

A Verified High-Performance Composable Object Library for Remote Direct Memory Access

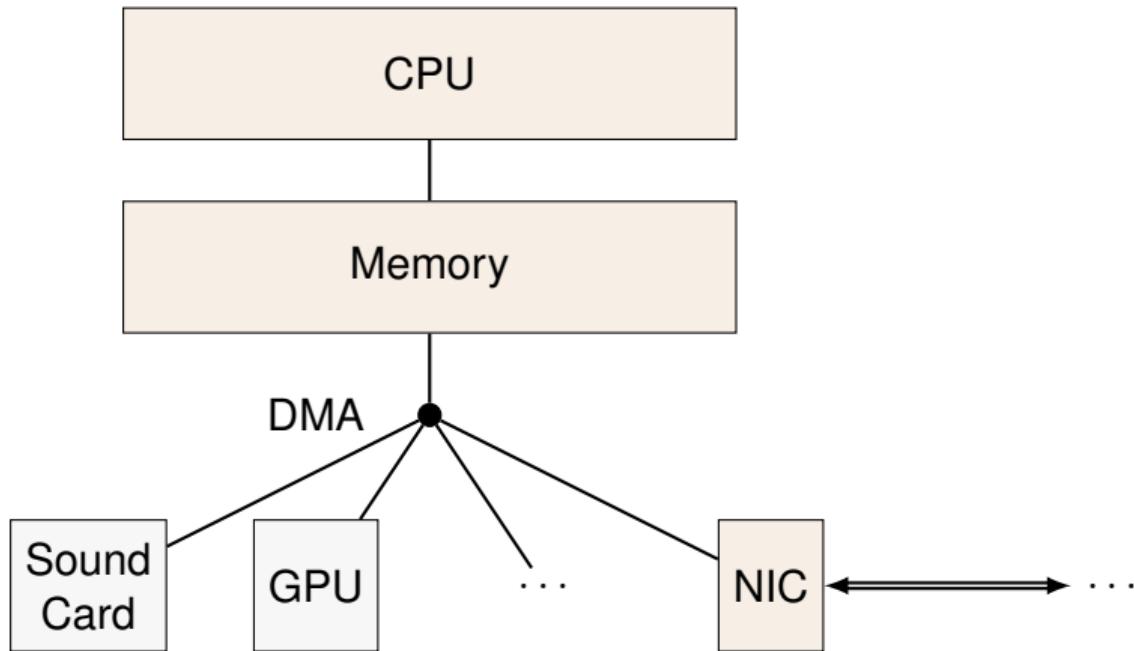
POPL 2026

Guillaume Ambal¹, George Hodgkins², Mark Madler², Gregory Chockler³,
Brijesh Dongol³, Joseph Izraelevitz², Azalea Raad¹, Viktor Vafeiadis⁴

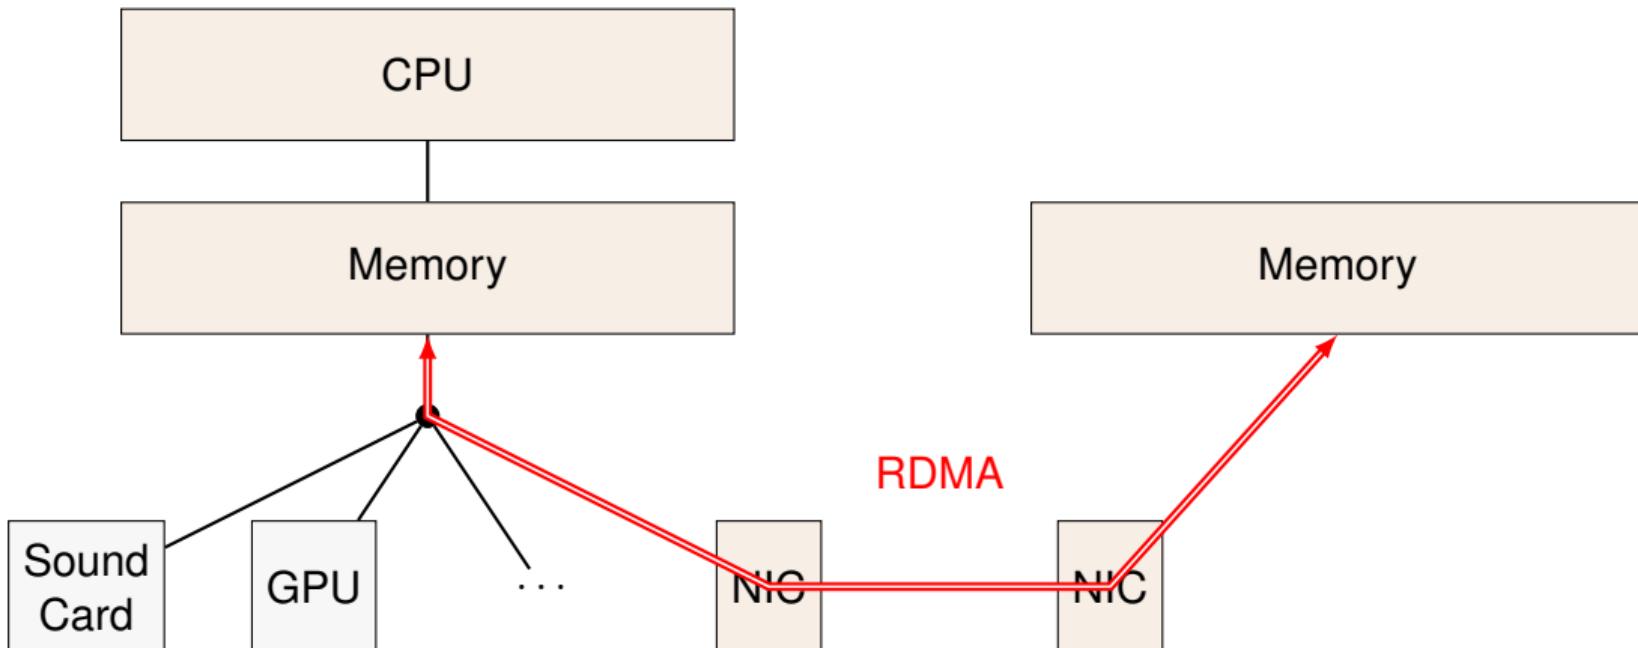
¹Imperial College London, ²University of Colorado, Boulder, ³University of Surrey, ⁴MPI-SWS

January 16, 2026

Remote Direct Memory Access (RDMA)



Remote Direct Memory Access (RDMA)



RDMA: Data-transfer protocol for High-Performance Computing

Low Latency: $\sim \mu\text{s}$

RDMA^{TSO} [OOPSLA 24]

Two nodes, one thread

$x = 0$	$z = 0$
$\tilde{z} := x$	
$x := 1$	

$z = 0$ ✓ $z = 1$?

Can the output be $z = 1$?

Two nodes, one thread

$x = 0$	$z = 0$
$\tilde{z} := x$	
$x := 1$	

$z = 0$ ✓ $z = 1$ ✓

Can the output be $z = 1$?

Yes! like this:

- CPU offloads “ $\tilde{z} := x$ ” to the NIC
- CPU executes “ $x := 1$ ”
- NIC picks up “ $\tilde{z} := x$ ” and reads 1
- ...

RDMA^{TSO} [OOPSLA 24]

$x = 0$	$z = 0$
$\tilde{z} := x$	
$x := 1$	

$z = 1$ ✓

$x = 0$	$z = 0$
$\tilde{z} := x$	
Pol1()	
$x := 1$	

$z = 1$ ✗



Can prevent this!

RDMA^{TSO} [OOPSLA 24]

$x = 0$	$z = 0$
$\tilde{z} := x$	
$x := 1$	

$z = 1$ ✓

$x = 0$	$z = 0$
$\tilde{z} := x$	
Poll()	
$x := 1$	

$z = 1$ ✗

$x = 0$	$z = 0$
⋮	
$\tilde{z} := x$	
Poll()	
$x := 1$	

$z = 1$ ✓

Can prevent this!

But it's context-dependent...

Problems and Contributions

Problems with current RDMA programs:

- Low-level non-modular code (no libraries)
- No formalisation or verification

Problems and Contributions

Problems with current RDMA programs:

- Low-level non-modular code (no libraries)
- No formalisation or verification

Our contributions:

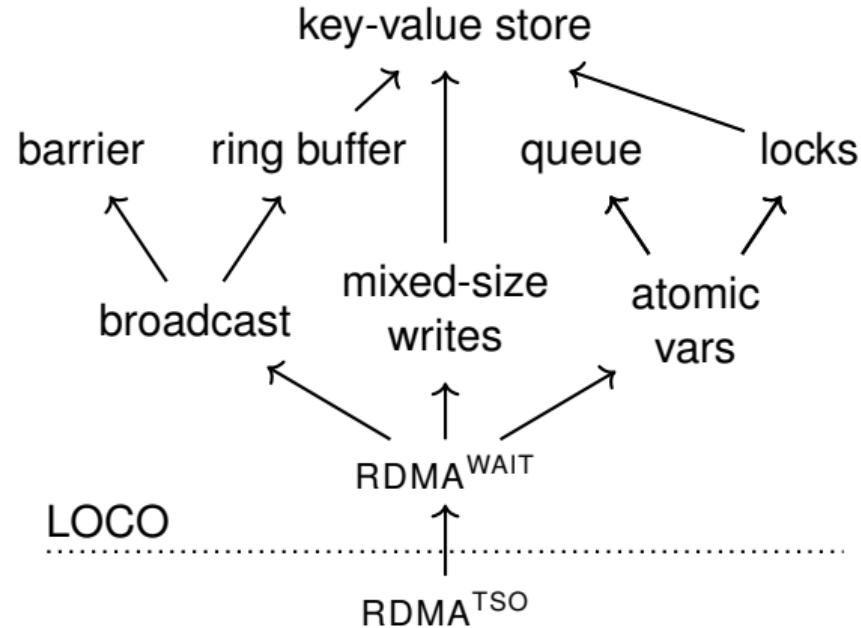
- Generalised RDMA^{TSO} to RDMA^{WAIT} as a modular alternative
- LOCO: High-performance Modular RDMA Libraries
- Mowgli: Framework for Libraries in (very) weak settings such as RDMA
- Local soundness result for simpler verification
- Specified and verified a large fragment of LOCO using Mowgli

LOCO: Overview

LOCO (Library of Composable Objects)

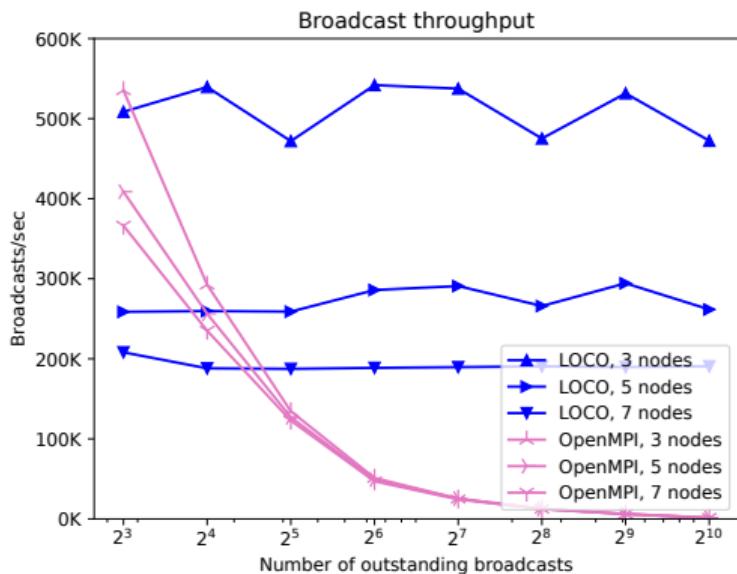
One manager thread talks to the hardware
⇒ RDMA^{WAIT} (context independent code)

Tower of libraries, allowing programming
RDMA similarly to shared-memory

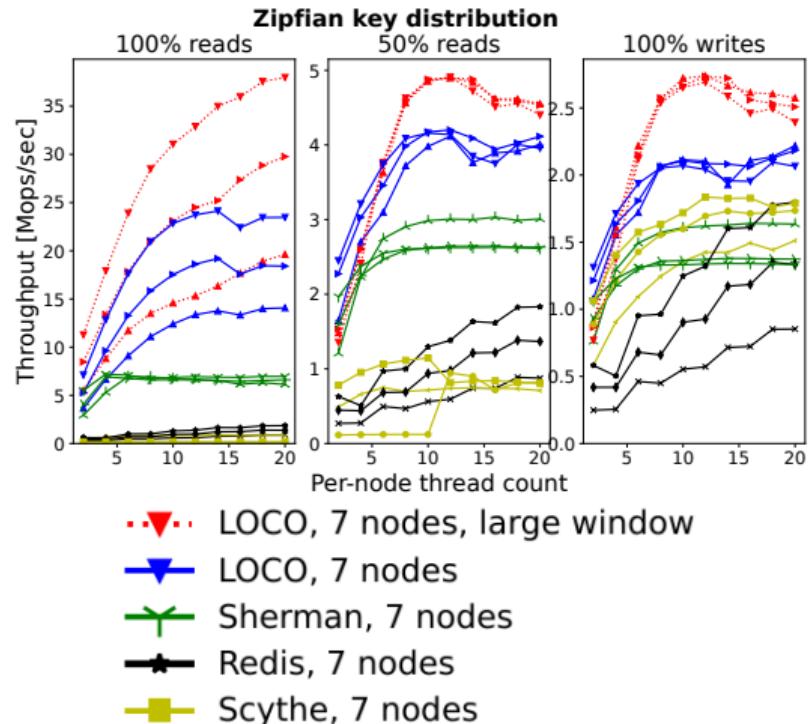


LOCO: High-Performance Libraries

Broadcast Benchmark



Key-Values Store Benchmark



Specification: Ring Buffer Library

Example 1: Ring Buffer (single-writer-multiple-reader, FIFO)

Send: $\text{Loc} \times \text{Val} \rightarrow ()$ and Recv: $\text{Loc} \rightarrow \text{Val} \cup \{\perp\}$

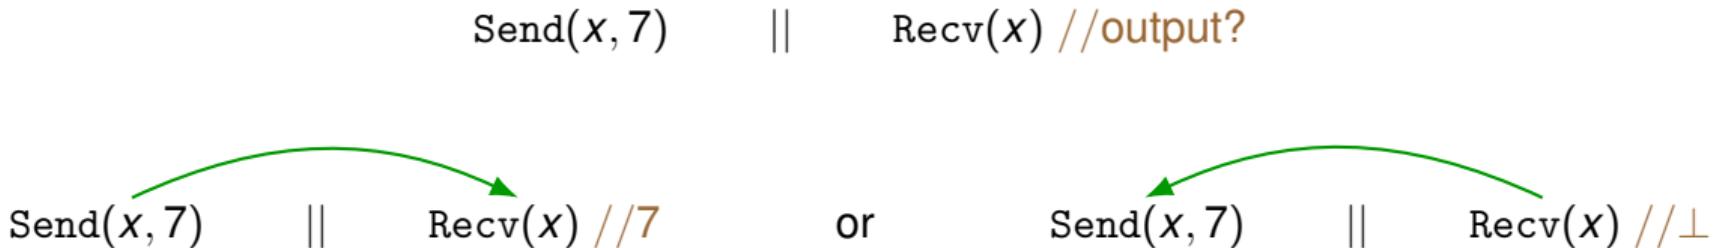
Specification: Ring Buffer Library

Example 1: Ring Buffer (single-writer-multiple-reader, FIFO)

Send: $\text{Loc} \times \text{Val} \rightarrow ()$ and Recv: $\text{Loc} \rightarrow \text{Val} \cup \{\perp\}$

$\text{Send}(x, 7) \parallel \text{Recv}(x) \text{ //output?}$

$\text{Send}(x, 7) \parallel \text{Recv}(x) \text{ //7}$ or $\text{Send}(x, 7) \parallel \text{Recv}(x) \text{ //}\perp$



Specification: Ring Buffer Library

Example 1: Ring Buffer (single-writer-multiple-reader, FIFO)

Send: $\text{Loc} \times \text{Val} \rightarrow ()$ and Recv: $\text{Loc} \rightarrow \text{Val} \cup \{\perp\}$

Send($x, 7$) || Recv(x) //output?

Send($x, 7$) || Recv(x) //7 or Send($x, 7$) || Recv(x) // \perp

Ring Buffer library spec (wishlist, first draft, version 0)

- Syntax (Send, Recv)
- Behaviour constraints (e.g. FIFO)
- Synchronisation dependencies (\rightarrow)

Specification: Barrier Library

Example 2: Barrier library

$$\text{Barr} : \text{Loc} \rightarrow ()$$

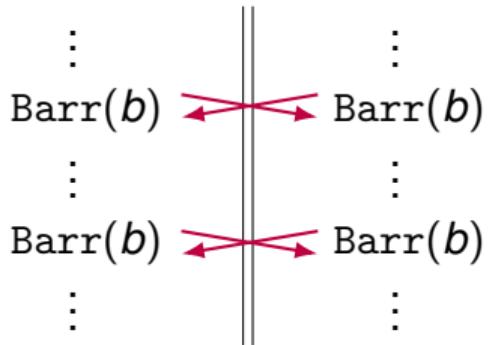
Specification: Barrier Library

Example 2: Barrier library

$$\text{Barr} : \text{Loc} \rightarrow ()$$

Thread synchronisation:

All threads reach the barrier before any can continue past.



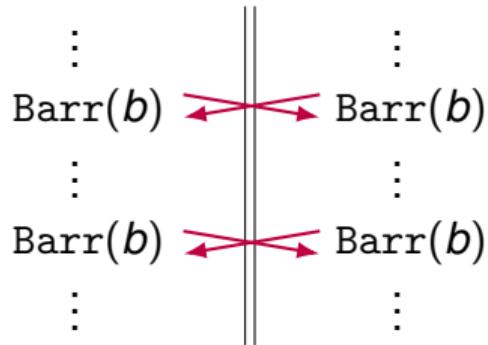
Specification: Barrier Library

Example 2: Barrier library

$$\text{Barr} : \text{Loc} \rightarrow ()$$

Thread synchronisation:

All threads reach the barrier before any can continue past.



Barrier library specification (wishlist, first draft, version 0)

- Syntax (Barr)
- Behaviour constraints (e.g. used same number of times by each thread)
- Synchronisation dependencies (\rightarrow)

Execution

Ring Buffer x , Barrier b	
Send($x, 7$) Barr(b)	Barr(b) $a := \text{Recv}(x)$

Claim: $a \neq \perp$

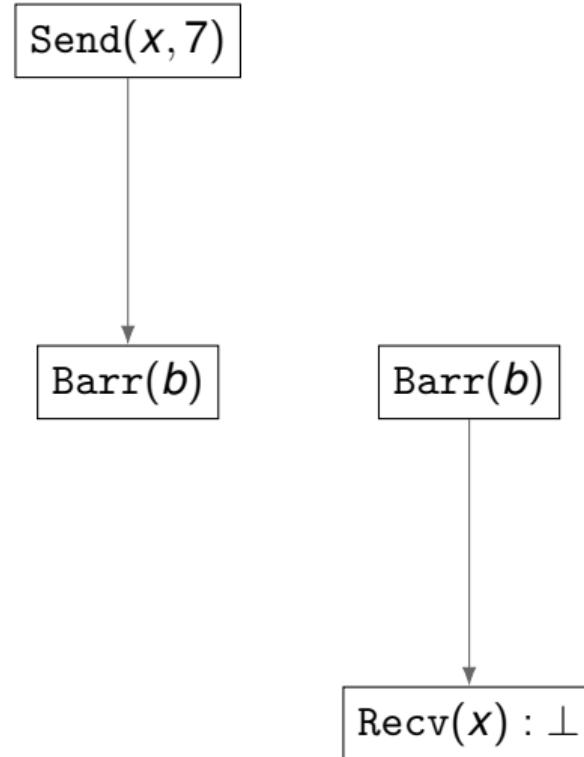
A global execution is allowed if it:

- respects the semantics of every library
- doesn't create dependency cycles

Execution

Ring Buffer x , Barrier b	
Send($x, 7$)	
Barr(b)	Barr(b) $a := \text{Recv}(x)$

Claim: $a \neq \perp$



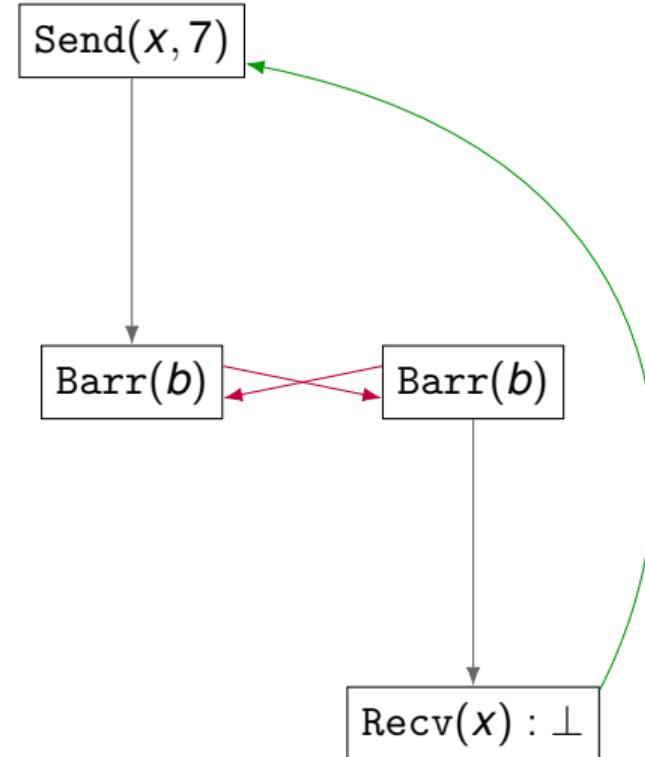
A global execution is allowed if it:

- respects the semantics of every library
- doesn't create dependency cycles

Execution

Ring Buffer x , Barrier b	
Send($x, 7$)	
Barr(b)	Barr(b) $a := \text{Recv}(x)$

Claim: $a \neq \perp$



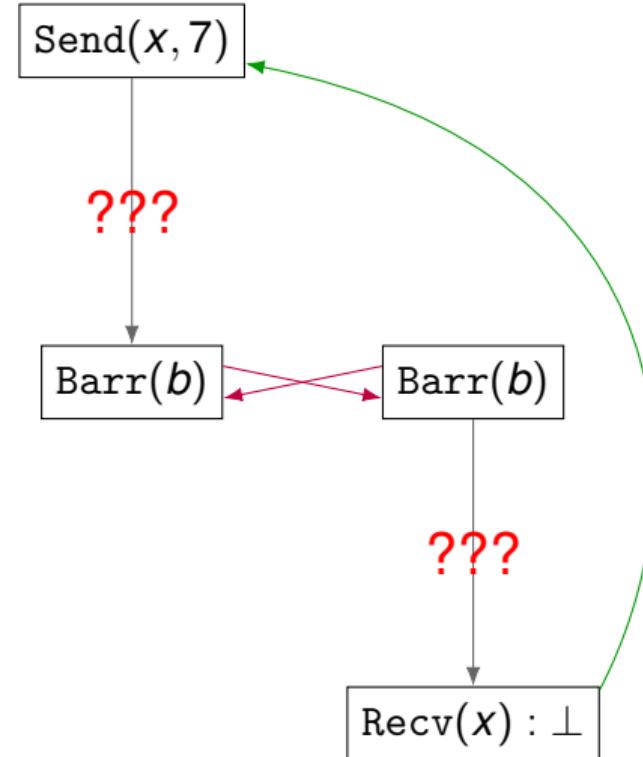
A global execution is allowed if it:

- respects the semantics of every library
- doesn't create dependency cycles

Execution (First Try)

Ring Buffer x , Barrier b	
Send($x, 7$)	
Barr(b)	Barr(b) $a := \text{Recv}(x)$

Claim: $a \neq \perp$



A global execution is allowed if it:

- respects the semantics of every library
- doesn't create dependency cycles

Program Order Dependencies???

Remember the first example:

$x = 0$	$z = 0$
$\tilde{z} := x$ $x := 1$	

$z = 1$ ✓

Program order is **not** always respected...

When is it respected?

Stamp ordering for RDMA

Previous documented hardware (re)orderings:

Ordered?		Second Stamp										
		single					families					
single	1	2	3	4	5	6	7	8	9	10	11	
	aCR	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	aCW	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	aCAS	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	aMF	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
families	aWT	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	aNLR _n	✗	✗	✗	✗	✗	SN	SN	SN	SN	SN	SN
	aNRW _n	✗	✗	✗	✗	✗	✗	SN	SN	SN	✗	SN
	aNRR _n	✗	✗	✗	✗	✗	✗	✗	✗	SN	SN	SN
	aNLW _n	✗	✗	✗	✗	✗	✗	✗	✗	SN	✗	SN
	aRF _n	✗	✗	✗	✗	✗	SN	SN	SN	SN	SN	SN
	aGF _n	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

11 stamps (categories of operations) define when ordering is preserved.

Program Order Dependencies

Remember the first example:

$x = 0$	$z = 0$
$\tilde{z} := x$	
$x := 1$	

$z = 1$ ✓

Program order is **not** always respected...

Reorderings for (independent) libraries?

Program Order Dependencies

Remember the first example:

$x = 0$	$z = 0$
$\tilde{z} := x$	
$x := 1$	

$z = 1$ ✓

Program order is **not** always respected...

Reorderings for (independent) libraries?

Mowgli solution:

- Use hardware behaviours as categories (stamps).
- Libraries specify the effects of each function.

Mowgli Ring Buffer library specification

- like before: syntax, semantics, →
- Send is in the “NIC write” category
- Recv is in the “CPU read” category

Mowgli Barrier library specification

- like before: syntax, semantics, →
- Barr has Fence and Read behaviours

Mowgli Execution (Fixed)

Ring Buffer x , Barrier b	
Send($x, 7$)	
Barr(b)	Barr(b) $a := \text{Recv}(x)$

Claim: $a \neq \perp$

- like before: syntax, semantics, \rightarrow
- Send has a NIC write behaviour
- Recv has a CPU read behaviour

- like before: syntax, semantics, \rightarrow
- Barr has Fence and Read behaviours

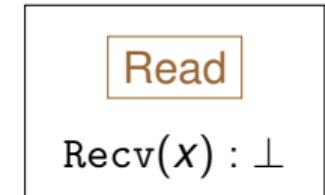
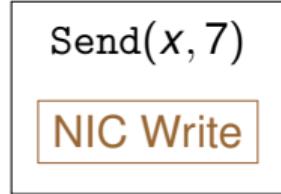
Mowgli Execution (Fixed)

Ring Buffer x , Barrier b	
Send($x, 7$)	
Barr(b)	Barr(b) $a := \text{Recv}(x)$

Claim: $a \neq \perp$

- like before: syntax, semantics, →
- Send has a NIC write behaviour
- Recv has a CPU read behaviour

- like before: syntax, semantics, →
- Barr has Fence and Read behaviours



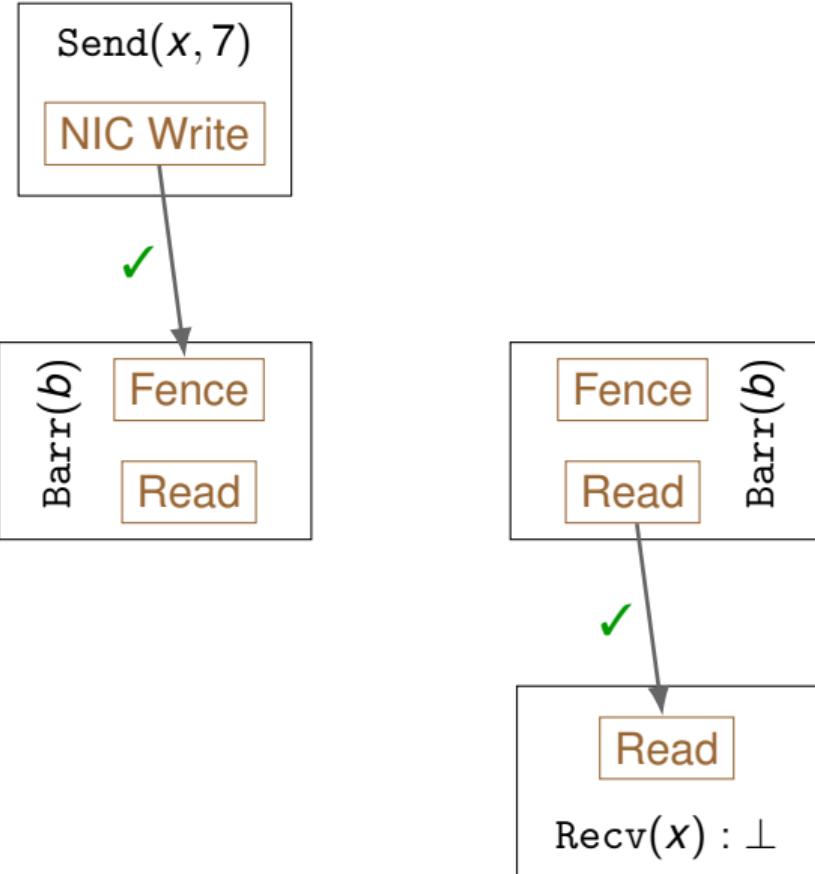
Mowgli Execution (Fixed)

Ring Buffer x , Barrier b	
Send($x, 7$)	
	Barr(b)
	$a := \text{Recv}(x)$

Claim: $a \neq \perp$

- like before: syntax, semantics, \rightarrow
- Send has a NIC write behaviour
- Recv has a CPU read behaviour

- like before: syntax, semantics, \rightarrow
- Barr has Fence and Read behaviours



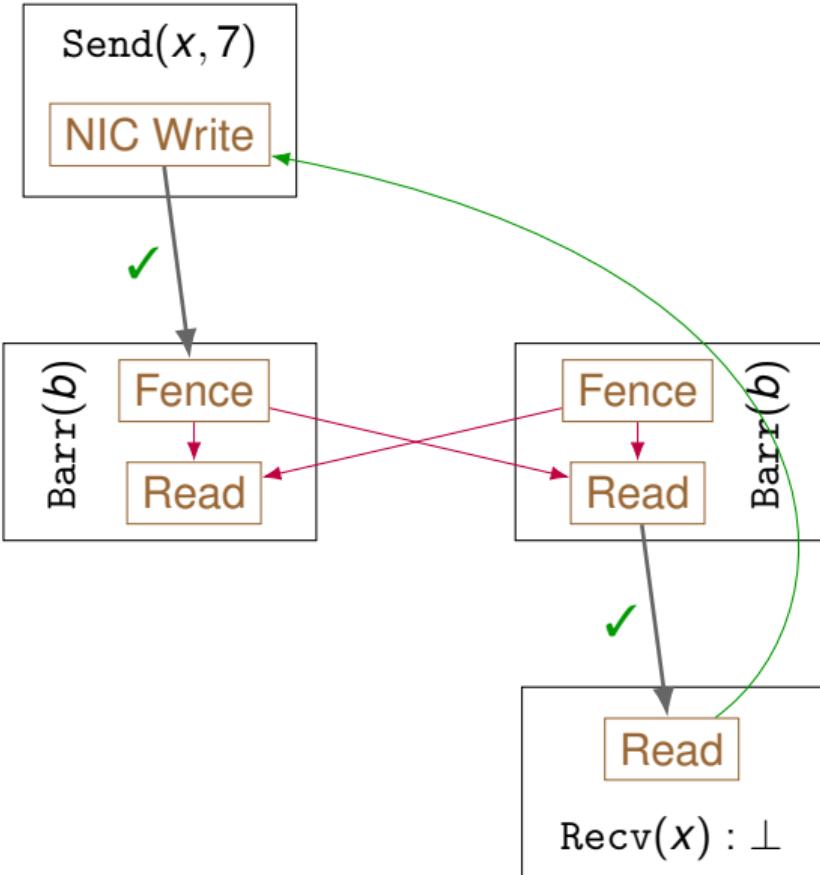
Mowgli Execution (Fixed)

Ring Buffer x , Barrier b	
Send($x, 7$)	
	Barr(b) $a := \text{Recv}(x)$

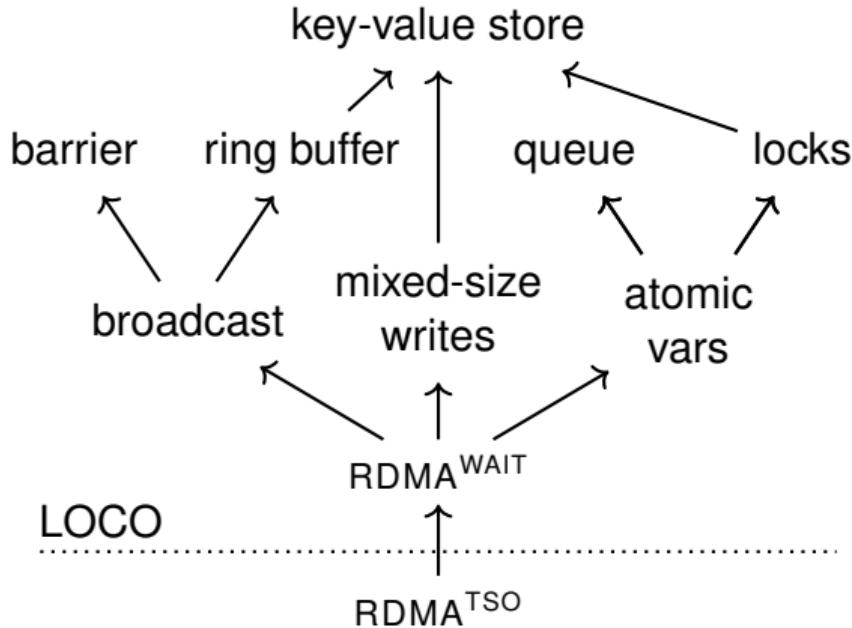
Claim: $a \neq \perp$

- like before: syntax, semantics, 
- Send has a NIC write behaviour
- Recv has a CPU read behaviour

- like before: syntax, semantics, 
- Barr has Fence and Read behaviours



Verifying Library Implementations



Libraries can be implemented using other libraries

Programs can use a set Λ of libraries

Inlining an implementation l (of library L) should be fine

Implementation Soundness

How do we prove correct an implementation I for a library L ?

- **Soundness** [general, hard to show]:

Outcomes of the implementation are valid, in **all** contexts:

$$\forall \Lambda, P. \text{outcome}_{\Lambda}(I(P)) \subseteq \text{outcome}_{\Lambda \cup \{L\}}(P)$$

Implementation Soundness

How do we prove correct an implementation I for a library L ?

- **Soundness** [general, hard to show]:

Outcomes of the implementation are valid, in **all** contexts:

$$\forall \Lambda, P. \text{outcome}_{\Lambda}(I(P)) \subseteq \text{outcome}_{\Lambda \cup \{L\}}(P)$$

- **Local soundness** [library specific, easier to show]:

Promises of the library specification are kept (formal definition is a bit technical).

Considers I and L in isolation.

- **Mowgli Locality Theorem:**

Local soundness implies soundness.

Implementation Soundness

How do we prove correct an implementation I for a library L ?

- **Soundness** [general, hard to show]:

Outcomes of the implementation are valid, in **all** contexts:

$$\forall \Lambda, P. \text{outcome}_{\Lambda}(I(P)) \subseteq \text{outcome}_{\Lambda \cup \{L\}}(P)$$

- **Local soundness** [library specific, easier to show]:

Promises of the library specification are kept (formal definition is a bit technical).

Considers I and L in isolation.

- **Mowgli Locality Theorem:**

Local soundness implies soundness.

Programmers prove local soundness, Mowgli gives soundness for free.

Conclusion

Summary:

- Generalised RDMA^{TSO} to RDMA^{WAIT} as a modular alternative
- LOCO: High-performance Modular RDMA Libraries
- Mowgli: Framework for Libraries in (very) weak settings such as RDMA
- Local soundness result for simpler verification
- Specified and verified a large fragment of LOCO using Mowgli

Next: extension to RDMA Atomic operations (ESOP 2026)

Conclusion

Summary:

- Generalised RDMA^{TSO} to RDMA^{WAIT} as a modular alternative
- LOCO: High-performance Modular RDMA Libraries
- Mowgli: Framework for Libraries in (very) weak settings such as RDMA
- Local soundness result for simpler verification
- Specified and verified a large fragment of LOCO using Mowgli

Next: extension to RDMA Atomic operations (ESOP 2026)

Thanks for listening!

Poll vs Wait Semantics

Base RDMA (RDMA^{TSO}):

$x = 0$	$z = 0$
$\tilde{z} := x$	
$x := 1$	

$z = 1$ ✓

$x = 0$	$z = 0$
$\tilde{z} := x$	
poll()	
$x := 1$	

$z = 1$ ✗

$x = 0$	$z = 0$
⋮	
$\tilde{z} := x$	
poll()	
$x := 1$	

$z = 1$ ✓

Poll vs Wait Semantics

Base RDMA (RDMA^{TSO}):

$x = 0$	$z = 0$
$\tilde{z} := x$	
$x := 1$	

$z = 1$ ✓

$x = 0$	$z = 0$
$\tilde{z} := x$	
Poll()	
$x := 1$	

$z = 1$ ✗

$x = 0$	$z = 0$
\vdots	
$\tilde{z} := x$	
Poll()	
$x := 1$	

$z = 1$ ✓

LOCO (RDMA^{WAIT}):

$x = 0$	$z = 0$
$\tilde{z} := x$	
$x := 1$	

$z = 1$ ✓

$x = 0$	$z = 0$
$\tilde{z} :=^d x$	
Wait(d)	
$x := 1$	

$z = 1$ ✗

$x = 0$	$z = 0$
\vdots	
$\tilde{z} :=^d x$	
Wait(d)	
$x := 1$	

$z = 1$ ✗