

Konstantin Taranov, Rodrigo Bruno, Gustavo Alonso, Torsten Hoefler

Naos: Serialization-free RDMA networking in Java



Sending and Receiving Objects in Java

```
Person person = new Person(18, "Mike");
```

```
1 /* Send an object with Kryo */
2 /* Prepare a send buffer */
3 ByteBuffer buffer = ByteBuffer.allocate(512);
4 /* Serialization */
5 kryo.register(Person.class);
6 Output out = new Output(buffer);
7 kryo.writeObject(out, person);
8 /* Networking */
9 connection.write(buffer);
```

- ← Memory allocation
- ← Class registration
- ← Serialization
- ← Writing buffer to the network

```
1 /* Receive an object with Kryo */
2 /* Prepare a receive buffer */
3 ByteBuffer buffer = ByteBuffer.allocate(512);
4 /* Networking */
5 connection.read(buffer);
6 /* De-serialization */
7 kryo.register(Person.class);
8 Input in = new Input(buffer);
9 Object obj = kryo.readObject(in, Person.class);
10 Person person = (Person)obj;
```

- ← Memory allocation
- ← Reading buffer from the network
- ← Class registration
- ← De-serialization

Serialization process

```
Person person = new Person(18, "Mike");
```

```
1 public class Person {
2     int age;
3     char[] name;
4 }
```

JVM specific

JVM memory




Portable across JVMs

Java serialization format



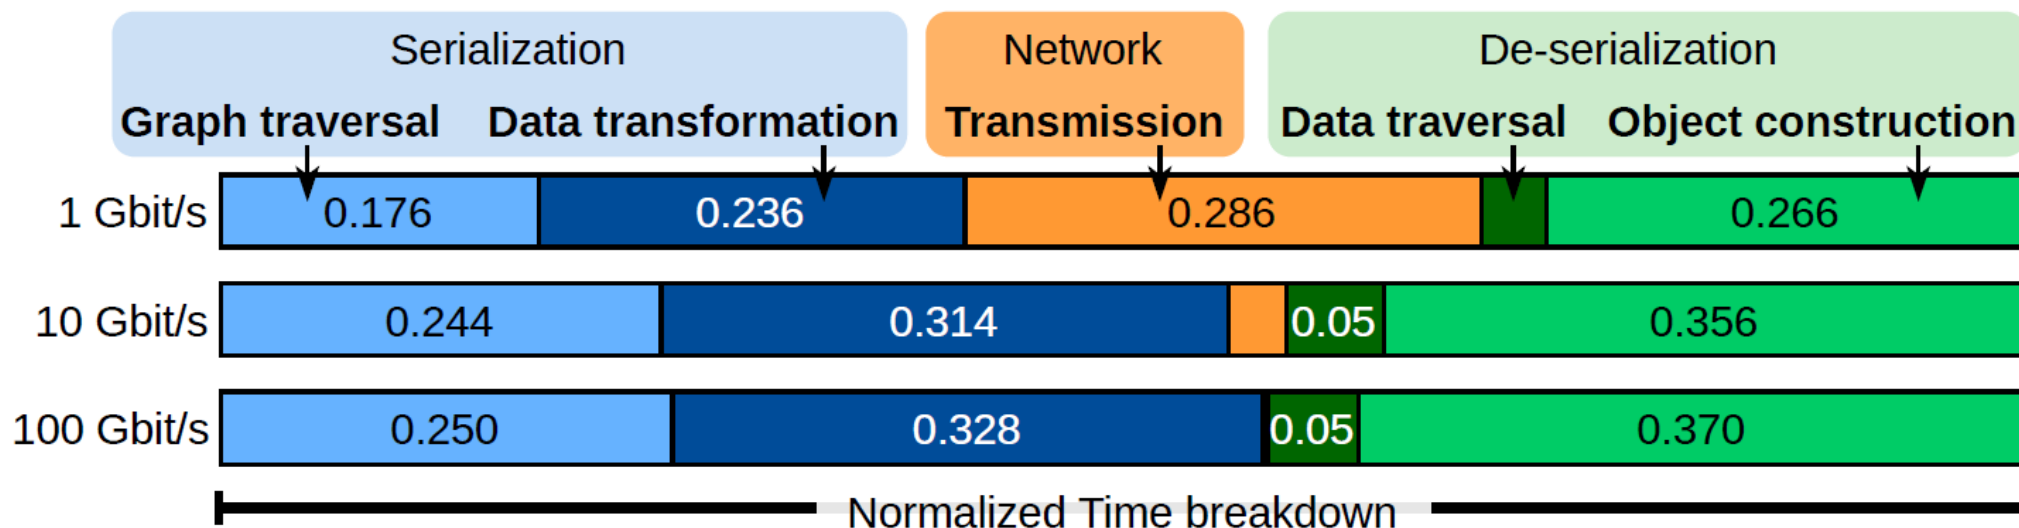
Registered types

Kryo + type registration



Impact of network bandwidth on sending objects between JVMs

- We are sending an array of 1.28M objects.



The network is not a bottleneck anymore!

Serialization already accounts for 6% of total CPU cycles at Google datacenters*.

Naos – serialization-free networking library

```

1 /* Send an object with Kryo */
2 /* Prepare a send buffer */
3 ByteBuffer buffer = ByteBuffer.allocate(512);
4 /* Serialization */
5 kryo.register(Person.class);
6 Output out = new Output(buffer);
7 kryo.writeObject(out, person);
8 /* Networking */
9 connection.write(buffer);
    
```

naos

```
connection.writeObject(person);
```

```

1 /* Receive an object with Kryo */
2 /* Prepare a receive buffer */
3 ByteBuffer buffer = ByteBuffer.allocate(512);
4 /* Networking */
5 connection.read(buffer);
6 /* De-serialization */
7 kryo.register(Person.class);
8 Input in = new Input(buffer);
9 Object obj = kryo.readObject(in, Person.class)
10 Person person = (Person)obj;
    
```

naos

```
Person person = (Person)connection.readObject();
```

Naos supports TCP and RDMA

API


```

void writeObject (Object)
Object readObject ()
boolean isReadable ()
long writeObjectAsync (Object)
int waitHandle (long)
int testHandle (long)
    
```

JVM
memory



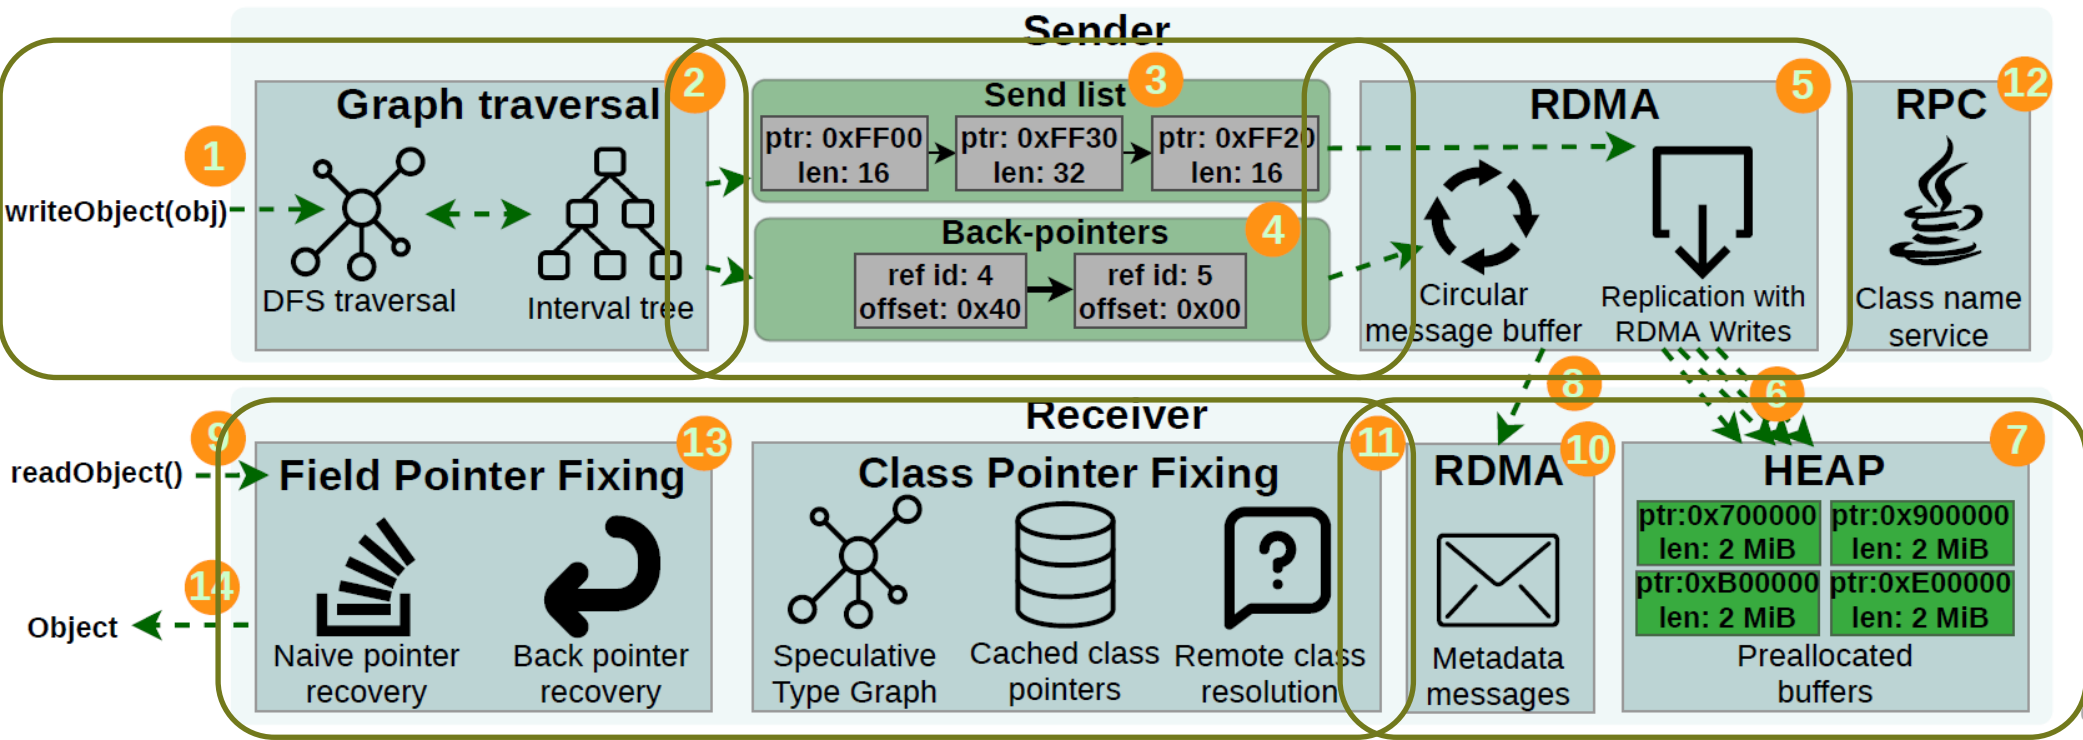
Naos



Naos sends objects directly from local JVM memory to remote JVM memory *without data transformations*.

Core challenges

- Naos cannot modify sender's memory, as the data is sent directly from the heap
- Garbage collector can concurrently touch/move objects
- RDMA-capable NICs can concurrently access the memory




Sending object with Naos

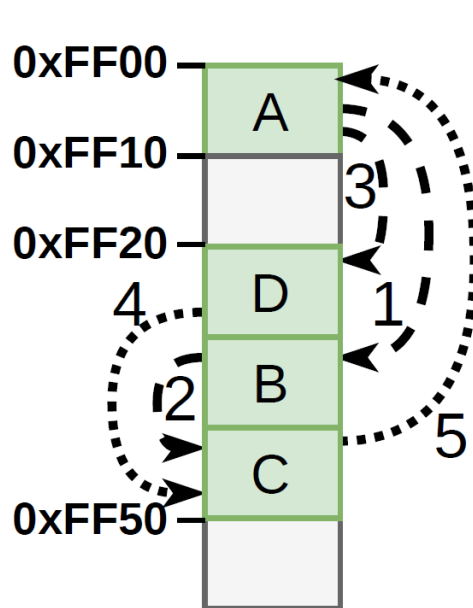


Naos proposes a novel way of the graph traversal:

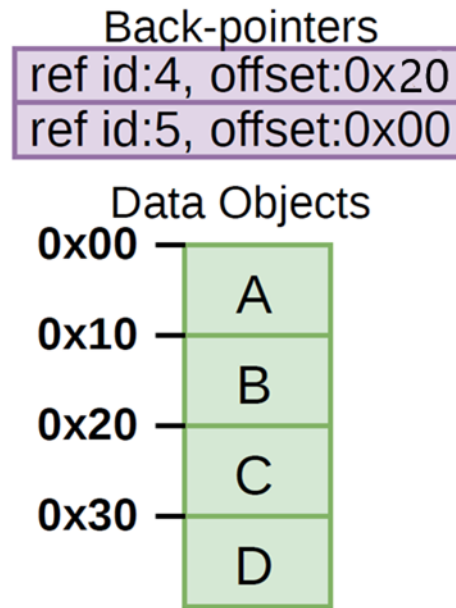
- It traverses object graphs in depth-first search (DFS) order on both sender and receiver
- It does not modify JVM memory and sends zero metadata about "trivial" pointers



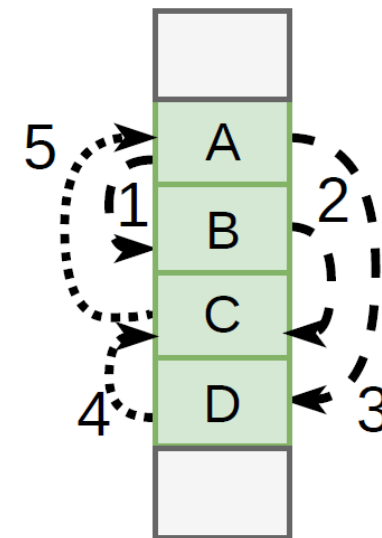
Logical View



Sender Memory



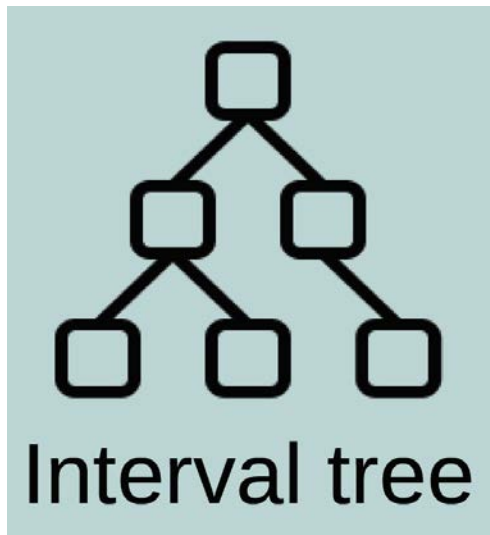
Network




Receiver Memory

Related Optimizations

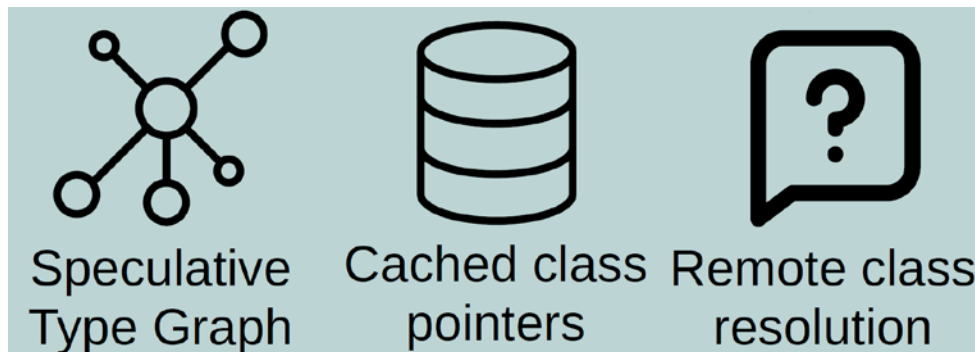
For back-pointer detection



To recover object pointers at the receiver

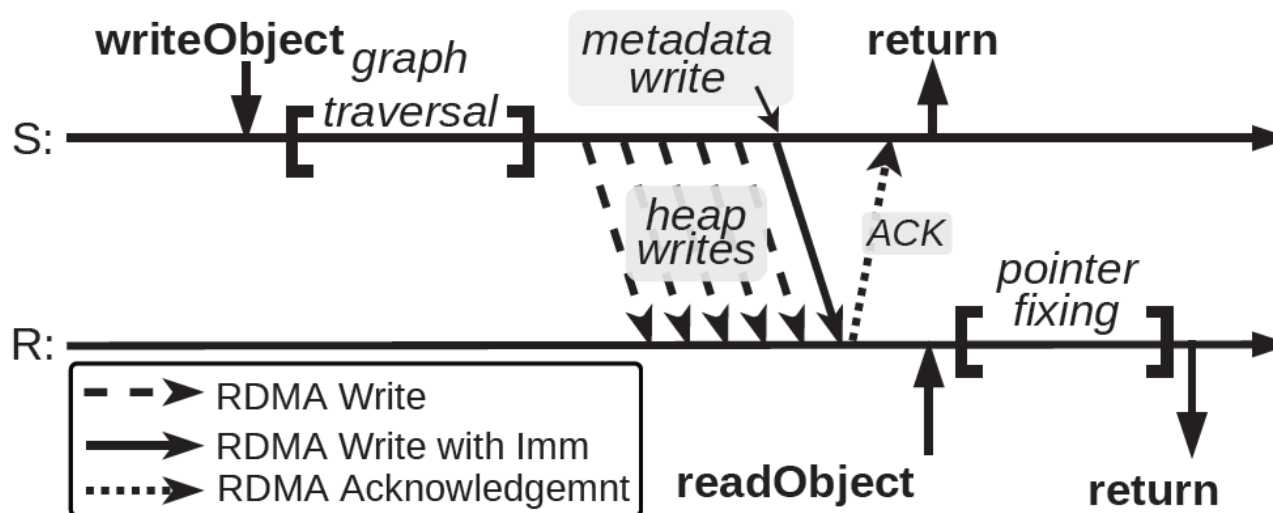


To recover class pointers at the receiver



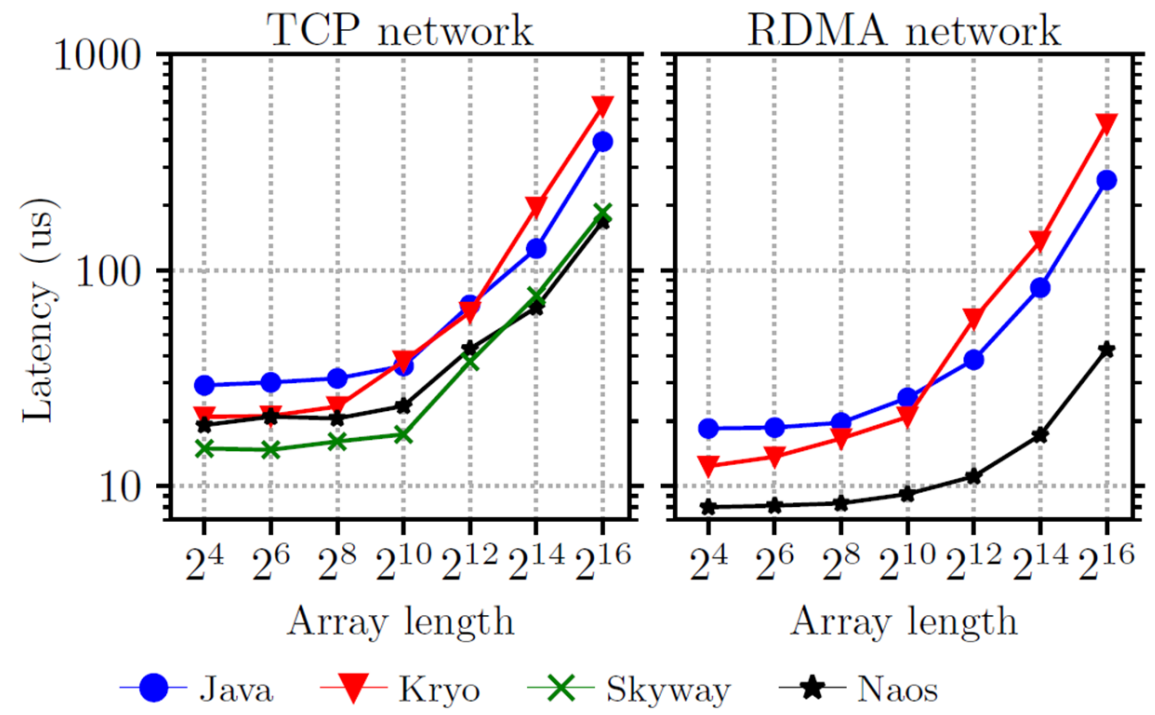
RDMA design

- The receiver reserves heap memory for receiving objects
- The sender “silently” writes object to the remote heap using RDMA writes
- Then the sender sends metadata that indicates the completion of a send

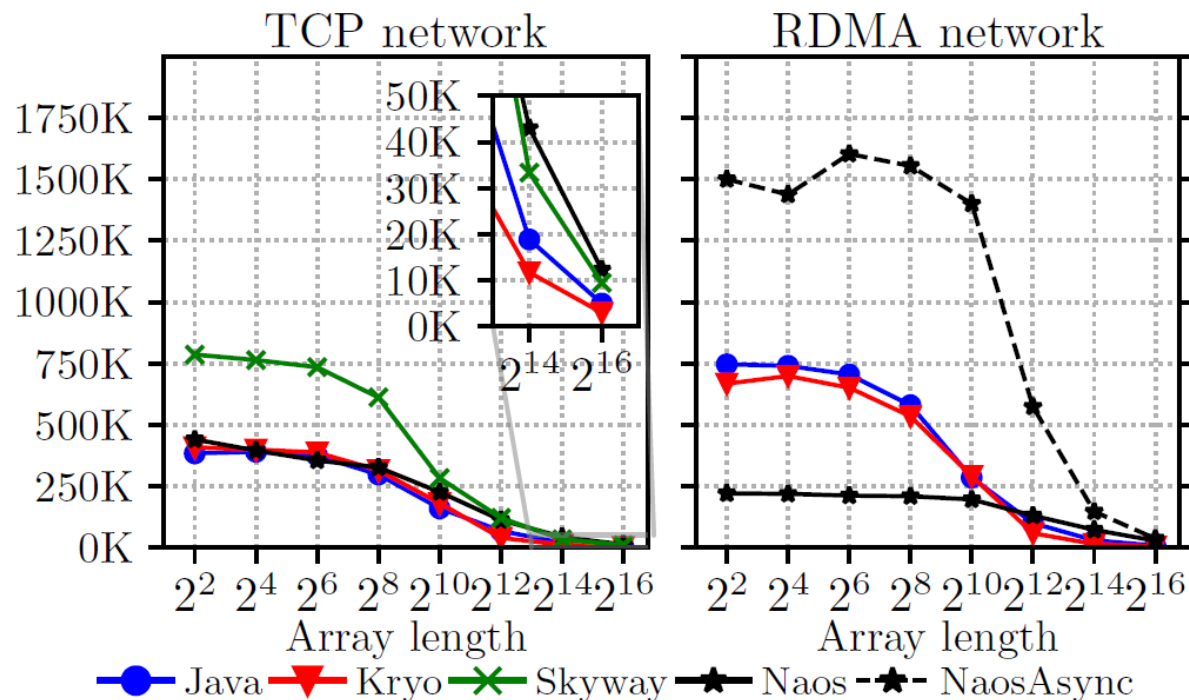


Evaluation

RTT latency for sending float[]



Throughput for sending float[]



Thank you for your attention!

- Naos is the first serialization-free communication library for JVMs
- Naos does not perform excessive data copies during communication
- Naos can directly send objects across Java heaps
- Naos supports RDMA networking



Naos implementation:



Speaker:

Konstantin Taranov

For questions:

konstantin.taranov@inf.ethz.ch