

# Aneris: A Mechanised Logic for Modular Reasoning about Distributed Systems

Morten Krogh-Jespersen<sup>1</sup>    Amin Timany<sup>2</sup>  
Marit Edna Ohlenbusch<sup>1</sup>    Simon Oddershede Gregersen<sup>1</sup>  
Lars Birkedal<sup>1</sup>

<sup>1</sup>Aarhus University, Aarhus, Denmark

<sup>2</sup>imec-DistriNet KU Leuven, Leuven, Belgium

October 28, 2019

Iris workshop,  
Aarhus University

# Introduction

Distributed systems are ubiquitous

Some applications are critical, e.g., online banking

Hence, there is a need for verification

Efficient implementations often use advanced features like **node-local concurrency** and **higher-order memory**

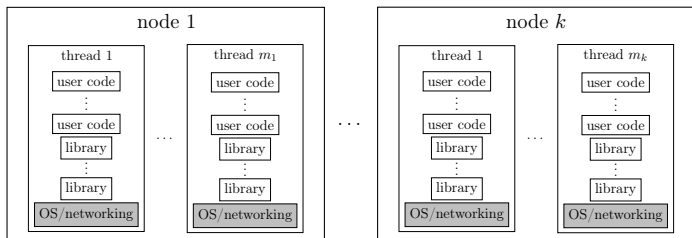
It is well known that reasoning about distributed programs is **difficult**

Traditionally, most works focus on verifying **a high-level model** of the system

We introduce **Aneris**: a program logic for **modular** verification of **safety** of distributed systems' **code**

# In this talk

- ▶ **AnerisLang**: an advanced programming language for **programming distributed systems**
- ▶ **Aneris logic**: a program logic for **modular** reasoning about AnerisLang programs
  - ▶ **Horizontal modularity**: nodes, and threads, are verified in isolation and the proofs are composed
  - ▶ **Vertical modularity**: library code is verified separately and library clients are verified against the library specs.



## In this talk

- ▶ **AnerisLang**: an advanced programming language for programming distributed systems
- ▶ **Aneris logic**: a program logic for modular reasoning about AnerisLang programs
  - ▶ **Horizontal modularity**: nodes, and threads, are verified in isolation and the proofs are composed
  - ▶ **Vertical modularity**: library code is verified separately and library clients are verified against the library specs.
- ▶ Concurrent-separation-logic-style specs for distributed systems
- ▶ Examples of distributed network specifications

In this work, we introduce **AnerisLang** and the **Aneris** logic.

**AnerisLang**: an untyped ML-style programming language with

- ▶ UDP-like network primitives
  - ▶ unordered messages
  - ▶ possibility of dropped packets
  - ▶ duplicate protection
- ▶ Concurrency (multiple threads on each node)
- ▶ Higher-order memory (can store code in memory)
- ▶ Primitive types and support of (de)serialization to strings
- ▶ Can (almost) be directly extracted to running OCaml code

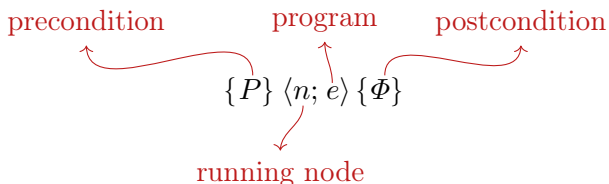
The operational semantics keeps track of

- ▶ A heap for each node
- ▶ A mapping from socket handles to socket addresses (ip and port) for each node
- ▶ A message soup: a collection of sent messages

# Aneris logic

Aneris logic: a program logic based on Iris for distributed systems

Hoare triples:



Hoare triples guarantee safety:

If  $\{P\} \langle n; e \rangle \{\Phi\}$  holds then

- given that  $P$  holds for the initial state,
- $e$  is **safe**, *i.e.* the program (distributed system) will not crash,
- and if it terminates with a final value  $v$ ,  $\Phi(v)$  holds.

# Modularity in Aneris logic

- ▶ Thread-local reasoning:

$$\frac{\{P_1\} \langle n; e_1 \rangle \{v.Q_1\} \quad \{P_2\} \langle n; e_2 \rangle \{v.Q_2\}}{\{P_1 * P_2\} \langle n; e_1 \parallel e_2 \rangle \{v.\exists v_1, v_2.v = (v_1, v_2) * Q_1[v_1/v] * Q_2[v_2/v]\}}$$

- ▶ Node-local reasoning:

$$\frac{\{P_1 * \text{IsNode}(n_1) * \text{FreePorts}(ip_1, \mathfrak{P})\} \langle n_1; e_1 \rangle \{\text{True}\} \quad \{P_2 * \text{IsNode}(n_2) * \text{FreePorts}(ip_2, \mathfrak{P})\} \langle n_2; e_2 \rangle \{\text{True}\}}{\{P_1 * P_2 * \text{Freelp}(ip_1) * \text{Freelp}(ip_2)\} \langle \mathfrak{S}; (n_1; ip_1; e_1) \parallel (n_2; ip_2; e_2) \rangle \{\text{True}\}}$$

- ▶ Reasoning about network communications (socket protocols)

$$\Phi : \text{Message} \rightarrow iProp \quad a \Vdash \Phi$$

- ▶ if  $a \Vdash \Phi$  and  $a \Vdash \Psi$  then  $\Phi$  and  $\Psi$  are equivalent.
- ▶  $a \Vdash \Phi \dashv\vdash a \Vdash \Phi * a \Vdash \Phi$ .

# Modularity in Aneris logic

- ▶ Thread-local reasoning:

$$\frac{\{P_1\} \langle n; e_1 \rangle \{v.Q_1\} \quad \{P_2\} \langle n; e_2 \rangle \{v.Q_2\}}{\{P_1 * P_2\} \langle n; e_1 \parallel e_2 \rangle \{v.\exists v_1, v_2.v = (v_1, v_2) * Q_1[v_1/v] * Q_2[v_2/v]\}}$$

- ▶ Node-local reasoning:

$$\frac{\{P_1 * \text{IsNode}(n_1) * \text{FreePorts}(ip_1, \mathfrak{P})\} \langle n_1; e_1 \rangle \{\text{True}\} \quad \{P_2 * \text{IsNode}(n_2) * \text{FreePorts}(ip_2, \mathfrak{P})\} \langle n_2; e_2 \rangle \{\text{True}\}}{\{P_1 * P_2 * \text{Freelp}(ip_1) * \text{Freelp}(ip_2)\} \langle \mathfrak{S}; (n_1; ip_1; e_1) \parallel (n_2; ip_2; e_2) \rangle \{\text{True}\}}$$

- ▶ Reasoning about network communications (socket protocols)

$$\Phi : \text{Message} \rightarrow iProp \quad a \Vdash \Phi$$

- ▶ if  $a \Vdash \Phi$  and  $a \Vdash \Psi$  then  $\Phi$  and  $\Psi$  are equivalent.
- ▶  $a \Vdash \Phi \dashv\vdash a \Vdash \Phi * a \Vdash \Phi$ .



# Modularity in Aneris logic

- ▶ Thread-local reasoning:

$$\frac{\{P_1\} \langle n; e_1 \rangle \{v.Q_1\} \quad \{P_2\} \langle n; e_2 \rangle \{v.Q_2\}}{\{P_1 * P_2\} \langle n; e_1 \parallel e_2 \rangle \{v.\exists v_1, v_2.v = (v_1, v_2) * Q_1[v_1/v] * Q_2[v_2/v]\}}$$

- ▶ Node-local reasoning:

$$\frac{\{P_1 * \text{IsNode}(n_1) * \text{FreePorts}(ip_1, \mathfrak{P})\} \langle n_1; e_1 \rangle \{\text{True}\} \quad \{P_2 * \text{IsNode}(n_2) * \text{FreePorts}(ip_2, \mathfrak{P})\} \langle n_2; e_2 \rangle \{\text{True}\}}{\{P_1 * P_2 * \text{Freelp}(ip_1) * \text{Freelp}(ip_2)\} \langle \mathfrak{S}; (n_1; ip_1; e_1) \parallel (n_2; ip_2; e_2) \rangle \{\text{True}\}}$$

- ▶ Reasoning about network communications (socket protocols)

$$\Phi : \text{Message} \rightarrow iProp \quad a \Vdash \Phi$$

- ▶ if  $a \Vdash \Phi$  and  $a \Vdash \Psi$  then  $\Phi$  and  $\Psi$  are equivalent.
- ▶  $a \Vdash \Phi \dashv\vdash a \Vdash \Phi * a \Vdash \Phi$ .

## Specifying socket protocols

$$\Phi : Message \rightarrow iProp \quad a \Vdash \Phi$$

Given

- a predicate  $P : String \rightarrow iProp$  for message contents,
- and a predicate  $Q : Message \rightarrow iProp$

we specify a socket protocol  $\Phi$  as follows:

$$\Phi(m) \triangleq \exists \Psi. \text{from}(m) \Vdash \Psi * P(\text{body}(m)) * (\forall m'. Q(m') \multimap \Psi(m'))$$

Socket protocols restrict messages and, if necessary, the protocol of the sender!

This is possible because of Iris's **impredicativity**:

given any  $\Phi : Message \rightarrow iProp$ ,  $a \Vdash \Phi : iProp$

## Concurrent-Separation-logic-style specifications

Our specifications for distributed systems are inspired by concurrent separation logic

This style of specification lends itself quite well to modular reasoning

To illustrate this point we see the specs for a distributed lock server

## CSL specs for a lock

Lock specifications:

$\exists \text{isLock}.$

$\wedge \quad \forall v, K. \text{isLock}(v, K) \dashv\vdash \text{isLock}(v, K) * \text{isLock}(v, K)$

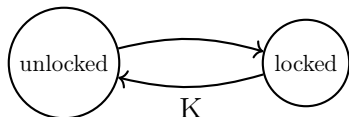
$\wedge \quad \forall v, K. \text{isLock}(v, K) \vdash K * K \Rightarrow \text{False}$

$\wedge \quad \{\text{True}\} \text{newLock } () \{v. \exists K. \text{isLock}(v, K)\}$

$\wedge \quad \forall v. \{\text{isLock}(v, K)\} \text{acquire } v \{v.K\}$

$\wedge \quad \forall v. \{\text{isLock}(v, K) * K\} \text{release } v \{\text{True}\}$

Intuitively we think of the following state transition system:



## Aneris specs for a distributed lock

$$\Phi_{lock}(m) \triangleq \exists \Psi. \text{from}(m) \Rightarrow \Psi * (\text{acq}(m, \Psi) \vee \text{rel}(m, \Psi))$$

$$\begin{aligned} \text{acq}(m, \Psi) \triangleq & (\text{body}(m) = \text{"LOCK"}) * \\ & \forall m'. (\text{body}(m') = \text{"NO"}) \vee (\text{body}(m') = \text{"YES"} * K) \\ & -* \Psi(m') \end{aligned}$$

$$\begin{aligned} \text{rel}(m, \Psi) \triangleq & (\text{body}(m) = \text{"RELEASE"}) * K * \\ & \forall m'. (\text{body}(m') = \text{"RELEASED"}) -* \Psi(m') \end{aligned}$$

# Sockets and binding

- Socket creation:

SOCKET  
{IsNode( $n$ )}  $\langle n; \text{socket } () \rangle \{z. z \hookrightarrow_n \text{None}\}$

- Binding to a static (fixed/primordial) address

SOCKETBIND-STATIC  
{Fixed( $A$ ) \*  $a \in A$  \* FreePort( $a$ ) \*  $z \hookrightarrow_n \text{None}\}$   
 $\langle n; \text{socketbind } z a \rangle$   
{ $x. x = 0$  \*  $z \hookrightarrow_n \text{Some } a$ }

- Binding to a dynamic address

SOCKETBIND-DYNAMIC  
{Fixed( $A$ ) \*  $a \notin A$  \* FreePort( $a$ ) \*  $z \hookrightarrow_n \text{None}\}$   
 $\langle n; \text{socketbind } z a \rangle$   
{ $x. x = 0$  \*  $z \hookrightarrow_n \text{Some } a * a \mapsto \Phi$ }

# Sending and receiving over sockets

- Sending over a socket

**SENDTO**  
 $\{z \hookrightarrow_n \text{Some } from * to \Rightarrow \Phi * \Phi((from, to, msg, \text{SENT}))\}$   
 $\langle n; \text{sendto } z \text{ msg } to \rangle$   
 $\{x. x = |msg| * z \hookrightarrow_n \text{Some } from\}$

- Receiving from a socket

**RECEIVEFROM**  
 $\{z \hookrightarrow_n \text{Some } to * to \Rightarrow \Phi\}$   
 $\langle n; \text{receivefrom } z \rangle$   
 $\left\{ \begin{array}{l} x. z \hookrightarrow_n \text{Some } to * \\ (x = \text{None} \vee ( \exists m. x = \text{Some } (\text{body}(m), \text{from}(m)) * \Phi(m) * \text{R}(m) )) \end{array} \right\}$

# Adequacy

Let

- $\mathcal{P} \subseteq Ip$ ,
- and  $A \subseteq Address$ .

Given a primordial socket protocol  $\Phi_a$  for each  $a \in A$ , suppose that the Hoare triple

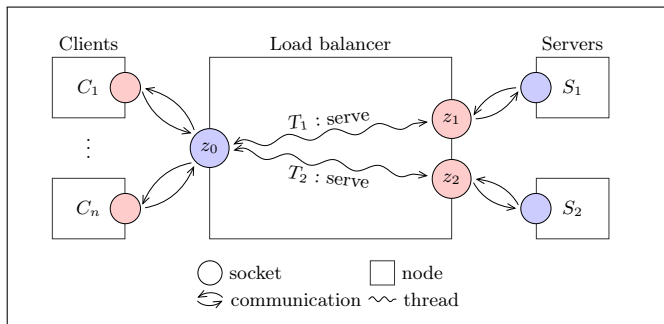
$$\{\text{Fixed}(A) * \bigstar_{a \in A} a \Rightarrow \Phi_a * \bigstar_{i \in \mathcal{P}} \text{Freelp}(i)\} \langle n; e \rangle \{v.\text{True}\}$$

is derivable in Aneris.

Then,  $e$  is **safe**, *i.e.* it will not crash.

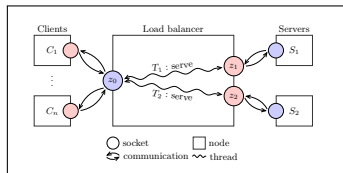


## Example: load balancer



$$\Phi_{rel}(P_{val}, P_{in}, P_{out})(m) \triangleq \exists \Psi, v. \text{from}(m) \Rightarrow \Psi * P_{in}(m, v) * P_{val}(v) * (\forall m'. P_{val}(v) * P_{out}(m', v) \multimap \Psi(m'))$$

# Example: load balancer



$$\Phi_{rel}(P_{val}, P_{in}, P_{out})(m) \triangleq \exists \Psi, v. \text{from}(m) \Rightarrow \Psi * P_{in}(m, v) * P_{val}(v) * \\ (\forall m'. P_{val}(v) * P_{out}(m', v) \rightarrow * \Psi(m'))$$

$$\left\{ \text{Static}((ip, p), A, \phi_{rel}(\lambda \_ . \text{True}, P_{in}, P_{out})) * \text{IsNode}(n) * \right. \\ \left. \left( \begin{array}{l} * \\ \text{Dynamic}((ip, p'), A) \end{array} \right) * \right. \\ \left. \left( \begin{array}{l} * \\ \exists v. \text{LB}(1, s, v) * s \Rightarrow \phi_{rel}(\lambda v. \text{LB}(\frac{1}{2}, s, v), P_{in}, P_{out}) \end{array} \right) \right\} \\ \langle n; \text{load\_balancer } ip \ p \ srvs \rangle \\ \{\text{True}\}$$

## More examples

- ▶ We use our load balancer specs to verify an addition service:
  - ▶ individual servers for adding numbers
  - ▶ a load balancing server that distributes the load between worker servers
  - ▶ demonstrates **horizontal** modularity
- ▶ We prove correctness of a two phase commit (TPC) library
  - ▶ the implementation is parameterized by functions for:
    - ▶ voting
    - ▶ finalizing (aborting/committing)
  - ▶ We use this to implement a replicated logging service:
    - ▶ TPC specs guarantee that:
      - ▶ either **all** log servers commit or they **all** reject the change
    - ▶ demonstrates **vertical** modularity
- ▶ We prove correctness of a distributed bag data structure
  - ▶ clients can store items in the bag or request items
  - ▶ multiple threads on the server respond to requests
  - ▶ server-side bag data structure is protected by a lock
  - ▶ demonstrates **thread-local** reasoning

## More examples

- ▶ We use our load balancer specs to verify an addition service:
  - ▶ individual servers for adding numbers
  - ▶ a load balancing server that distributes the load between worker servers
  - ▶ demonstrates **horizontal** modularity
- ▶ We prove correctness of a two phase commit (TPC) library
  - ▶ the implementation is parameterized by functions for:
    - ▶ voting
    - ▶ finalizing (aborting/committing)
  - ▶ We use this to implement a replicated logging service:
    - ▶ TPC specs guarantee that:  
either **all** log servers commit or they **all** reject the change
    - ▶ demonstrates **vertical** modularity
- ▶ We prove correctness of a distributed bag data structure
  - ▶ clients can store items in the bag or request items
  - ▶ multiple threads on the server respond to requests
  - ▶ server-side bag data structure is protected by a lock
  - ▶ demonstrates **thread-local** reasoning

## More examples

- ▶ We use our load balancer specs to verify an addition service:
  - ▶ individual servers for adding numbers
  - ▶ a load balancing server that distributes the load between worker servers
  - ▶ demonstrates **horizontal** modularity
- ▶ We prove correctness of a two phase commit (TPC) library
  - ▶ the implementation is parameterized by functions for:
    - ▶ voting
    - ▶ finalizing (aborting/committing)
  - ▶ We use this to implement a replicated logging service:
    - ▶ TPC specs guarantee that:  
either **all** log servers commit or they **all** reject the change
    - ▶ demonstrates **vertical** modularity
- ▶ We prove correctness of a distributed bag data structure
  - ▶ clients can store items in the bag or request items
  - ▶ multiple threads on the server respond to requests
  - ▶ server-side bag data structure is protected by a lock
  - ▶ demonstrates **thread-local** reasoning

Thanks!