

Verifying a higher-order, concurrent, stateful library

Kasper Svendsen, Lars Birkedal and Matthew Parkinson

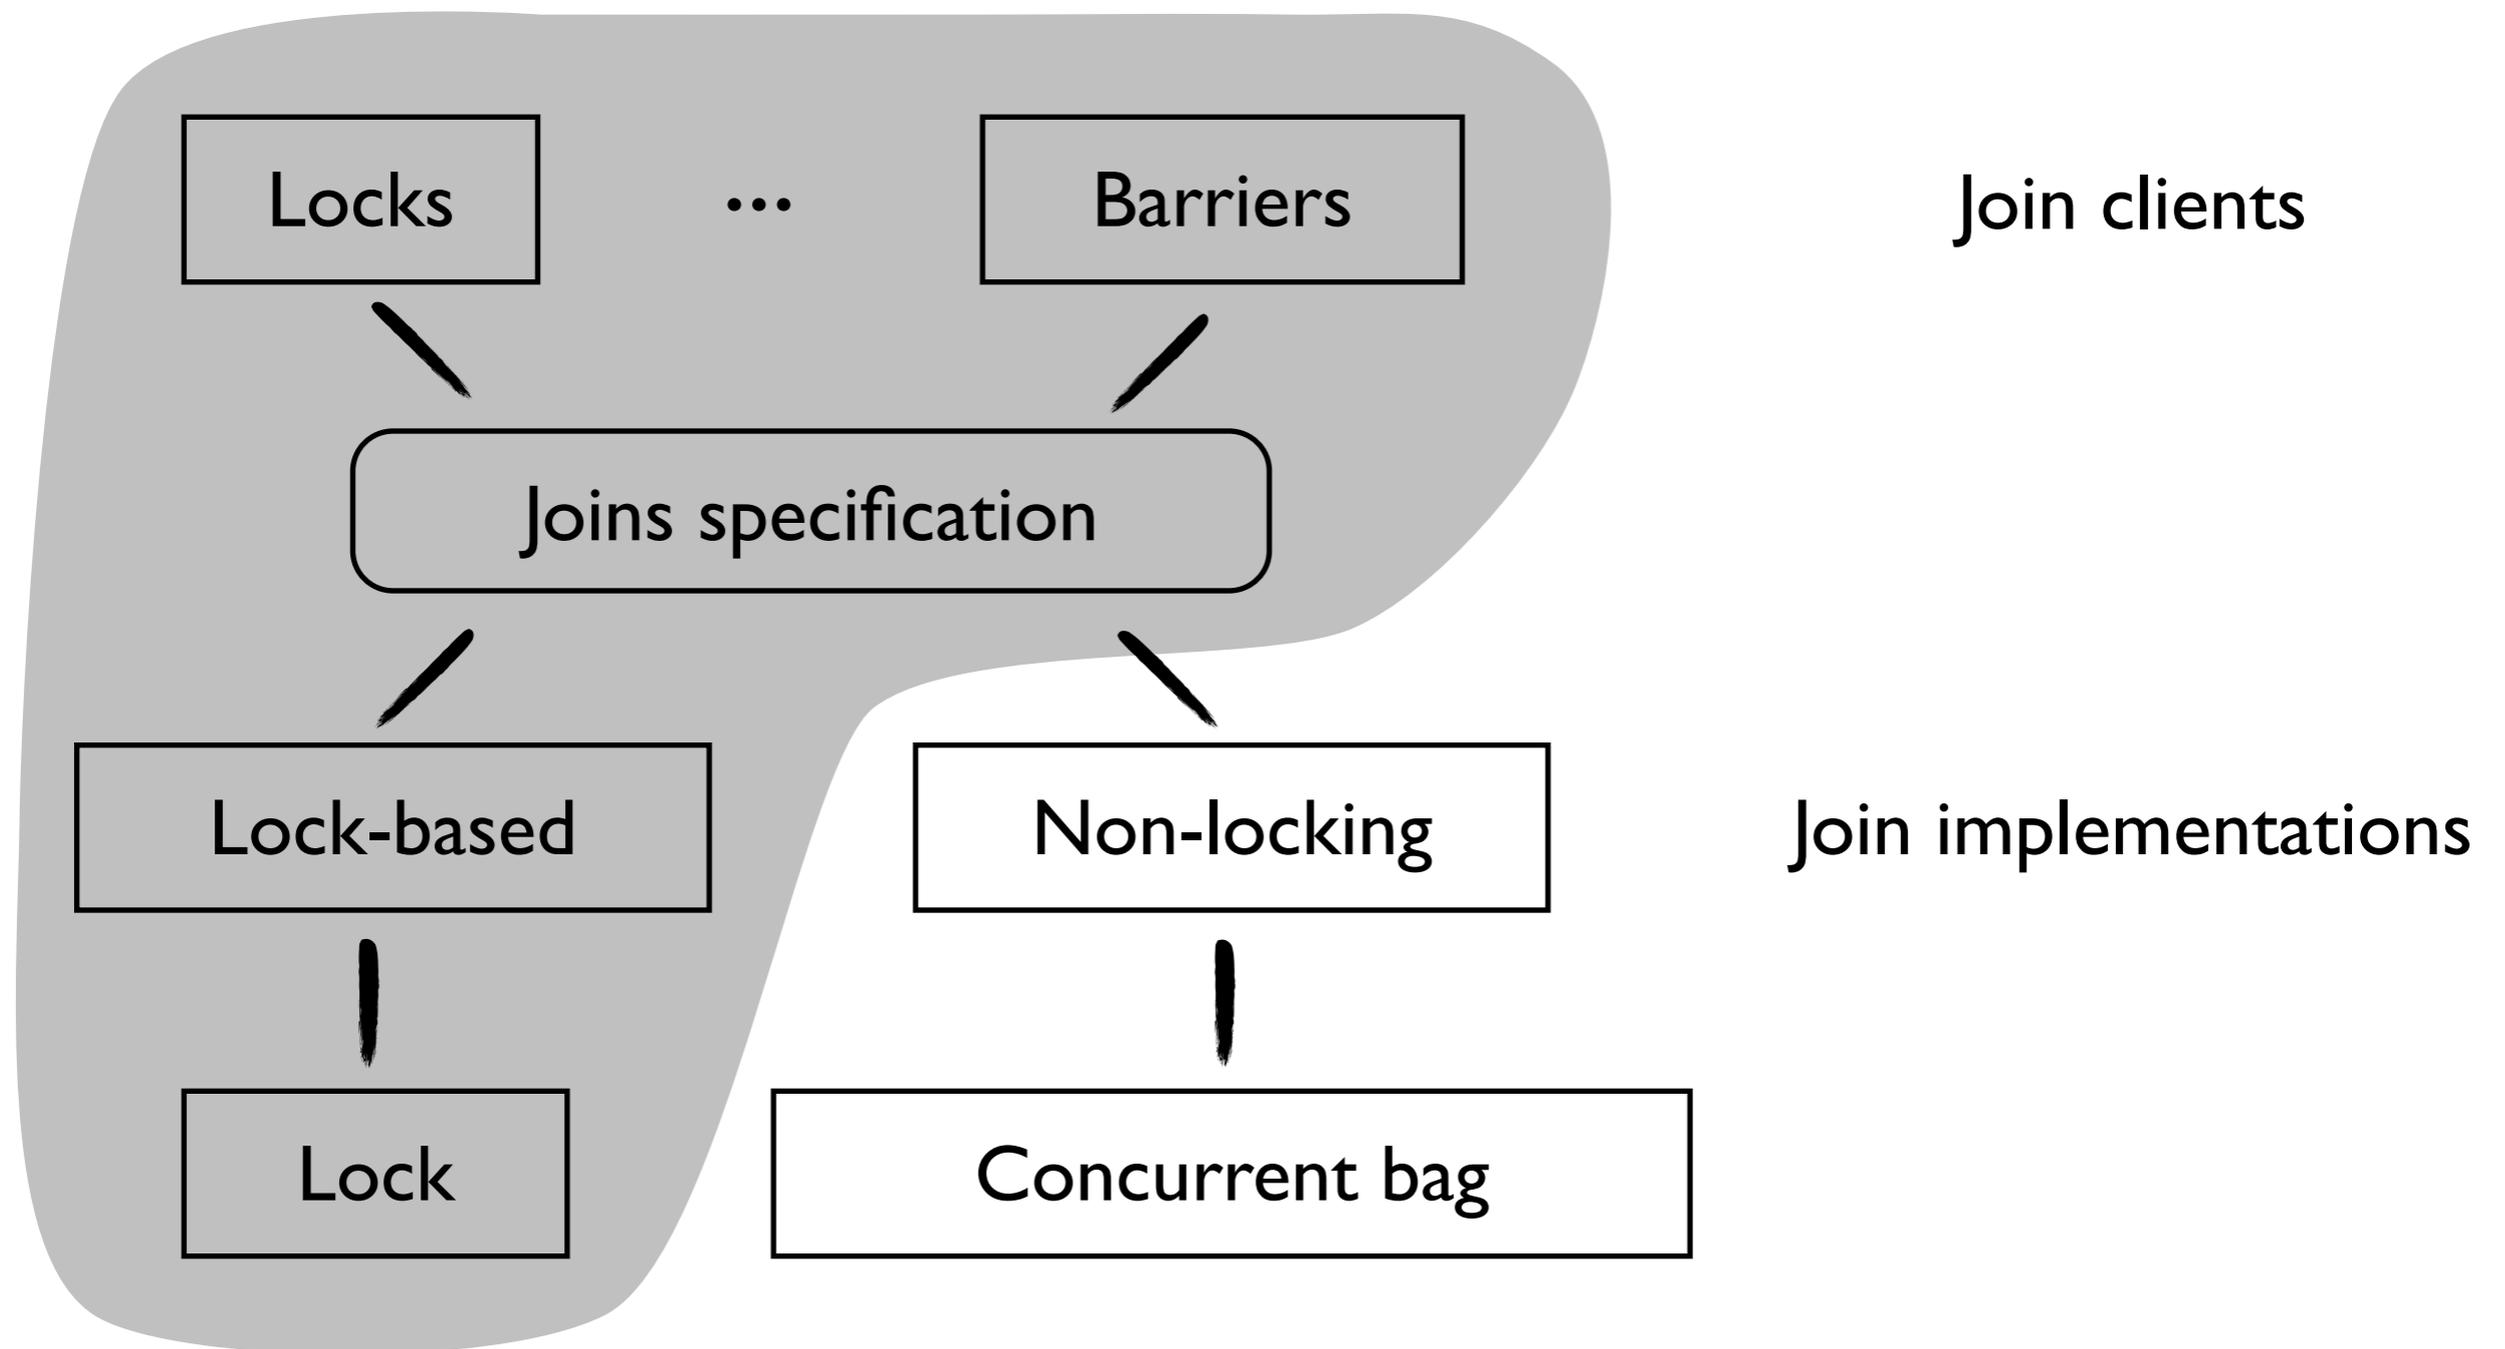
September 9, 2012

HOPE 2012

A case study ...

- C# Joins library [Russo, Turon & Russo]
 - declarative way of defining synchronization primitives, based on the join calculus [Fournet & Gonthier]
 - combines higher-order features with state, concurrency, recursion through the store and fine-grained synchronization
 - small (150 lines of C#) realistic library

A case study in modularity



Joins example

```
class RWLock {  
    public SyncChannel acqR, acqW, relR, relW;  
    private AsyncChannel unused, shared, writer;  
    private int readers = 0;
```

channels

```
    public RWLock() {  
        Join join = new Join();  
        // ... initialize channels ...
```

pattern

```
        join.When(acqR).And(unused).Do(() => { readers++; shared(); });  
        join.When(acqR).And(shared).Do(() => { readers++; shared(); });  
        join.When(acqW).And(unused).Do(() => { writer(); });  
        join.When(relW).And(writer).Do(() => { unused(); });  
        join.When(relR).And(shared).Do(() => {  
            if (--readers == 0) unused() else shared(); });
```

chord

```
        unused();
```

send a message on
the unused channel

continuation

A reader/writer lock

synchronous channels to
acquire and release the lock

asynchronous channels
encode the state of the lock

```
class RWLock {  
  public SyncChannel acqR, acqW, relR, relW;  
  private AsyncChannel unused, shared, writer;  
  private int readers = 0;
```

```
  public RWLock() {  
    Join join = new Join();  
    // ... initialize channels ...
```

each chord matches and sends
exactly one asynchronous message

```
    join.When(acqR).And(unused).Do(() => { readers++; shared(); });  
    join.When(acqR).And(shared).Do(() => { readers++; shared(); });  
    join.When(acqW).And(unused).Do(() => { writer(); });  
    join.When(relW).And(writer).Do(() => { unused(); });  
    join.When(relR).And(shared).Do(() => {  
      if (--readers == 0) unused() else shared(); });
```

```
    unused();
```

initially, there is exactly one
pending asynchronous message

```
  }  
}
```

Verification challenges

```
class RWLock {
  public SyncChannel acqR, acqW, relR, relW;
  private AsyncChannel unused, shared, writer;
  private int readers = 0;

  public RWLock() {
    Join join = new Join();
    // ... initialize channels ...

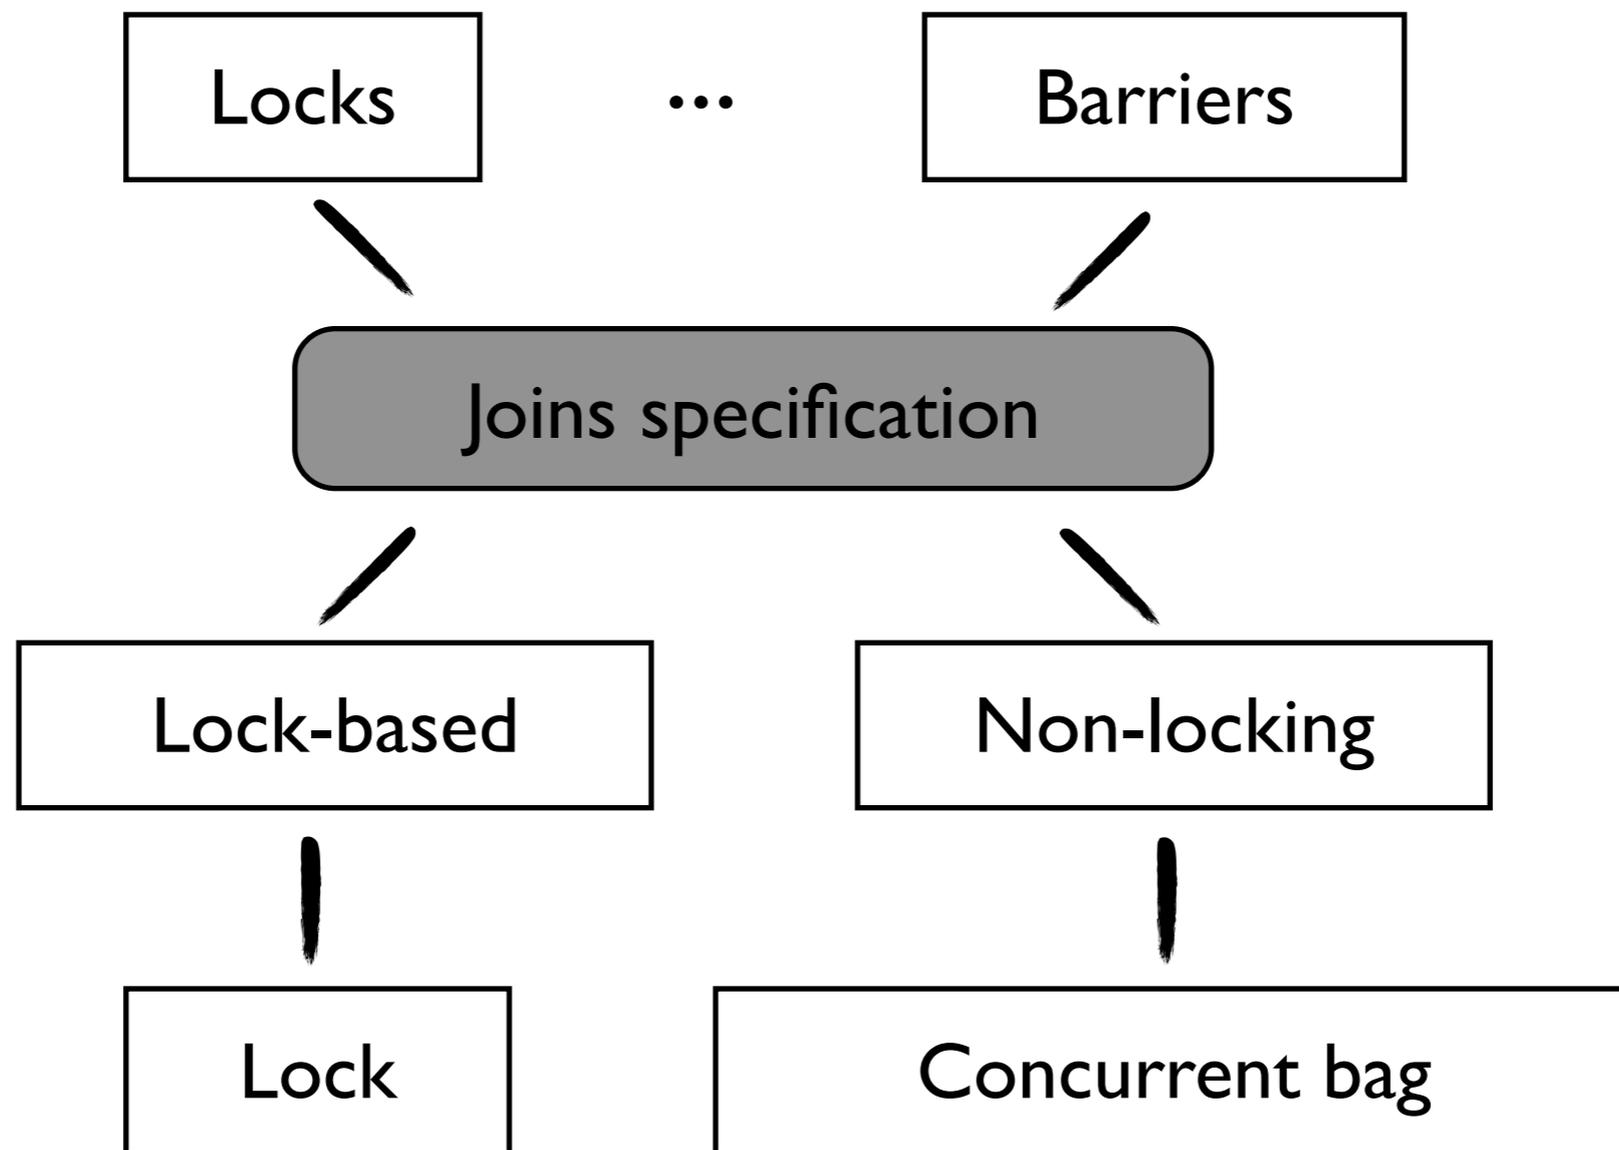
    join.When(acqR).And(unused).Do(() => { readers++; shared(); });
    join.When(acqR).And(shared).Do(() => { readers++; shared(); });
    join.When(acqW).And(unused).Do(() => { writer(); });
    join.When(relW).And(writer).Do(() => { unused(); });
    join.When(relR).And(shared).Do(() => {
      if (--readers == 0) unused() else shared(); });

    unused();
  }
}
```

reentrant continuation

state effect

Joins specification



Specification

- Requirements:
 - Ownership transfer
 - Stateful reentrant continuations
- Restrict attention to non-self-modifying clients

Ideas

- Let clients pick an ownership protocol for each channel
 - The channel pre-condition describes the resources the sender is required to transfer to the recipient upon sending a message
 - The channel post-condition describes the resources the recipient is required to transfer to the sender upon receiving the message
 - The channel post-condition of asynchronous channels must be emp
- Prove chords obey the ownership protocol, assuming channels obey the ownership protocol (to support reentrancy)

Specification

- Send a message on channel c (async or sync)

family of channel pre- and post-conditions, indexed by channels

transfer channel pre-condition from client to join instance

$$\{\text{join}(P, Q, j) * \text{chan}(c, j) * P(c)\}$$

$c()$

$$\{\text{join}(P, Q, j) * \text{chan}(c, j) * Q(c)\}$$

if c is an asynchronous channel, then channel post-condition must be emp

transfer channel post-condition from join instance to client

Specification

- Register a new chord with pattern p and continuation b

pattern p matches the multiset of channels X

$$\left\{ \begin{array}{l} \text{join}_{\text{init-pat}}(P, Q, j) * \text{pattern}(p, j, X) \\ * b \mapsto \left\{ \begin{array}{l} \underline{\otimes_{x \in X} P(x)} * \text{join}(P, Q, j) \\ \underline{\otimes_{x \in X} Q(x)} * \text{join}(P, Q, j) \end{array} \right\} \end{array} \right\}$$

$p.\text{Do}(b)$

$\{\text{join}_{\text{init-pat}}(P, Q, j)\}$

resources senders must transfer to recipient

resources recipient must transfer to senders

Specification

- Register a new chord with pattern p and continuation b

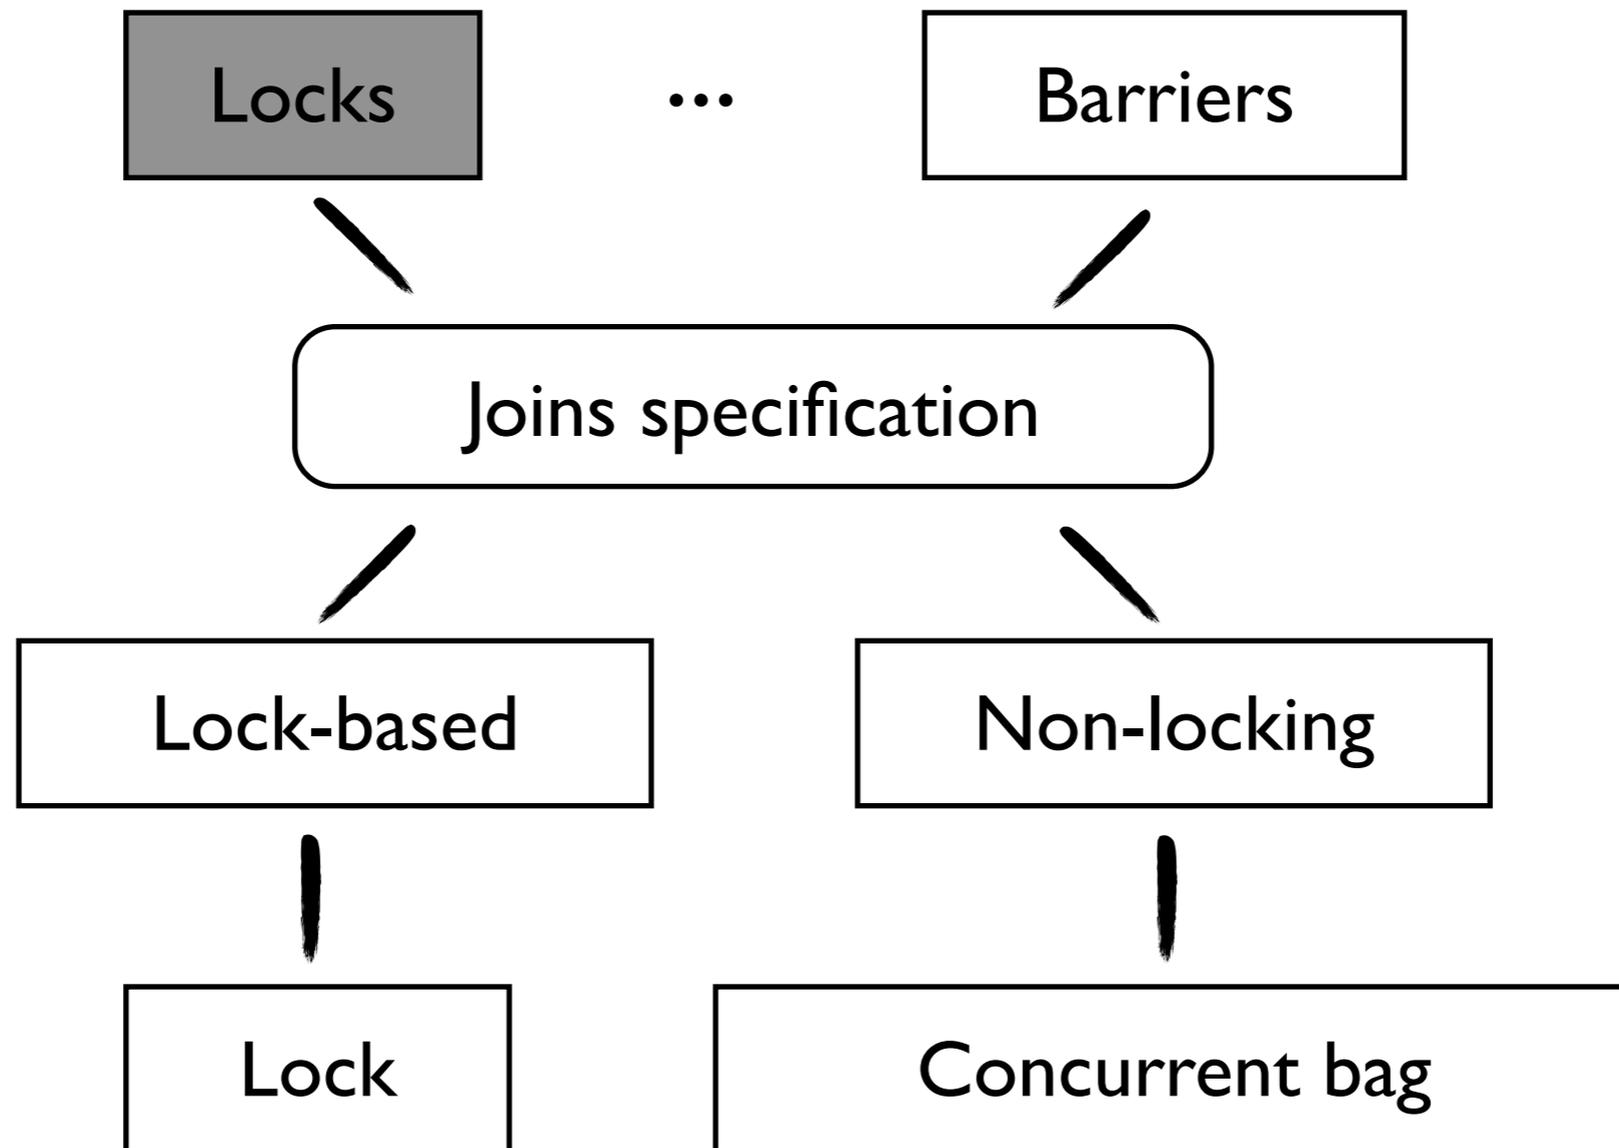
$$\left\{ \begin{array}{l} \text{join}_{\text{init-pat}}(P, Q, j) * \text{pattern}(p, j, X) \\ * b \mapsto \left\{ \begin{array}{l} \bigotimes_{x \in X} P(x) * \text{join}(P, Q, j) \\ \bigotimes_{x \in X} Q(x) * \text{join}(P, Q, j) \end{array} \right\} \end{array} \right\}$$

$p.\text{Do}(b)$

$\{\text{join}_{\text{init-pat}}(P, Q, j)\}$

the continuation is allowed to assume channels obey their ownership protocol

Verifying Clients



Reader/Writer lock

- Given resource invariants R and R_{ro} (picked by client) s.t.

$$\forall n \in \mathbf{N}. R(n) \Leftrightarrow R_{ro} * R(n + 1)$$

- R_{ro} : read permission to underlying resource
- $R(0)$: write permission to underlying resource
- $R(n)$: resource after splitting off n read permissions
- The reader/writer lock satisfies the following specification

$$\begin{array}{llll} \{emp\} & \text{acqR}() & \{R_{ro}\} & \{R_{ro}\} \text{relR}() \{emp\} \\ \{emp\} & \text{acqW}() & \{R(0)\} & \{R(0)\} \text{relW}() \{emp\} \end{array}$$

- Assign pre-conditions to asynchronous channels

$$P(\text{unused}) = \text{readers} \mapsto 0 * R(0)$$

$$P(\text{shared}) = \exists n \in \mathbb{N}. \text{readers} \mapsto n * R(n) * n > 0$$

$$P(\text{writer}) = \text{readers} \mapsto 0$$

- Assign pre- and post-conditions to synchronous channels

$$P(\text{acqR}) = \text{emp}$$

$$Q(\text{acqR}) = R_{ro}$$

$$P(\text{acqW}) = \text{emp}$$

$$Q(\text{acqW}) = R(0)$$

$$P(\text{relR}) = R_{ro}$$

$$Q(\text{relR}) = \text{emp}$$

$$P(\text{relW}) = R(0)$$

$$Q(\text{relW}) = \text{emp}$$

- Prove chords obey channel ownership protocol

```
class RWLock {  
    ...  
    public int readers = 0;  
  
    public RWLock() {  
        ...  
        join.When(acqR).And(UNUSED).Do(() => { readers++; shared(); });  
        ...  
    }  
}
```

$\{P(\text{acqR}) * P(\text{UNUSED}) * \text{join}(P, Q, j)\}$

readers++

shared();

$\{Q(\text{acqR}) * Q(\text{UNUSED}) * \text{join}(P, Q, j)\}$

- Prove chords obey channel ownership protocol

```

class RWLock {
  ...
  public int readers = 0;

  public RWLock() {
    ...
    join.When(acqR).And(UNUSED).Do(() => { readers++; shared(); });
    ...
  }
}

```

$\{\text{readers} \mapsto 0 * R(0) * \text{join}(P, Q, j)\}$

readers++

$\{\text{readers} \mapsto 1 * R(1) * R_{ro} * \text{join}(P, Q, j)\}$

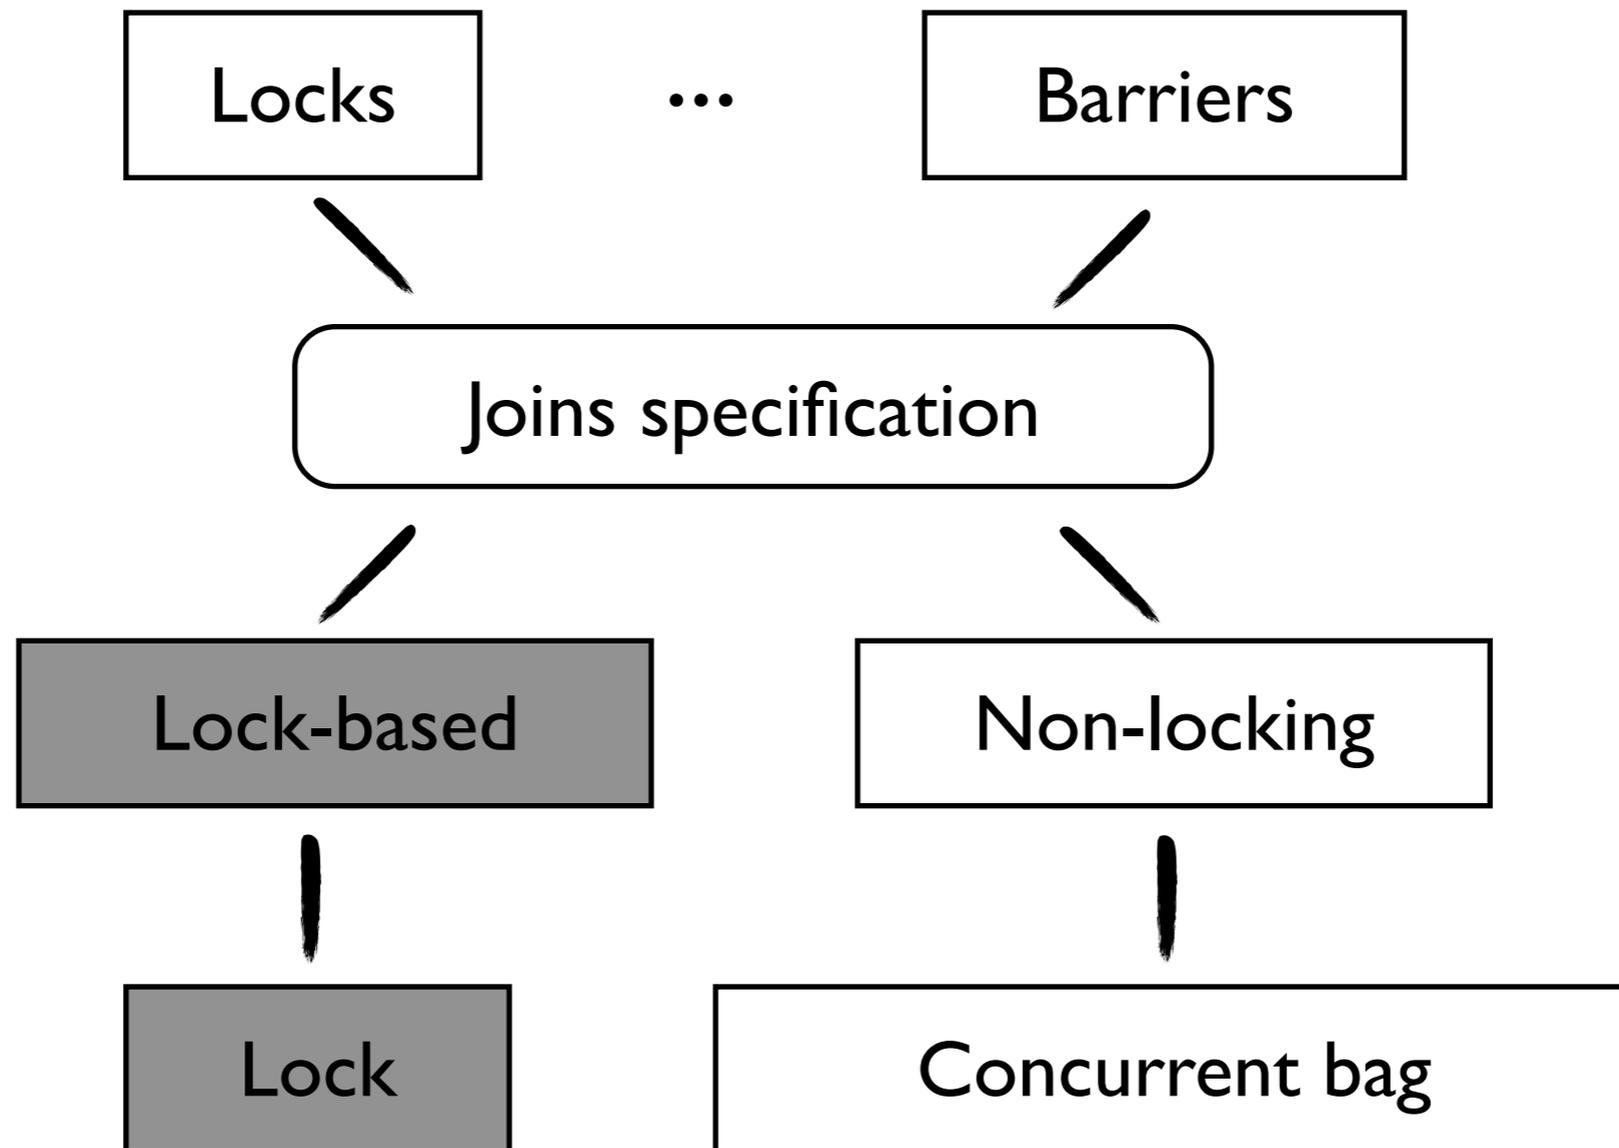
shared();

$\{R_{ro} * \text{join}(P, Q, j)\}$

$P(\text{shared}) = \exists n \in \mathbb{N}_+.$

$\text{readers} \mapsto n * R(n)$

Verifying an Implementation



Verifying an Implementation

- Challenges:
 - High-level join primitives implemented using shared mutable state
 - Definition of recursive representation predicates

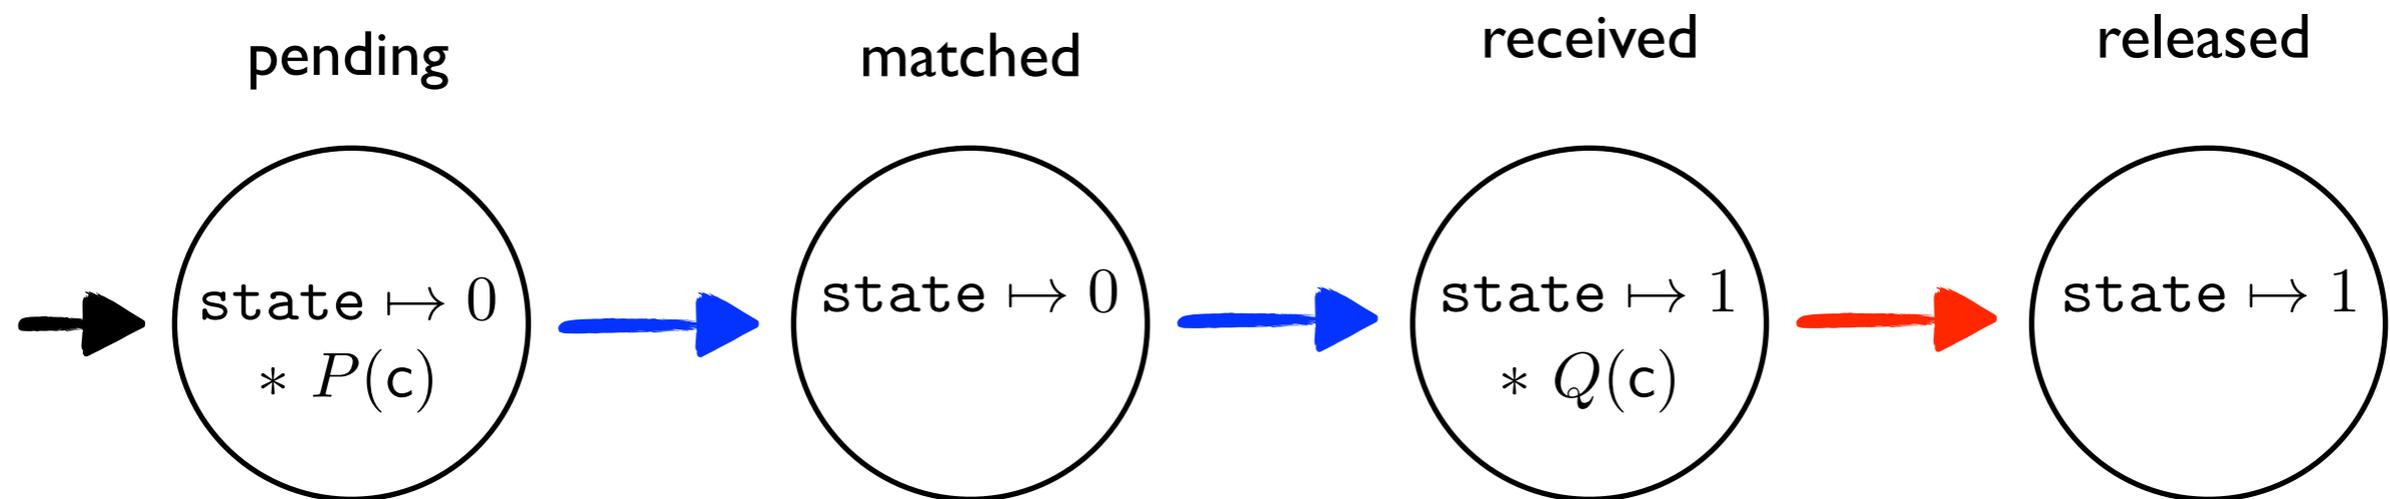
guarded recursion & step-indexed model

Messages

```
class Message {  
    public int state;  
  
    public Message() {  
        state = 0;  
    }  
  
    public void Receive() {  
        state = 1;  
    }  
}
```

Messages

- Assume channel pre- and post-conditions P and Q
- Imagine a message on channel c

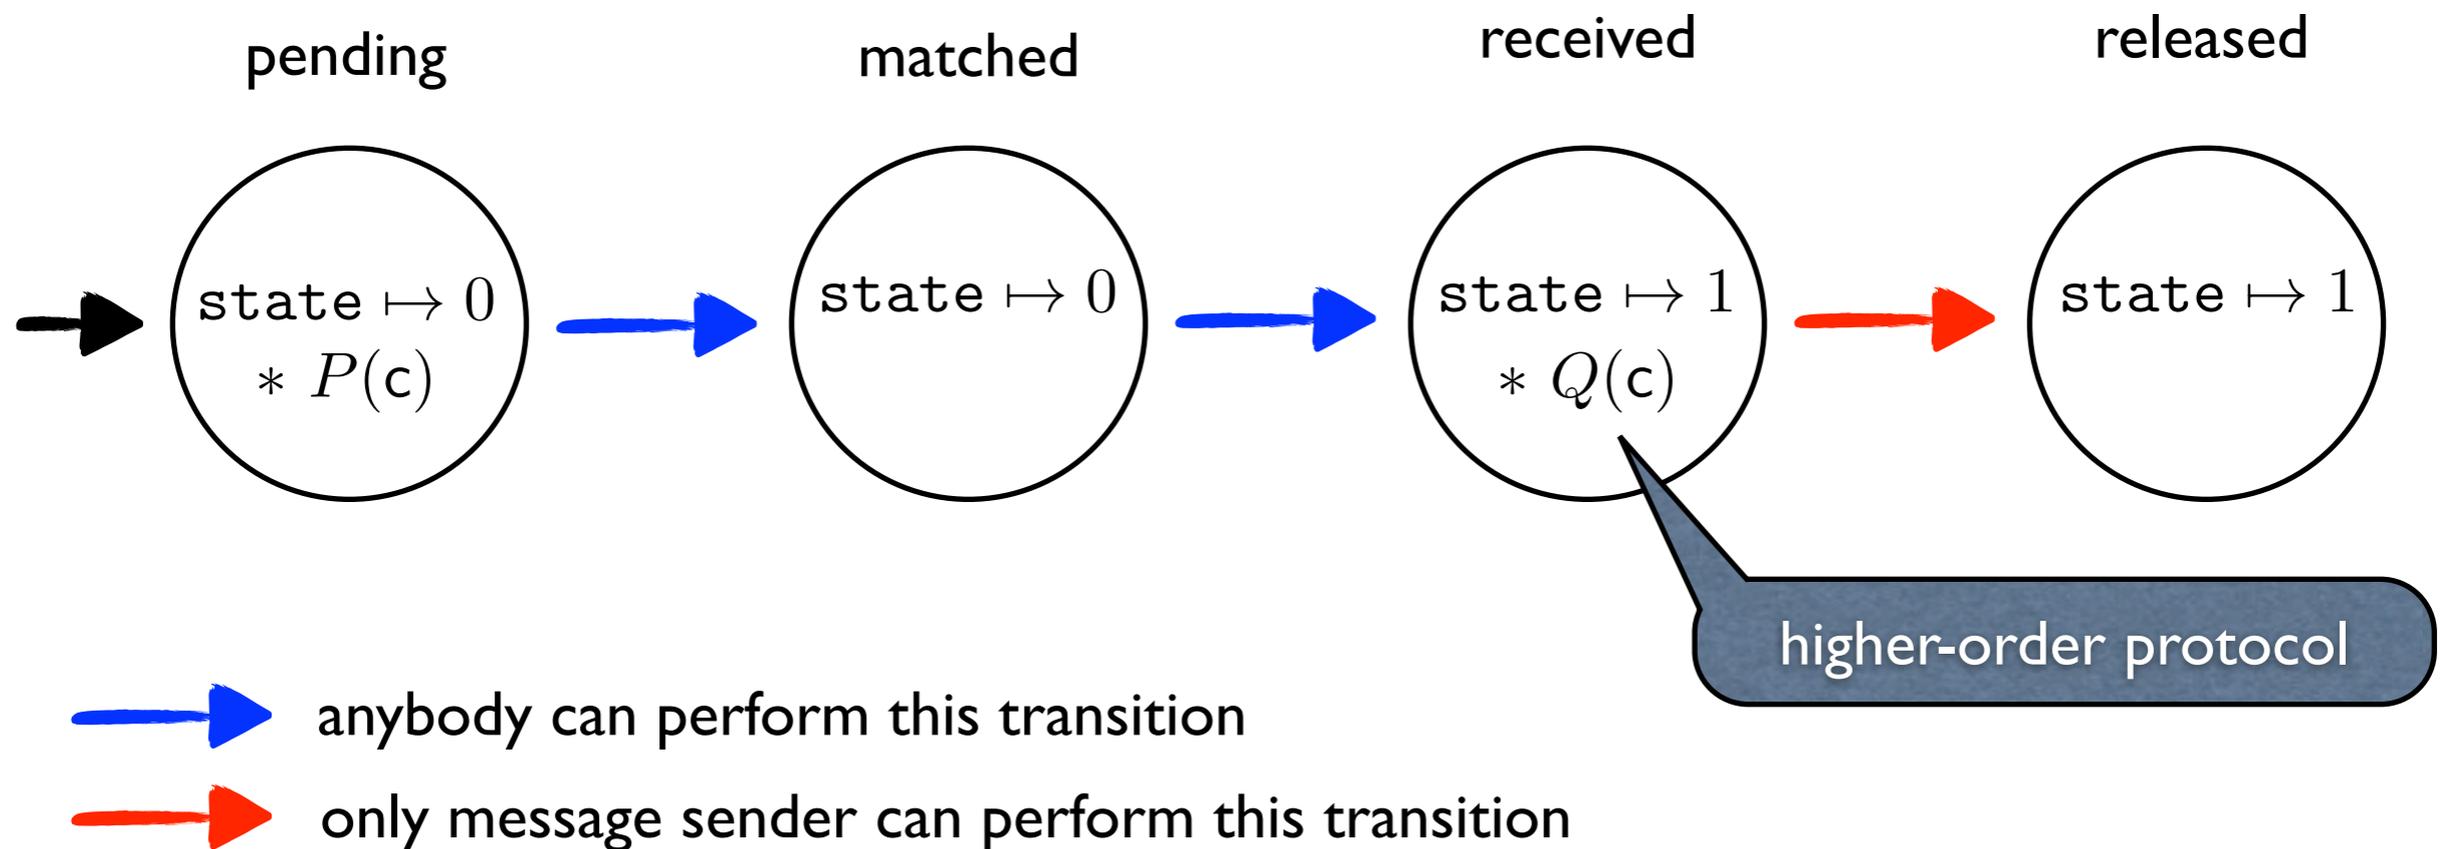


 anybody can perform this transition

 only message sender can perform this transition

Messages

- Use Concurrent Abstract Predicates [Dinsdale-Young et. al.] to impose this low-level protocol on messages



HOCAP

- Higher-order protocols are difficult; the previous proposal [Dodds et. al.] from POPL11 is unsound!
- We restrict attention to state-independent higher-order protocols. An assertion P is expressible using state-independent protocols (SIPs) iff

$\exists R, S : Prop. \text{ valid } (P \Leftrightarrow R * S) \wedge \text{ noprotocol}(R) \wedge \text{ nostate}(S)$

invariant under arbitrary
changes to protocols

invariant under arbitrary
changes to the state

- We require all channel pre- and post-conditions to be expressible using SIPs

Summary

- Verified the lock-based joins implementation against the high-level joins specification
- Verified a couple of classic synchronization primitives using the high-level joins specification
- Given a logic and model for HOCAP with support for state-independent higher-order protocols
- TRs available at www.itu.dk/~kasv

Questions?

Higher-order protocols in CAP

Let

$$P \stackrel{\text{def}}{=} (x \mapsto 0 * (\boxed{y \mapsto 0}_I^r \vee \boxed{y \mapsto 0}_J^r)) \vee \\ (x \mapsto 1 * \boxed{y \mapsto 0}_J^r)$$

where

$$I[\alpha] : y \mapsto 1 \rightsquigarrow y \mapsto 2$$

$$J[\alpha] : y \mapsto 1 \rightsquigarrow y \mapsto 3$$

$$K[\alpha] : P \rightsquigarrow P$$

then P is stable, but $\boxed{P}_K^{r'}$ is not