



# Programming and Proving with Distributed Protocols

Disel: Distributed Separation Logic



$$\vdash \{P\} c \{Q\}$$



<http://distributedcomponents.net>

Ilya Sergey

**James R. Wilcox**

Zach Tatlock



PAUL G. ALLEN SCHOOL  
OF COMPUTER SCIENCE & ENGINEERING

# Distributed Systems



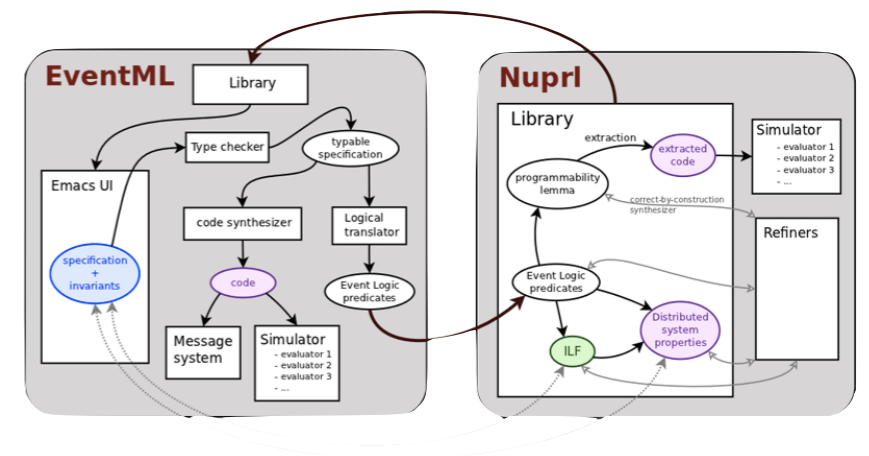

# Distributed *Infrastructure*



# Distributed *Applications*




# Verified Distributed Systems





# Verified Distributed *Infrastructure*

$\text{holds}(\Phi, S, \rightsquigarrow_1) \rightarrow$   
 $\text{holds}(\text{transfer}(\Phi), T(S), \rightsquigarrow_2)$




# Verified Distributed *Infrastructure*

$\text{holds}(\Phi, S, \rightsquigarrow_1) \rightarrow$   
 $\text{holds}(\text{transfer}(\Phi), T(S), \rightsquigarrow_2)$



# Verified Distributed *Applications*


$\text{holds}(\Phi, S, \rightsquigarrow_1) \rightarrow$   
 $\text{holds}(\text{transfer}(\Phi), T(S), \rightsquigarrow_2)$





# Verified Distributed *Applications*


$\text{holds}(\Phi, S, \rightsquigarrow_1) \rightarrow$   
 $\text{holds}(\text{transfer}(\Phi), T(S), \rightsquigarrow_2)$



Veri-



Wow



Iron



-Cert



# Verified Distributed *Applications*

Challenging to verify apps in terms of infra.  
*starting from scratch is unacceptable*

Indicates deeper problems with composition  
*one node's client is another's server!*

A yellow starburst graphic containing the word "Iron".

Iron

A brown square logo with the letters "HP" in white.

-Cert

# Challenges

Client reasoning

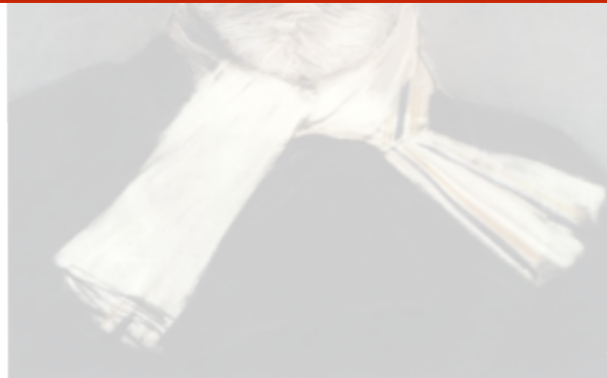
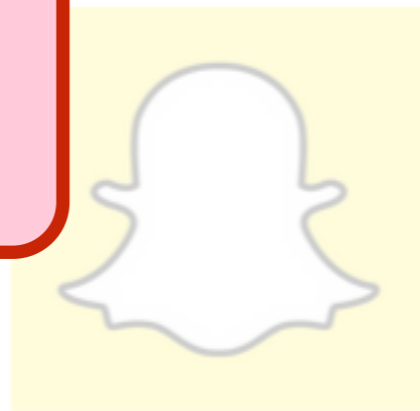
Invariants

Separation

uted *Applications*



Wow



Iron



-Cert

## Challenges

Client reasoning

Invariants

Separation

## Solutions

Protocols

WITHINV rule


FRAME rule/Hooks

Disel:



$\vdash \{P\} c \{Q\}$

# Outline



Protocols and running example

Logical mechanisms



*programming with protocols*

*invariants*


*framing and hooks*

Implementation and future work


# Cloud Compute



# Cloud Compute




# Cloud Compute





# Cloud Compute



# Cloud Compute: Server

```
while true:
```

```
    (from, n) <- recv Req
```

```
    send Resp(n, factors(n)) to from
```

# Cloud Compute: Server

```
while true:
```

```
    (from, n) <- recv Req
```


```
    send Resp(n, factors(n)) to from
```

Traditional specification:


messages from server have correct factors

Proved by finding an invariant of the system

# Cloud Compute: Server



# Cloud Compute: Client



# Cloud Compute: Client

```
send Req(21) to server  
(_, ans) <- recv Resp  
assert ans == {3, 7}
```

# Cloud Compute: Client

```
send Req(21) to server  
(_, ans) <- recv Resp  
assert ans == {3, 7}
```

Start over with clients in system?

# Cloud Compute: Client


```
send Req(21) to server  
(_, ans) <- recv Resp  
assert ans == {3, 7}
```

Start over with clients in system?


In Diesel: use protocol to describe client interface




# Protocols



# Protocols




# Protocols



A protocol is an **interface** among nodes


# Protocols



A protocol is an **interface** among nodes


Enables compositional verification

# Cloud Compute Protocol



Messages:


# Cloud Compute Protocol



Messages:

State:

# Cloud Compute Protocol




Messages:

State:

Transitions:

# Cloud Compute Protocol



Messages:


State:

Transitions:

Sends: precondition and effect



# Cloud Compute Protocol



Messages:


State:

Transitions:

Sends: precondition and effect

Receives: effect

# Cloud Compute Protocol



Messages:


State:

Transitions:

Sends:

Receives:

# Cloud Compute Protocol



Messages:  $\text{Req}(n)$  |  $\text{Resp}(n, s)$


State:

Transitions:

Sends:

Receives:

# Cloud Compute Protocol



Messages:  $\text{Req}(n) \mid \text{Resp}(n, s)$


State:  $\text{outstanding} : \text{Set}\langle \text{Msg} \rangle$

Transitions:

Sends:

Receives:

# Cloud Compute Protocol



Messages:  $\text{Req}(n) \mid \text{Resp}(n, s)$

State:  $\text{outstanding} : \text{Set}\langle \text{Msg} \rangle$

Transitions:

Sends:

Req


Resp

Receives:

Req

Resp

# Cloud Compute




Send Req( $n$ )

Precondition: none

Effect: none


# Cloud Compute



Receive Req( $n$ )

Effect:            add (from,  $n$ ) to out

# Cloud Compute




Send  $\text{Resp}(n, l)$

Requires:  $l == \text{factors}(n)$   
 $(n, to)$  in out

Effect: removes  $(n, to)$  from out




# Cloud Compute



Recv Resp( $n, l$ )

Effect: none

# Cloud Compute Protocol



Messages:  $\text{Req}(n) \mid \text{Resp}(n, s)$

State:  $\text{outstanding} : \text{Set}\langle \text{Msg} \rangle$

Transitions:

Sends:

Req


Resp

Receives:

Req

Resp

# Outline



Protocols and running example

Logical mechanisms

*programming with protocols*

*invariants*

*framing and hooks*

Implementation and future work

# Cloud Compute: Server

```
while true:
```

```
    (from, n) <- recv Req
```

```
    send Resp(n, factors(n)) to from
```

# Cloud Compute: Server

```
while true:
```


```
    (from, n) <- recv Req
```

```
    send Resp(n, factors(n)) to from
```

Precondition on **send** requires correct factors

# Cloud Compute: Server

---



$\vdash \{ \quad \} \mathbf{send} \ m \ \mathbf{to} \ h \ \{ \quad \}$

while true:

    (from, n) <- **recv** Req

**send** Resp(n, factors(n)) **to** from

Precondition on **send** requires correct factors

# Cloud Compute: Server

---

  $\vdash \{ \quad \} \mathbf{send}_{\boxed{t}} m \mathbf{to} h \{ \quad \}$


while true:

(from, n)  $\leftarrow$  **recv** Req

**send** Resp(n, factors(n)) **to** from

Precondition on **send** requires correct factors

# Cloud Compute: Server



while true:


(from, n) <- **recv** Req

**send** Resp(n, factors(n)) **to** from

Precondition on **send** requires correct factors



# Cloud Compute: Server




while true:

(from, n)  $\leftarrow$  **rcv** Req

**send** Resp(n, factors(n)) **to** from

Precondition on **send** requires correct factors

# Cloud Compute: Server



---

  $\vdash \{Pre_t\} \mathbf{send}_t m \mathbf{to} h \{sent_t(m, h)\}$

while true:

(from, n)  $\leftarrow$  **rcv** Req

**send** Resp(n, factors(n)) **to** from

Precondition on **send** requires correct factors

# Cloud Compute: Client

```
send Req(21) to server  
(_, ans) <- recv Resp  
assert ans == {3, 7}
```

# Cloud Compute: Client

```
send Req(21) to server  
(_, ans) <- recv Resp  
assert ans == {3, 7}
```


**recv** doesn't ensure correct factors

# Cloud Compute: Client

```
send Req(21) to server  
(_, ans) <- recv Resp  
assert ans == {3, 7}
```

**recv** doesn't ensure correct factors

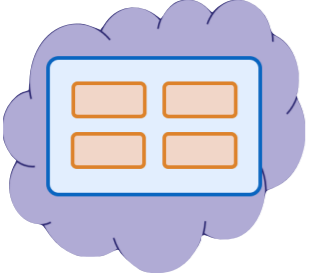
# Cloud Compute: Client



```
send Req(21) to server  
(_, ans) <- recv Resp  
assert ans == {3, 7}
```

**recv** doesn't ensure correct factors

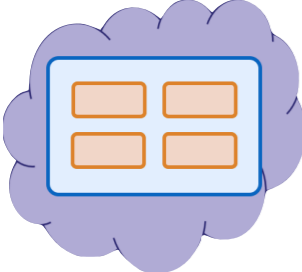
# Protocol Invariants


$$\vdash \{P\} c \{Q\} \quad I \text{ inductive}$$

---


$$\vdash \{P \wedge I\} c \{Q \wedge I\}$$

# Protocol Invariants


$$\vdash \{P\} c \{Q\} \quad I \text{ inductive}$$


---


$$\vdash \{P \wedge I\} c \{Q \wedge I\}$$

Protocol where every state satisfies  $I$



# Cloud Compute: Client



```
send Req(21) to server
(_, ans) <- recv Resp
assert ans == {3, 7}
```

Now **recv** ensures correct factors

# Cloud Compute: More Clients

```
send Req(21) to server1  
send Req(35) to server2  
(_, ans1) <- rcv Resp  
(_, ans2) <- rcv Resp  
assert ans1 ∪ ans2 == {3, 5, 7}
```

# Cloud Compute: More Clients

```
send Req(21) to server1  
send Req(35) to server2  
(_, ans1) <- rcv Resp  
(_, ans2) <- rcv Resp  
assert ans1 ∪ ans2 == {3, 5, 7}
```

Same protocol enables verification

# Frame rule

$$\text{[Icon]} \vdash \{P\} c \{Q\} \quad R \text{ stable}$$

---

$$\text{[Icon]} \vdash \{P * R\} c \{Q * R\}$$

# Frame rule

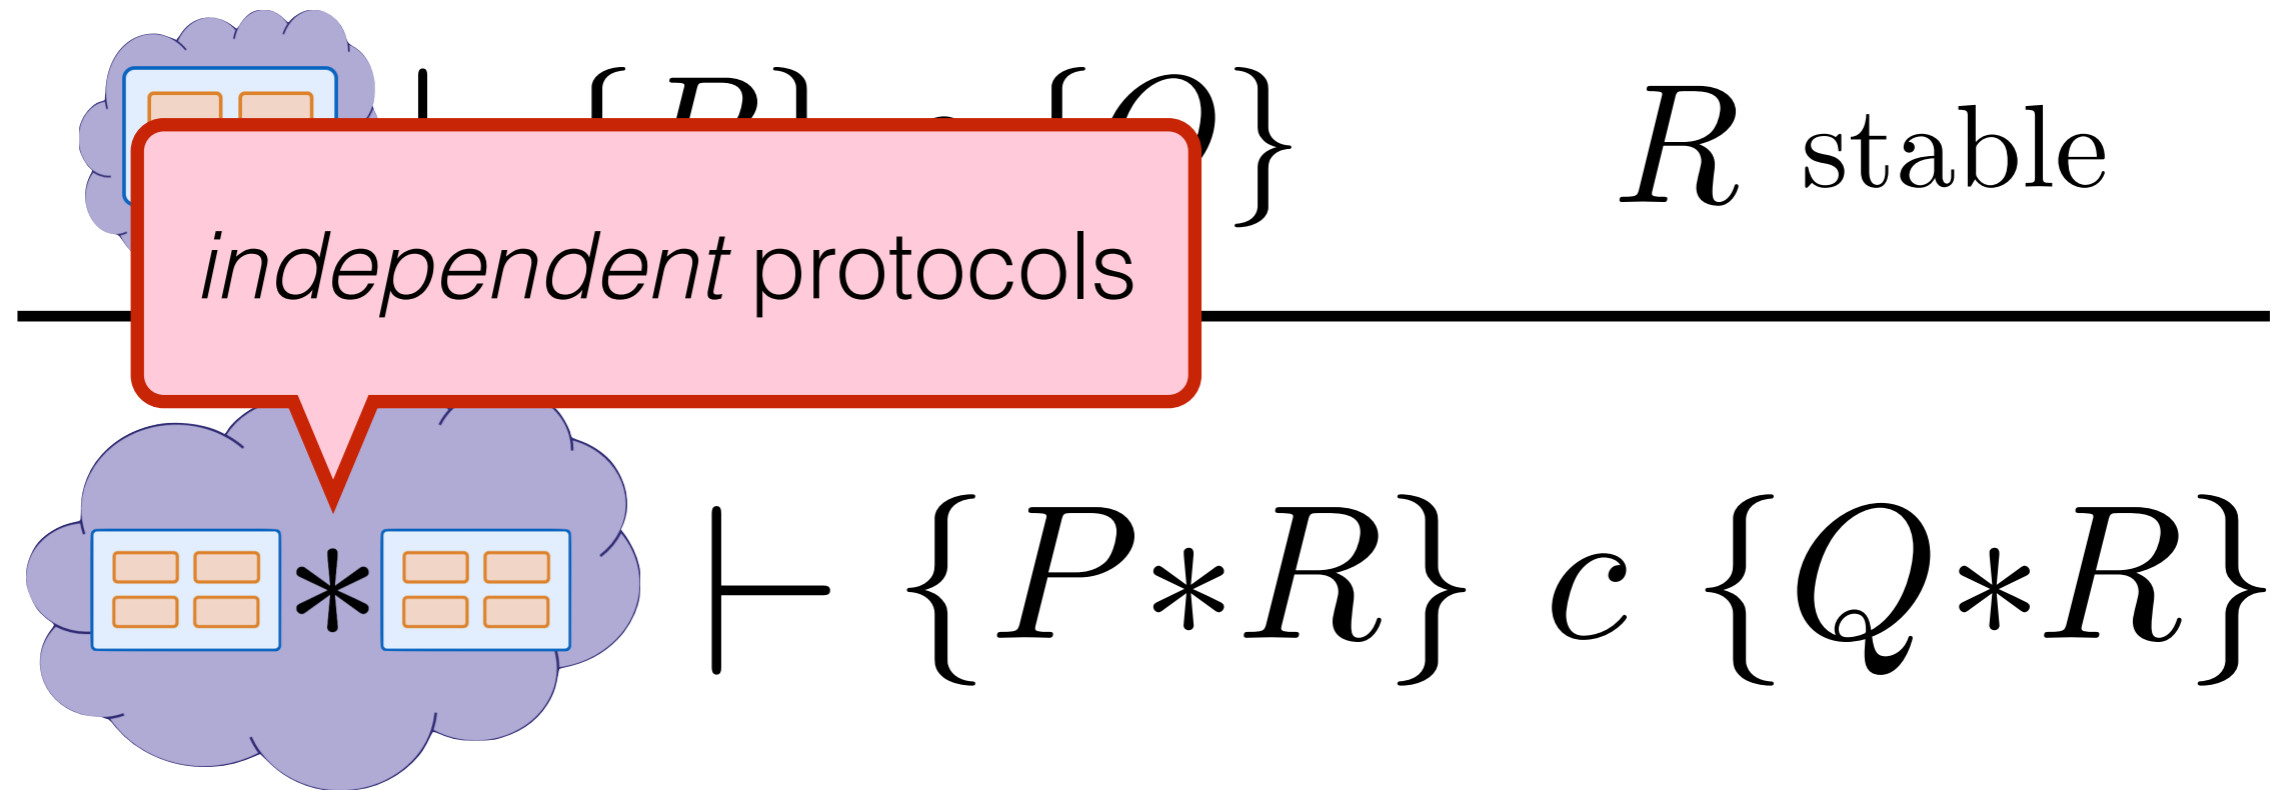
$$\text{[Component Icon]} \vdash \{P\} c \{Q\} \quad R \text{ stable}$$

---

$$\text{[Component Icon]} * \text{[Component Icon]} \vdash \{P * R\} c \{Q * R\}$$

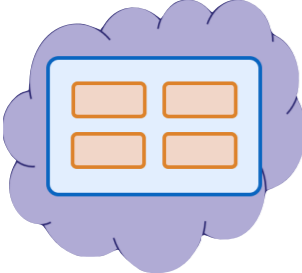
*Reuse invariants from component protocols*

# Frame rule

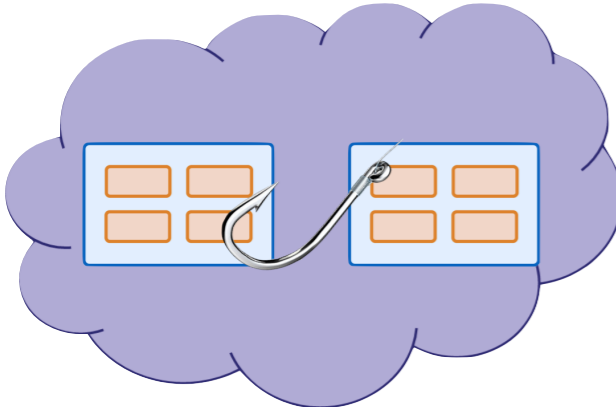


*Reuse invariants from component protocols*

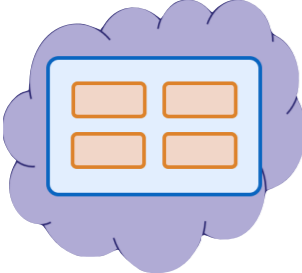
# Frame rule: Hooks


$$\vdash \{P\} \text{ c } \{Q\} \quad R \text{ stable}$$

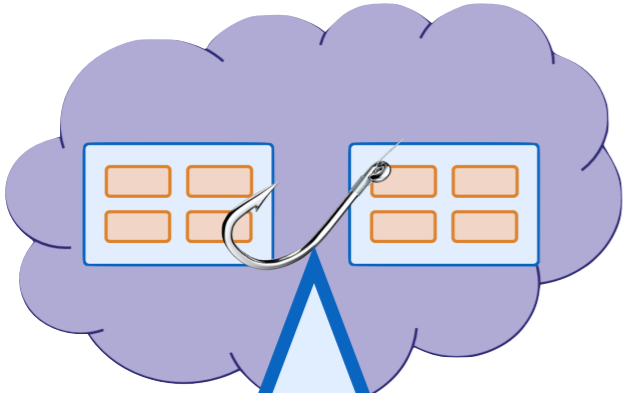
---


$$\vdash \{P * R\} \text{ c } \{Q * R\}$$

# Frame rule: Hooks


$$\vdash \{P\} c \{Q\} \quad R \text{ stable}$$


---


$$\vdash \{P * R\} c \{Q * R\}$$

Allows one protocol to restrict another



# Outline



Protocols and running example

Logical mechanisms

*programming with protocols*

*invariants*

*framing and hooks*

Implementation and future work

# Implementation

Shallowly embedded in Coq

*with full power of functional programming*



Executable via extraction to OCaml

*via trusted shim to implement semantics*

Case study: two-phase commit

*exercises all features of the logic*

# Related and Future Work




Concurrent separation logics

*Iris, FCSL, CAP, ...*


Adding other effects

*e.g. mutable heap, threads, failure...*

Composition: A way  
to make proofs harder



# Composition: A way to make proofs harder



“In 1997, the unfortunate reality is that engineers rarely specify and reason formally about the systems they build. It seems unlikely that reasoning about the composition of open-system specifications will be a practical concern within the next 15 years.”

## Challenges

Client reasoning

Invariants

Separation

## Solutions

Protocols

WITHINV rule

FRAME rule/Hooks

Disel:



$\vdash \{P\} c \{Q\}$