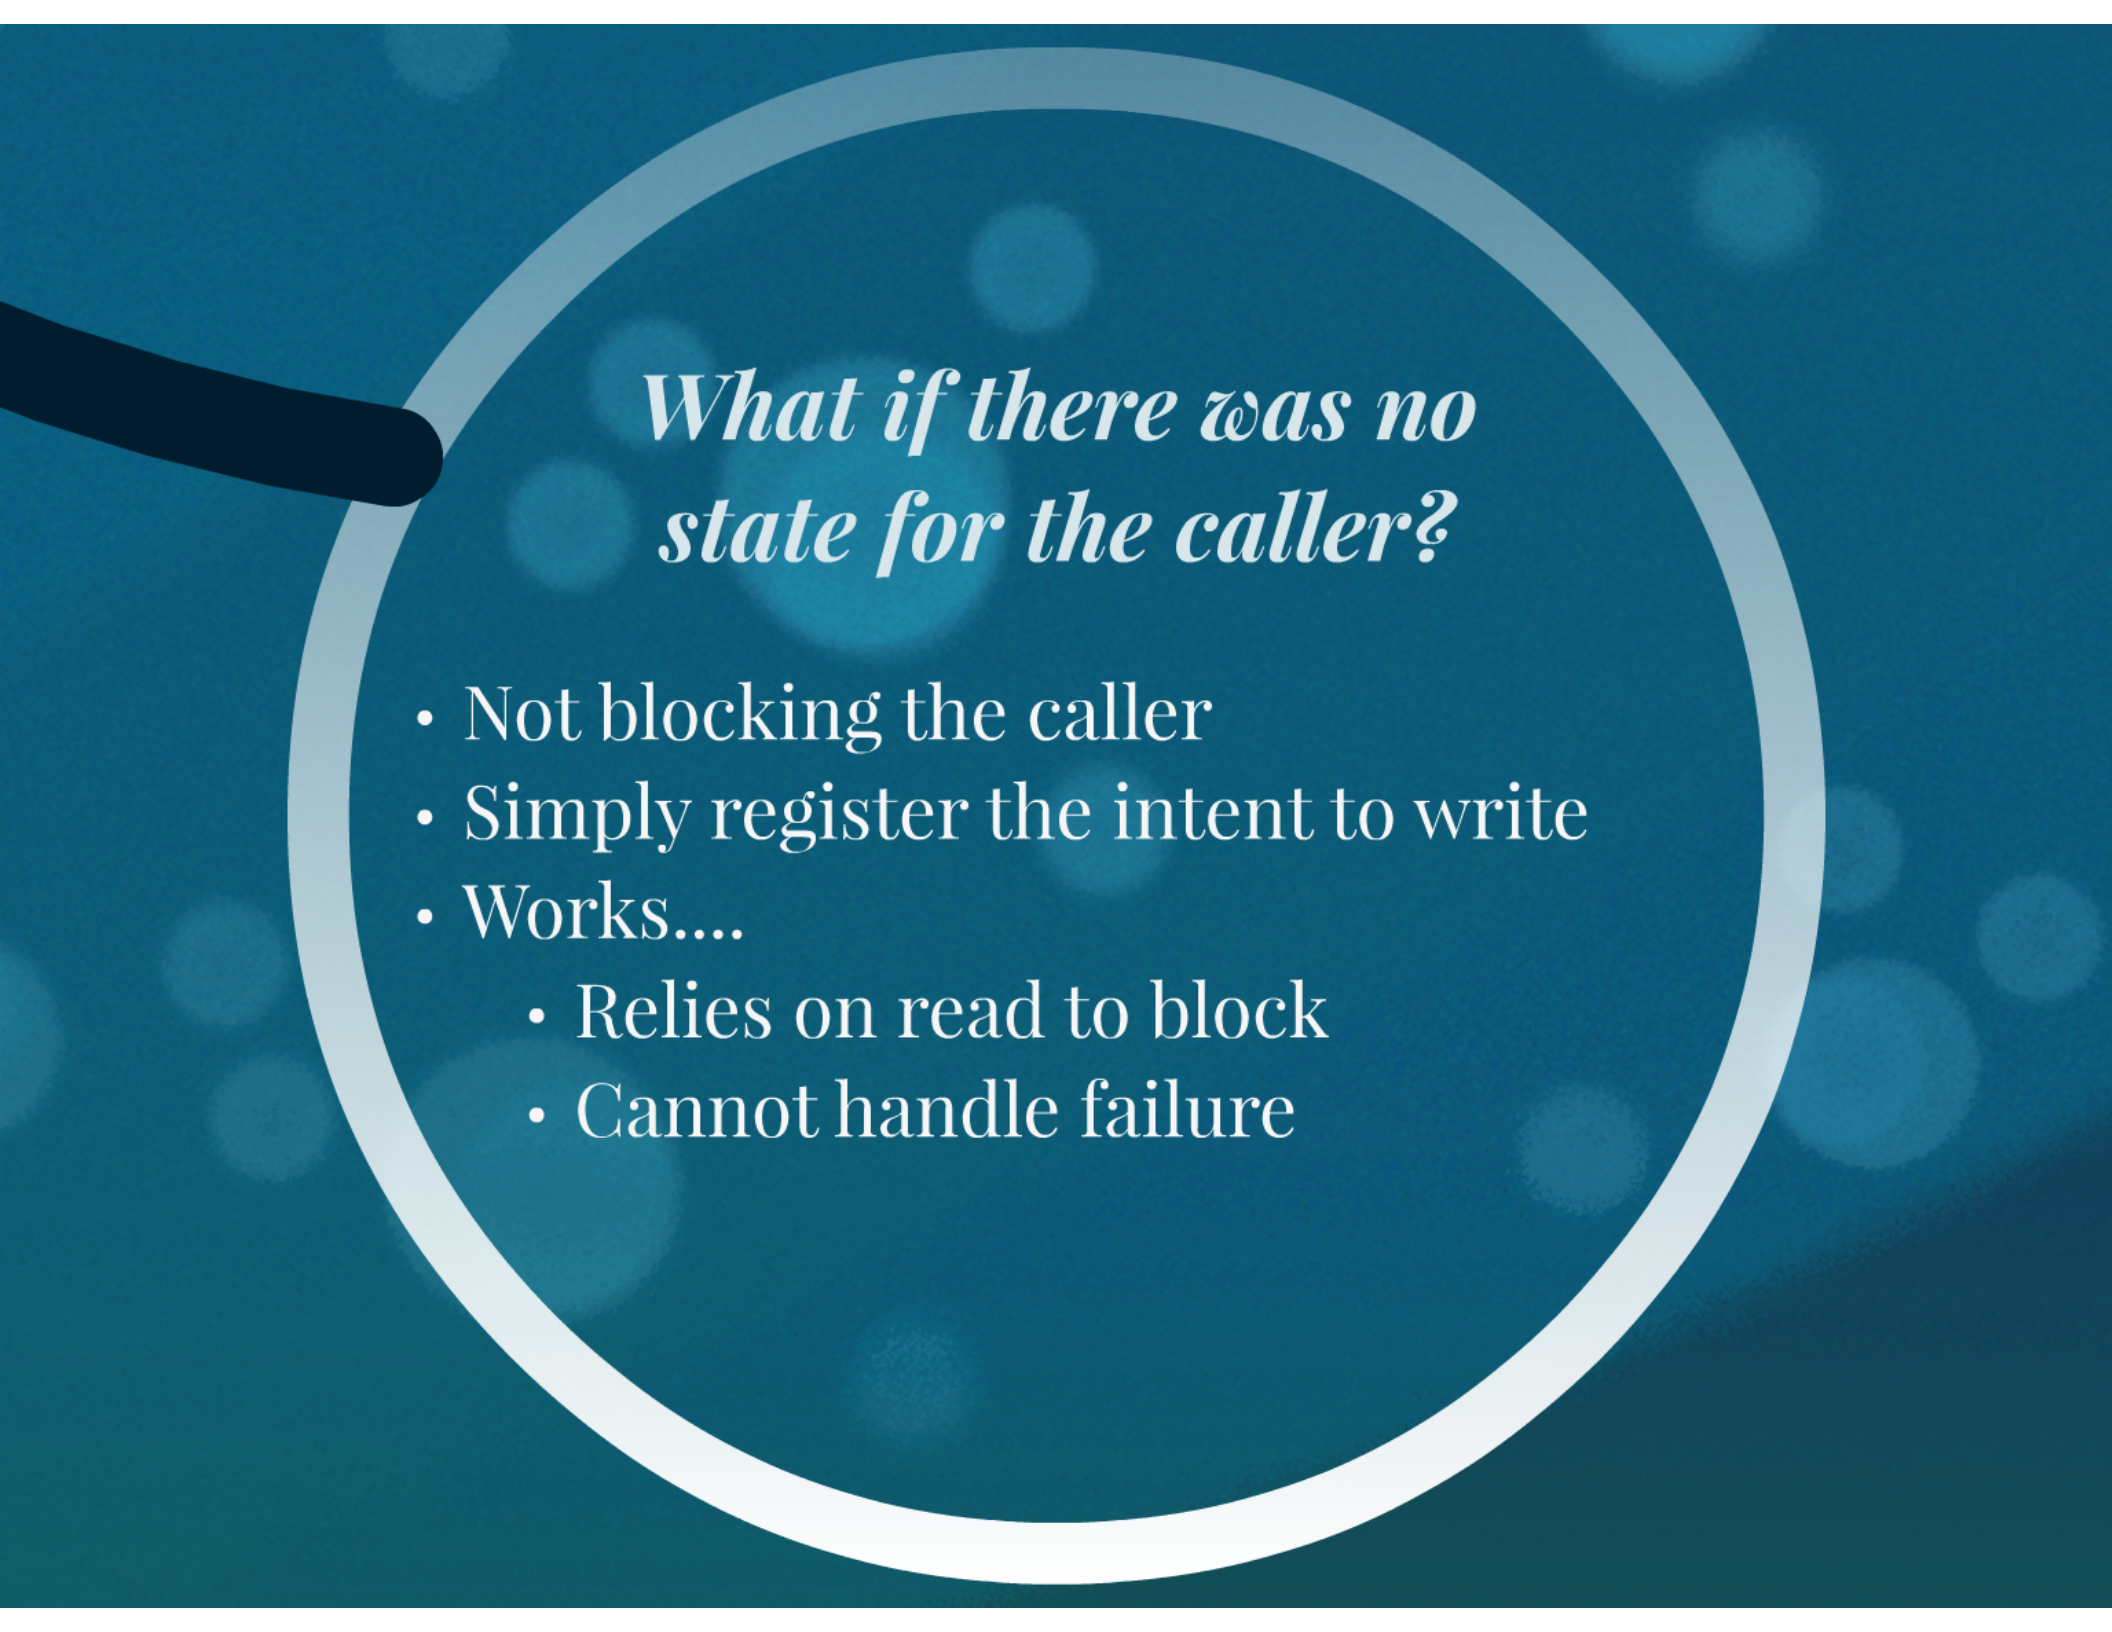
```
while(true)  
  out.write(in.read() + 1)
```



## *Problems with a simple approach*

- Too simple
  - Inefficient with sets of channels
- Needs a thread pr process
  - Upper limit is memory
  - Windows default stack 1MB
  - Minimum is page size = 4KB
  - Threads have communication overhead





*What if there was no  
state for the caller?*

- Not blocking the caller
- Simply register the intent to write
- Works....
  - Relies on read to block
  - Cannot handle failure



# *Callbacks to the rescue!*

Driving idea in node.js and other asynchronous models

Can be extended to deliver exceptions

```
void write(int value, Action callback) {  
    lock(m_lock) {  
        if (m_pendingReads.Count > 0) {  
            var cb = m_pendingReads.Pop();  
            cb(value);  
            callback();  
        } else {  
            m_pendingWrites.Push(value, callback);  
        }  
    }  
}
```

```
void read(Action<int> callback) {  
    lock(m_lock) {  
        if (m_pendingWrites.Count > 0) {  
            var cb, value = m_pendingWrites.Pop();  
            callback(value);  
            cb();  
        } else {  
            m_pendingReads.Push(callback);  
        }  
    }  
}
```



## *... but a messy rescue*

```
while(true) {  
    var v = in.read();  
    out1.write();  
    out2.write();  
}
```

```
void run() {  
    in.read(value => {  
        out1.write(value, () => {  
            out2.write(value, run)  
        })  
    });  
}
```

```
int m_lastRead;  
void run() {  
    onWriteOut2();  
}  
  
void onReadIn(int value) {  
    m_lastRead = value;  
    out1.write(value, onWriteOut1);  
}  
  
void onWriteOut1() {  
    out2.write(m_lastRead, onWriteOut2);  
}  
  
void onWriteOut2() {  
    in.read(onReadIn);  
}
```



# *Finite State Machine?*


```
while(true) {  
    var v = in.read();  
    out1.write();  
    out2.write();  
}
```

```
int m_state = 0;  
int m_value;  
  
void run() {  
    callback();  
}  
  
void callback(int? value) {  
    switch (m_state) {  
        case 0:  
            m_state = 1;  
            in.read(callback);  
            break;  
        case 1:  
            m_state = 2;  
            m_value = value;  
            out1.write(m_value, callback);  
            break;  
        case 2:  
            m_state = 0;  
            out2.write(m_value, callback);  
            break;  
    }  
}
```



# *Future, Promise, async ...*

```
void run() {  
  while(true) {  
    var v = in.read();  
    out1.write();  
    out2.write();  
  }  
}
```



```
void async run() {  
  while(true) {  
    var value = await in.read();  
    await out1.write(value);  
    await out2.write(value);  
  }  
}
```

```
int m_state = 0;  
int m_value;  
  
void callback(int? value) {  
  switch (m_state) {  
    case 0:  
      m_state = 1;  
      in.read(callback);  
      break;  
    case 1:  
      m_state = 2;  
      m_value = value;  
      out1.write(m_value, callback);  
      break;  
    case 2:  
      m_state = 0;  
      out2.write(m_value, callback);  
      break;  
  }  
}
```



## *Tasks solve it*

```
Task<T> read<T>() {
    lock(m_lock) {
        var task = Task.Create<T>();

        if (m_pendingWrites.Count > 0) {
            var wt = m_pendingWrites.Pop();
            task.SetResult(value);
            wt.SetResult(true);
        } else {
            m_pendingReads.Push(value, task);
        }
    }
}
```

# *Two-phase commit*

```
while (m_pendingWrites.Count > 0) {
    var wt = m_pendingWrites.Peek();
    var writerAccepts = wt.Offer();
    var readerAccepts = task.Offer();
    if (writerAccepts && readerAccepts) {
        m_pendingWrites.Pop();
        wt.Commit();
        task.Commit();
        task.SetResult(value);
        wt.SetResult(true);
        return;
    }

    if (writerAccepts) {
        wt.WithDraw();
        m_pendingWrites.Pop();
    }

    if (readerAccepts) {
        task.WithDraw();
        return;
    }
}
```

```
bool Offer(object caller) {
    Monitor.Enter(m_lock);

    if (!m_taken)
        return true;

    Monitor.Exit(m_lock);
    return false;
}

void Commit() {
    m_taken = true;
    Monitor.Exit(m_lock);
}

void WithDraw() {
    Monitor.Exit(m_lock);
}
```

Many optimizations implemented ...




# *External choice*

```
Task<T> readFromAny(IEnumerable<Channel<T>> channels) {  
    var twophase = new TwoPhaseCommit();  
    return Task.WhenAny(from c in channels select c.read(twophase));  
}
```

Timeouts are handled with a shared timer, skip is done with timeout=0


Fair select relies on a "usage counter"



# Performance

We tested on a MacBook Pro running OSX, and the Parallels virtual machine

We saw no artifacts from the virtual machine in the measurements






# *Performance*



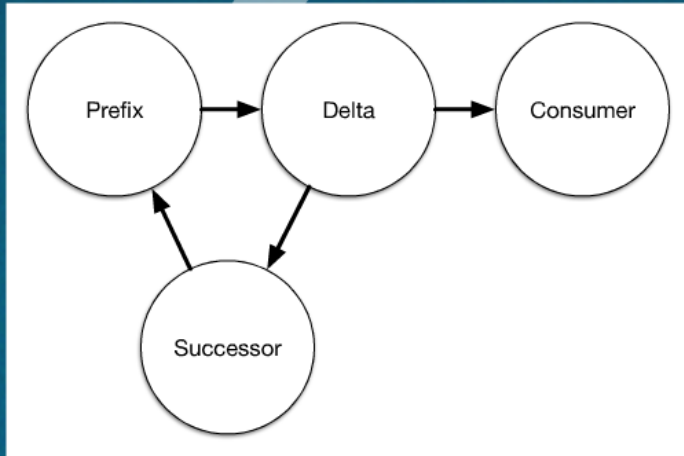
We tested on a MacBook Pro running OSX, and the Parallels virtual machine

We saw no artifacts from the virtual machine in the measurements


*StressedAlt*




# CommsTime















# Scaling Comms Time









# Mandelbrot








# *StressedAlt*







*All code, including examples / benchmarks:*

*<https://github.com/kenkendk/cocol>*

*Small code footprint:*

*Channel is appx 300 SLOC*

*Entire library is 1500 SLOC*