

Autonomic Solutions for Parallel and Distributed Data Stream Processing (Auto-DaSP) - Euro-Par 2017 Workshop

STREAMING IN THE PGAS ERA

Marco Aldinucci, Maurizio Drocco
University of Torino, Italy



OUTLINE

- Programming models
- Distributed Memories: the PGAS model
 - Some example
- Streaming
 - Some example
- Streaming with PGAS

LOW-LEVEL PARALLEL PROGRAMMING MODELS

- Message-Passing
 - Scalability and performance
 - Developer-based precise knowledge of code and overhead
 - Processes + communications (symmetric/collective, blocking/nonblocking)
- Shared-Memory
 - Productivity
 - Global and uniform vision of data layout
 - Threads + synchronisations mechanisms (mutex, atomics, transactions, ...)

MPI IS LIKE A CAR, YOU CAN DRIVE DATA WHERE YOU LIKE
–D.K. Panda, leader MVAPICH project at Ohio state Uni.



Car


Val D'Orcia, Tuscany, Italy



EXPECTATIONS



REALITY



ANYWAY, MPI IS NO. 1 IN HPC

Courtesy of P. Messina, director of ECP.
No. of software proposals in US ECP 2017

DISTRIBUTED SHARED (VIRTUAL) MEMORY — DSM OR DVSM

- Physically separated memories can be addressed as one logically shared address space
- Hardware or software. Conceptually similar to Virtual Memory
- Designed to distributed platform transparent to programmer i.e. **"simplify programming"**
- "Vanilla" API
 - **read(addr)**
 - **write(value, adds)**
 - **lock/unlock**

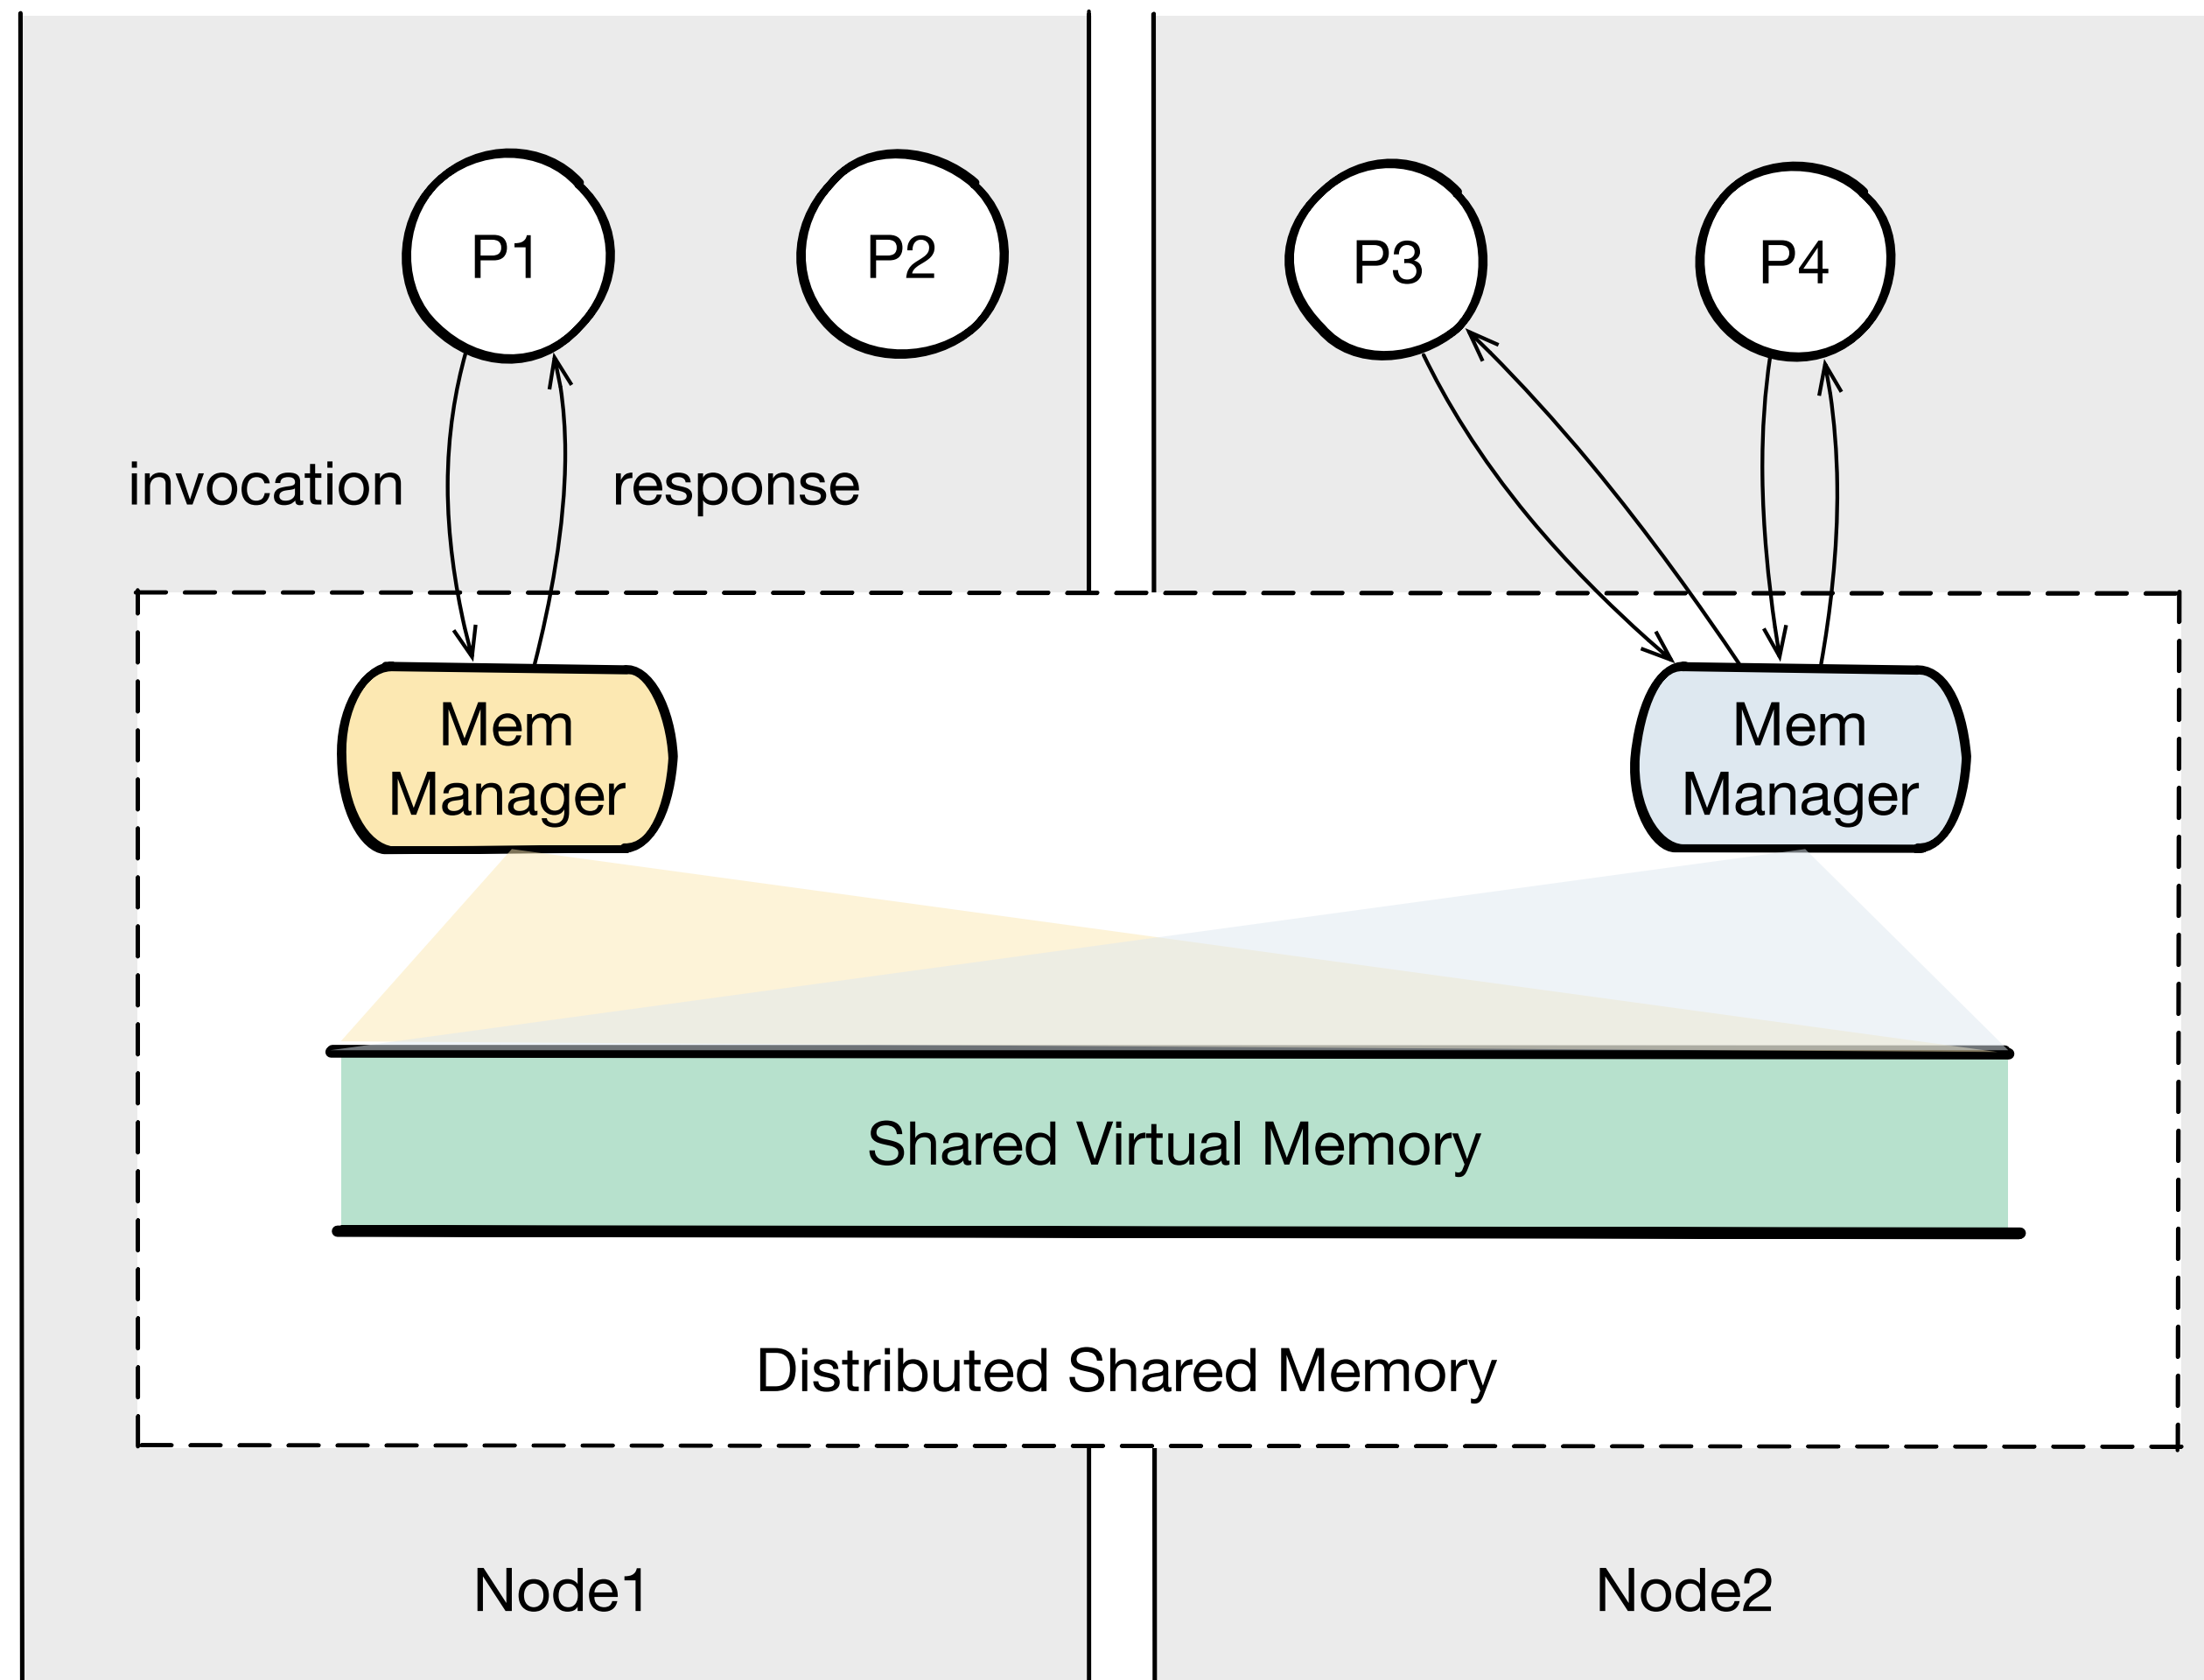


Table A. Software DSM implementations.

IMPLEMENTATION	TYPE OF IMPLEMENTATION	TYPE OF ALGORITHM	CONSISTENCY MODEL	GRANULARITY UNIT	COHERENCE POLICY
IVY	User-level library + OS modification	MRSW	Sequential	1 Kbyte	Invalidate
Mermaid	User-level library + OS modifications	MRSW	Sequential	1 Kbyte, 8 Kbytes	Invalidate
Munin	Runtime system + linker + library + preprocessor + OS modifications	Type-specific (SRSW, MRSW MRMW)	Release	Variable size objects	Type-specific (delayed update, invalidate) Update
Midway	Runtime system + compiler	MRMW	Entry, release, processor	4 Kbytes	Update
TreadMarks	User-level	MRMW	Lazy release	4 Kbytes	Update, invalidate
Blizzard	User-level + OS kernel modification	MRSW	Sequential	32-128 bytes	Invalidate
Mirage	OS kernel	MRSW	Sequential	512 bytes	Invalidate
Clouds	OS, out of kernel	MRSW	Inconsistent, sequential	8 Kbytes	Discard segment when unlocked
Linda	Language	MRSW	Sequential	Variable (tuple size)	Implementation- dependent
Orca	Language	MRSW	Synchronization dependent	Shared data object size	Update

Already mature 20 years ago,
now quite rotten

P. Jelica, M. Tomasevic, and V. Milutinovic. "Distributed shared memory: Concepts and systems." IEEE Parallel & Distributed Technology: Systems & Applications 4.2 (1996): 63-71.

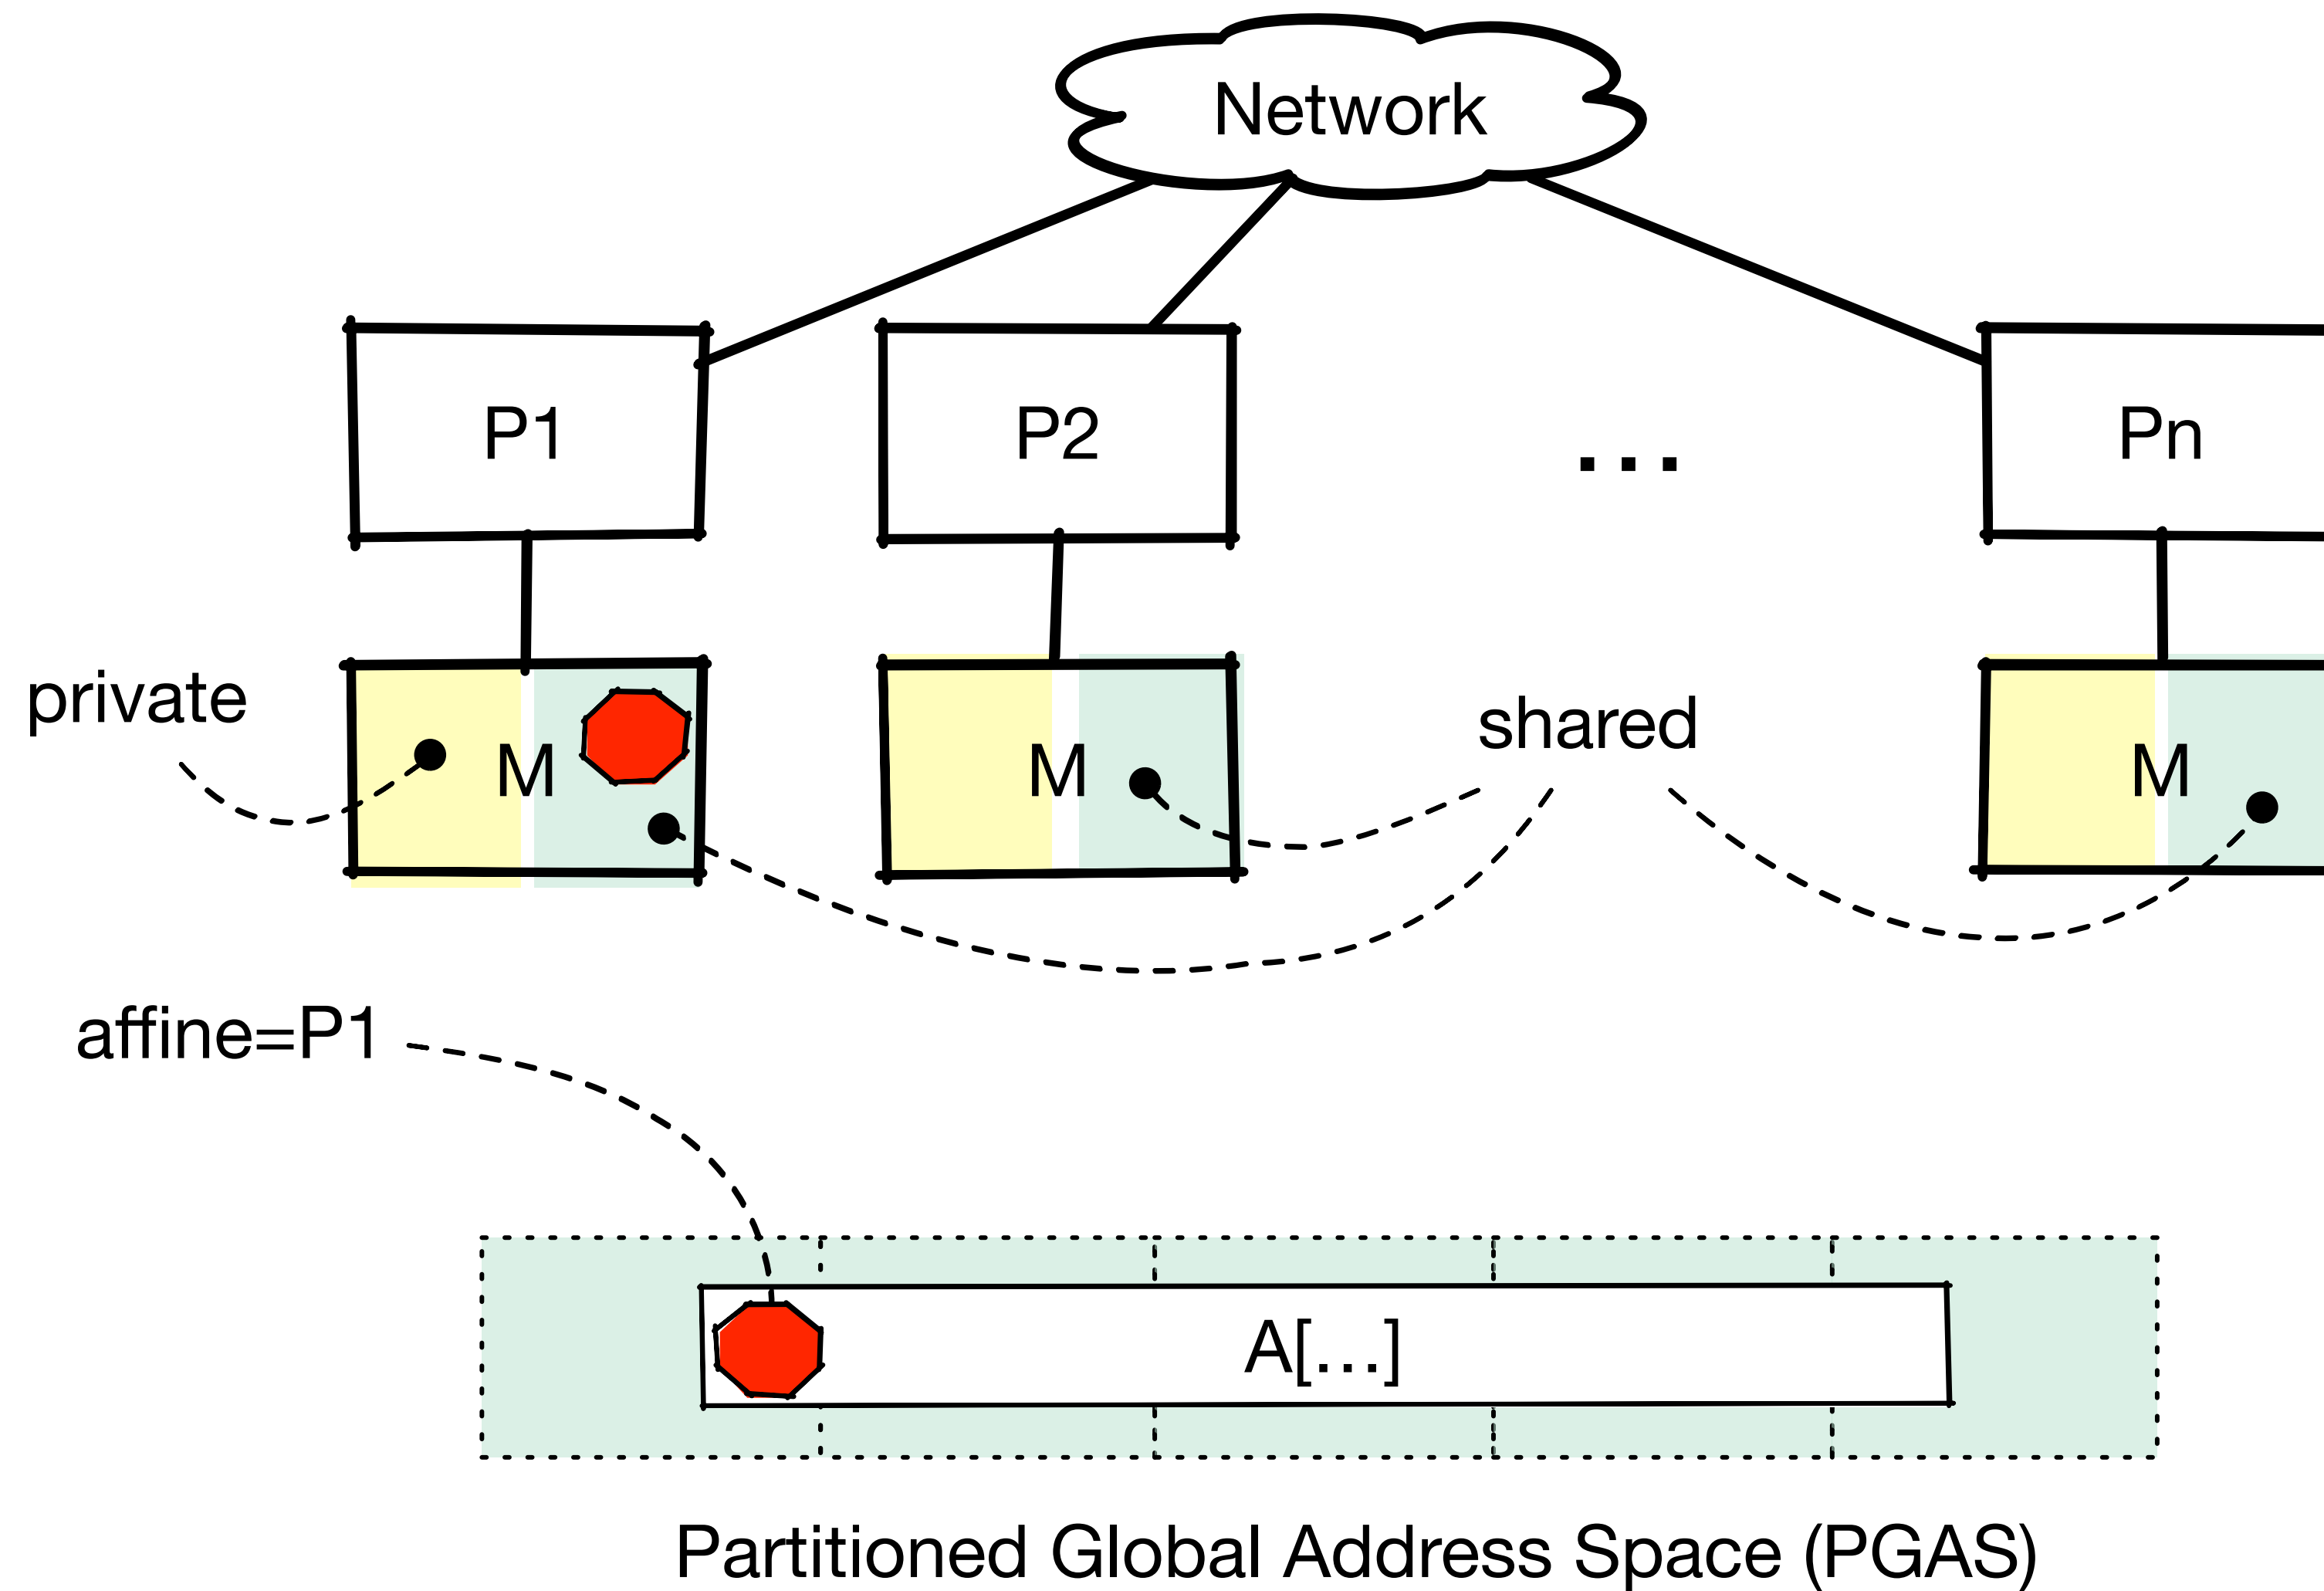
DSM: WHY THEY FAILED

- Started to simplify distributed code
- To make them efficient, we made the memory consistency model very complex
- This seriously affect the coding effort

PGAS PROGRAMMING MODEL (DSM EVOLVED)

- A set of processor, each with own local memory
- Part managed as private, part as shared
 - Sharing implemented HW or SW
- **Explicitly NUMA**
 - Each location has an affinity with a processor
 - Model differentiates between local and remote data partitions
- **Explicitly partitioned**
 - Collective synchronisations, i.e. barriers and fences

PGAS SYSTEM AT BARE BONES



APPROACHES

Languages

- Unified Parallel C (UPC)
- Co-Array Fortran (CAF)
- X10
- **Chapel**
- STAPL
- Titanium

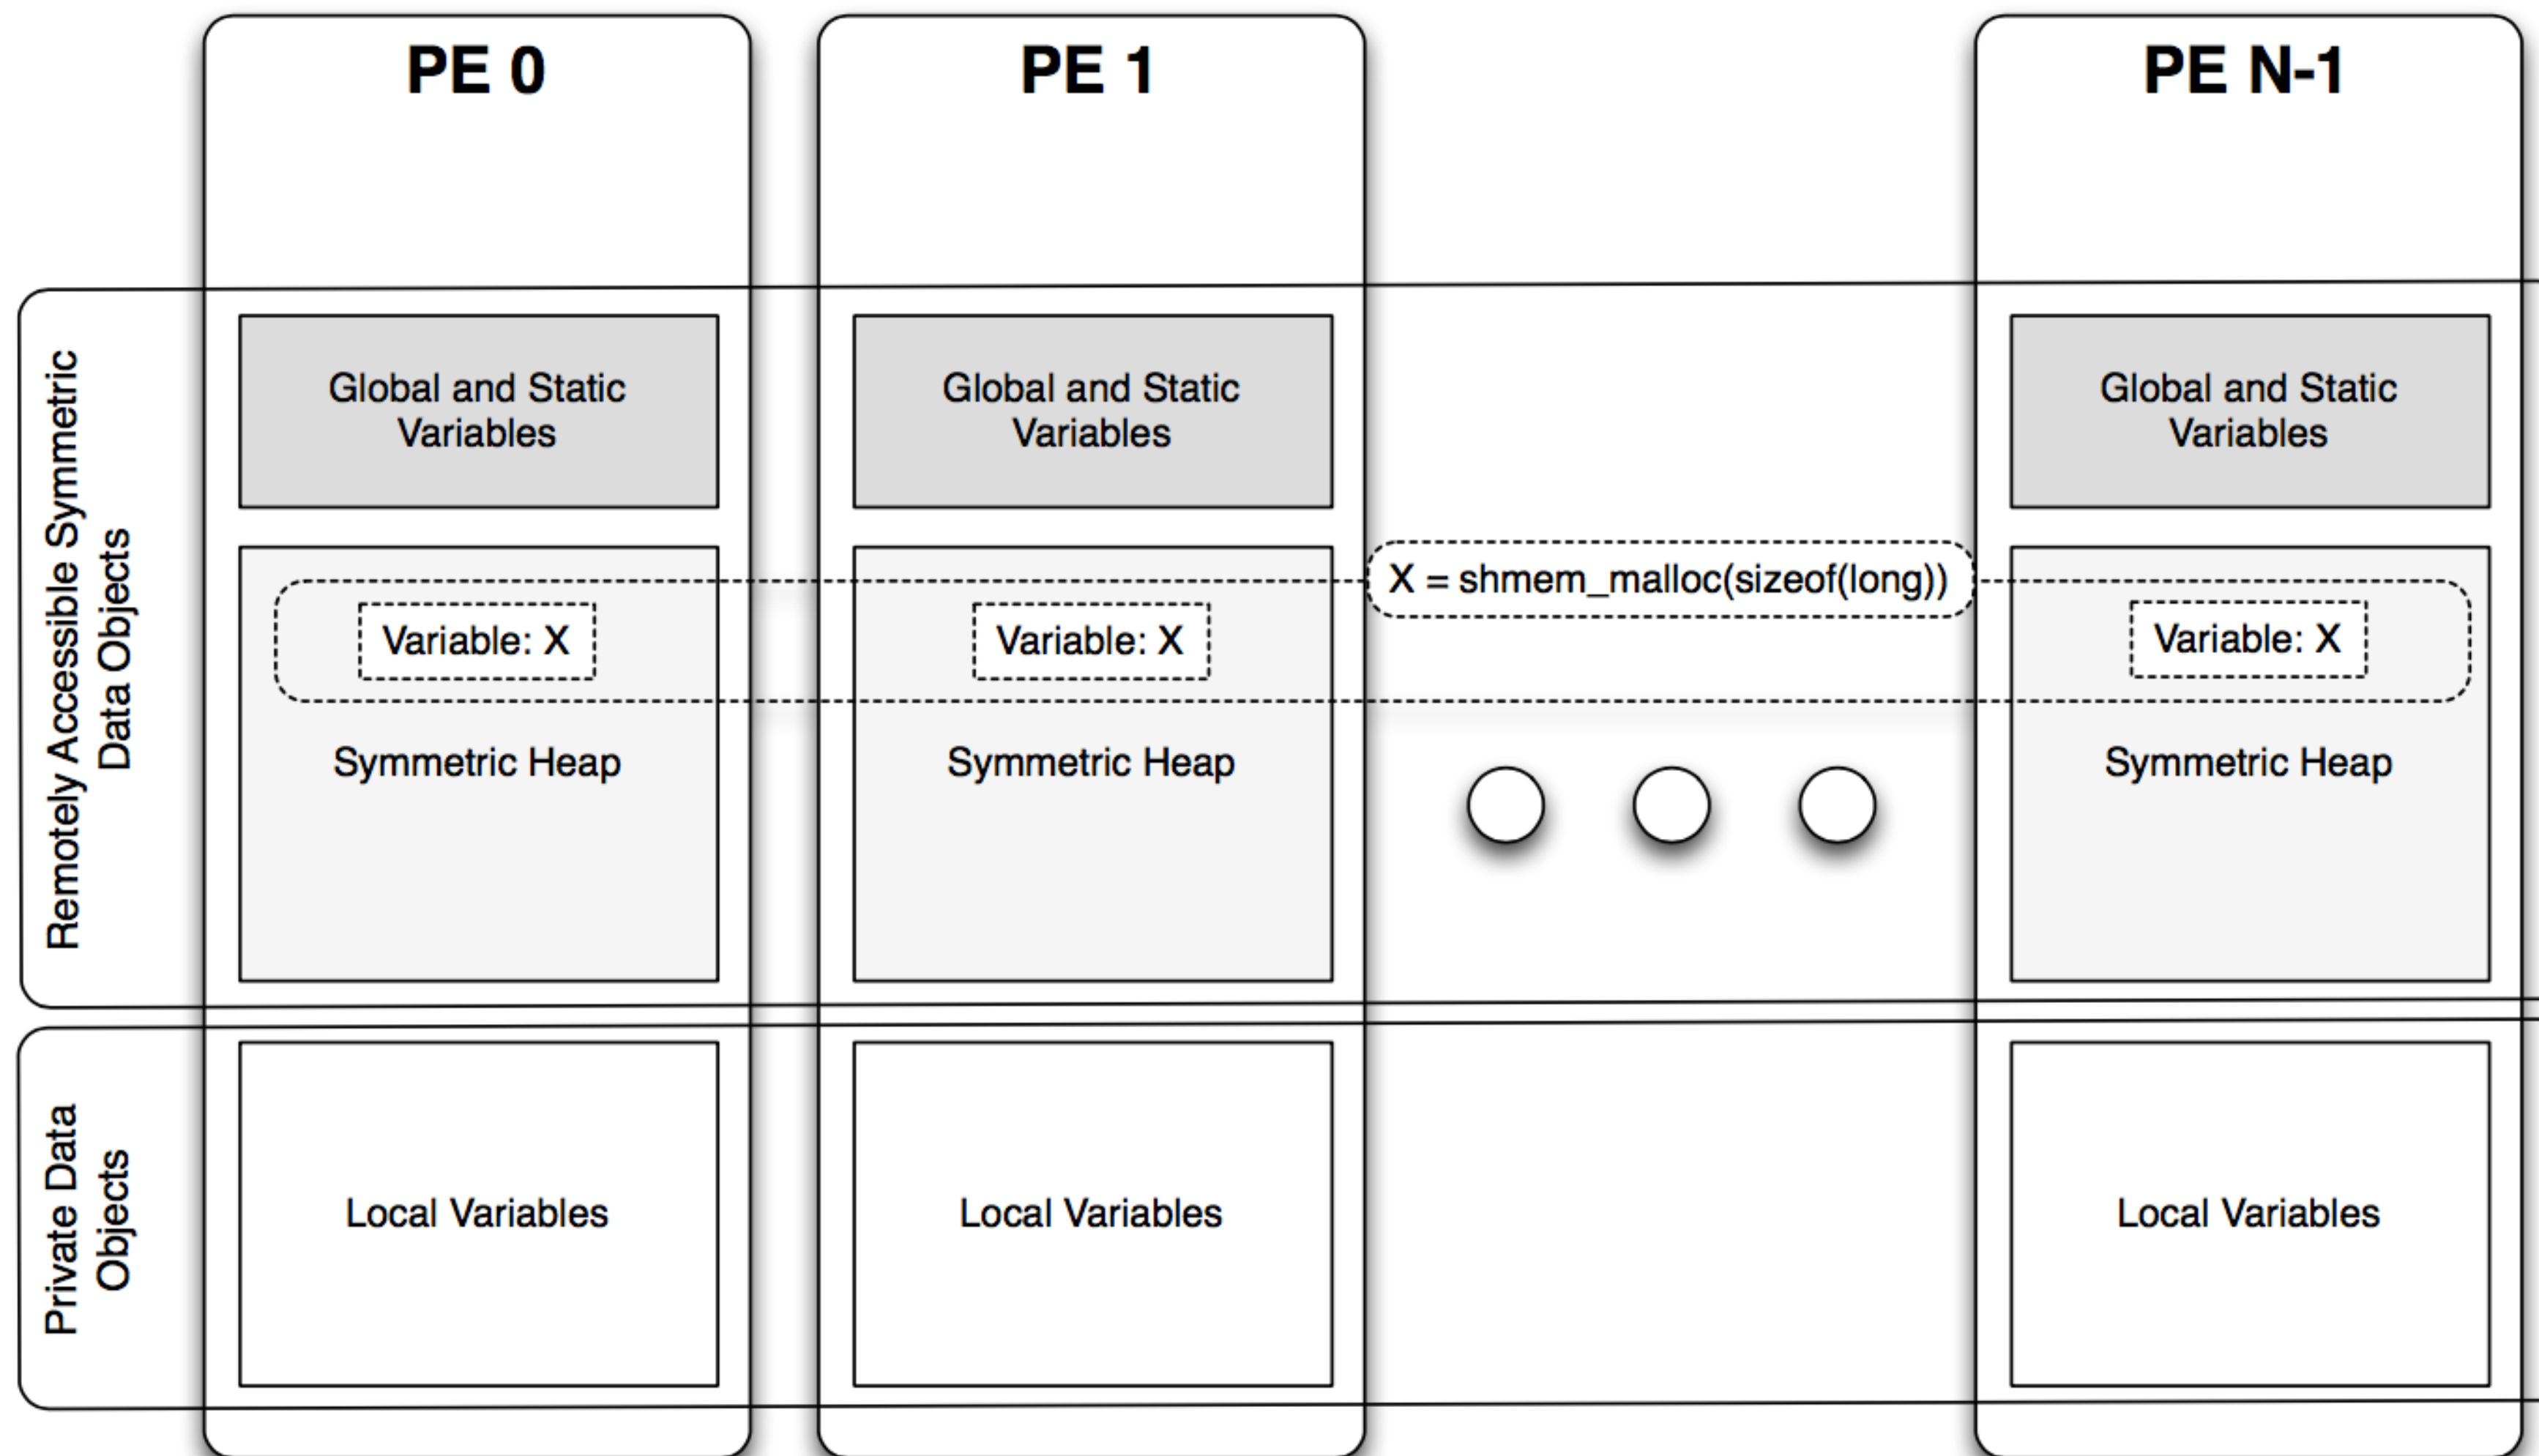
Libraries

- UPC++
- **OpenSHMEM**
- Global Arrays
- **DASH**
- ...

<code>#include <libdash.h></code>	1
	2
<code>int main(int argc, char* argv[]) {</code>	3
<code>dash::init(&argc, &argv);</code>	4
	5
<code>int myid = dash::myid();</code>	6
<code>int size = dash::size();</code>	7
	8
<code>dash::array<int> key(size);</code>	9
	10
<code>if(myid==0) {</code>	11
<code>for(i=0; i<size; i++) key[i]=compute_key(...);</code>	12
<code>}</code>	13
	14
<code>dash::barrier();</code>	15
	16
<code>cout<<"Hello_ from_ unit_"<<myid<<"_of_"</code>	17
<code><<size<<"_my_key_is"<<key[myid]<<endl;</code>	18
	19
<code>dash::finalize();</code>	20
<code>}</code>	21

DASH Hello World

K. Furlinger, C. Glass, J. Gracia, et al. DASH: Data Structures and Algorithms with Support for Hierarchical Locality. In Euro-Par Workshops, 2014.



OpenSHMEM “symmetric”
memory model

formerly know as Cray SHMEM (1993)

OPENSHPMEM

A OLD FASHIONED PGAS

- Perform computations in separate address spaces and explicitly pass data to and from different processes in the program
- API
 - Library Setup: init, query
 - **Symmetric** Data Object Management: allocation, deallocation, reallocation:
 - Remote Memory Access: **get** (private \leftarrow shared), **put** (shared \leftarrow private)
 - **Atomics**: swap, inc, add, was, fetch_add
 - **Synchronisation** and Ordering: **fence** (i.e. flush inflight ops), quiet (ensure other completion), **barrier**
 - Collective Communication: **broadcast**, **collection** (i.e. gather), **reduction**
 - Mutual Exclusion: **lock**, **testlock**, **unlock**

OPENSHPMEM HELLO WORLD!

```
#include <stdio.h>
#include <mpp/shmem.h>
int main (int argc, char **argv){

int me, npes;

start_pes (0); /*Library Initialization*/
me = _my_pe ();
npes = _num_pes ();
printf ("Hello World from node %d of %d\n", me, npes);
return 0;
}
```

OPENSHPMEM EXAMPLE

```
#include <stdio.h>
#include <shmem.h>

long pSync[_SHMEM_BARRIER_SYNC_SIZE];
int x = 10101;

int main(void)
{
    int i, me, npes;

    for (i = 0; i < _SHMEM_BARRIER_SYNC_SIZE; i += 1){
        pSync[i] = _SHMEM_SYNC_VALUE;
    }

    shmem_init();
    me = shmem_my_pe();
    npes = shmem_n_pes();

    if(me % 2 == 0){
        x = 1000 + me;
        /*put to next even PE in a circular fashion*/
        shmem_int_p(&x, 4, (me+2)%npes);
        /*synchronize all even pes*/
        shmem_barrier(0, 1, (npes/2 + npes%2), pSync);
    }
    printf("%d: x = %d\n", me, x);
    return 0;
}
```

CHAPEL LIST WALK

```
class Node {
  var data: real;
  var next: Node;
}

// List init (seq, with remote operations)
var head = new Node(0);
var current = head;
for i in 1..numLocales-1 do
  on Locales[i] {
    current.next = new Node(i);
    current = current.next;
  }

// List walk (seq)

current = head;
while current {
  writeln("node with data = ", current.data, " on locale
", current.locale.id);
  current = current.next;
}
writeln();
```

```
// Data-driven List walk (work done in parallel
// Each locale WRITE its own data (own-compute)

current = head;
while current {
  on current {
    writeln("node with data = ", current.data, " on locale
", here.id);
    current = current.next;
  }

// Deallocate (seq)

current = head;
while current {
  on current {
    var ptr = current;
    current = current.next;
    delete ptr;
  }
}
```


PGAS

- Overcomes some DSM criticality
- Aiming at productivity for large scale
 - To overcome MPI — not succeeding yet
 - Most of them build on top or interoperate with MPI
- Designed for weak scalability (own-compute) → Data Parallelism
 - Not for streaming

COROUTINES: SEMANTICS IN SEARCH OF A SYNTAX

by M. Douglas McIlroy


Oxford University and

Bell Telephone Laboratories, Incorporated.

ABSTRACT: Unlike subroutines, coroutines may be connected, and reconnected, in nonhierarchical arrangements. Coroutines are particularly useful for generating and processing data streams. Semantics for coroutines are developed and examples are given.

Streams!

Coroutines: McIlroy, 1968




J. Schwartz

M. Engeler

McCarthy

Paterson

F. Allen



Bahrs

Luckham

Ershov

Milner

Miller

Igarashi

Warren

Dennis

Hoare

Symposium on Theoretical Programming Novosibirsk – 1972



Gita

Arvind

Dennis

Data Flow Workshop MIT Endicott House – 1977



IEEE STC DATAFLOW AND BEYOND

HTTP://DFSTC.CAPSL.UDEL.EDU



STC Founding Member

These are the people that made possible the creation of this IEEE Special Technical Community.




- | | |
|--|---|
| • Arvind
Massachusetts Institute of Technology, USA | • Mei Hong
Beijing Institute of technology, China |
| • David Abramson
The University of Queensland, Australia | • Nahid Emad
University of Versailles, France |
| • Georgi Gaydadjiev
Maxeler, Imperial College London, UK | • Nelson Amaral
University of Alberta, Canada |
| • Guang R. Gao
University of Delaware, USA | • R. Govindarajan
Indian Institute of Science, India |
| • Hirarki Kei
The University of Tokyo, Japan | • Roberto Giorgi
Università di Siena, Italy |
| • Hironori Kasahara
Waseda University, Japan | • Skevos Evripidou
University of Cyprus, Cyprus |
| • Jack B. Dennis
Massachusetts Institute of Technology, USA | • Stephane Zuckerman
Michigan Technological University, USA |
| • Jean-Luc Gaudiot
University of California, Irvine, USA | • Vivek Sarkar
Rice University, USA |
| • Jin Hai
Huazhong University of Science and Technology, China | • Won Woo Ro
Yonsei University, South Korea |
| • Kuan-Ching Li
Providence University, Taiwan | • Zheng Weimin
Tsinghua University, China |
| • Mateo Valero
Universitat Politècnica de Catalunya, Spain | |



[Click here to see the STC founding members](#)

Members of our STC

- | | |
|-----------------------------------|--|
| • Dr. Mario Donato Marino | • Dr. Andres Marquez |
| • Dr Lorenzo Verdoscia | • Yoshitake Oki |
| • Marco Procaccini | • Professor Junqing Yu |
| • Dr. Chen Liu | • Professor wenguang chen |
| • Jin Yan | • Professor Sunita Chandrasekaran |
| • Tongsheng Geng | • Professor Dongrui Fan |
| • Dr. Sreepathi Pai | • Joshua Suetterlein |
| • Peng Qu | • Professor Bob Iannucci |
| • Dr. Albert Cohen | • Dr Long Zheng |
| • Shipeng Qi | • Dr. Hao Tu |
| • Dr. Toshiaki Kitamura | • Dr Tobias Becker |
| • Jose M Monsalve Diaz | • Siddhisanket Raskar |
| • Dr Erik Altman | • Professor Marco Aldinucci |
| • Dr. Karu Sankaralingam | • Professor Avi Mendelson |
| • Dr. Jidong Zhai | • Professor Xiaoming Li |
| • Professor Eduardo Juarez | • Dr. Xu Tan |



- Data Stream Processing (DaSP)

- Real-time processing of continuous data streams
- Processed on-the-fly with stringent Quality of Service (QoS)
- Potentially irregular flows of data must be timely processed
 - detect anomalies, real-time incremental responses, etc.



Stream processing applications

D. Turaga, H. Andrade, B. Gedik, et al. Design principles for developing stream processing applications. Softw. Pract. Exper. 2010; 40:1073–1104

INTERSECTING DIFFERENT CS AREAS

- Control Systems (and CPS)
 - Including network processors, FPGAs, etc
- Parallel Computing (shared memory)
 - StreamIt (2006), TBB, FastFlow (2009), RaftLib (2013), StreamBox (2016), ...
- BigData Analytics
 - Lambda, Kappa architectures
 - Apache + $x \in \{\text{Kafka, Flink, Storm, Apex, Spark, ... and counting}\}$
 - Tensorflow (2015)

[Download](#)[Libraries ▾](#)[Documentation ▾](#)[Examples](#)[Community ▾](#)[Developers ▾](#)[Apache Software Foundation ▾](#)[SQL and DataFrames](#)[Spark Streaming](#)[MLlib \(machine learning\)](#)[GraphX \(graph\)](#)[Third-Party Projects](#)

Spark Streaming makes it easy to build scalable fault-tolerant streaming applications.

Ease of Use

Build applications through high-level operators.

Spark Streaming brings Apache Spark's [language-integrated API](#) to stream processing, letting you write streaming jobs the same way you write batch jobs. It supports Java, Scala and Python.


```
TwitterUtils.createStream(...)
  .filter(_.getText.contains("Spark"))
  .countByWindow(Seconds(5))
```

Counting tweets on a sliding window

Fault Tolerance

Stateful exactly-once semantics out of the box.

Spark Streaming recovers both lost work and operator state (e.g. sliding windows) out of the box without any extra code on your part



Latest News

Spark 2.2.0 released (Jul 11, 2017)

Spark 2.1.1 released (May 02, 2017)

Spark Summit (June 5-7th, 2017, San Francisco) agenda posted (Mar 31, 2017)

Spark Summit East (Feb 7-9th, 2017, Boston) agenda posted (Jan 04, 2017)

[Archive](#)[Download Spark](#)


Built-in Libraries:

[SQL and DataFrames](#)[Spark Streaming](#)[MLlib \(machine learning\)](#)[GraphX \(graph\)](#)[Third-Party Projects](#)

all
ch

Counting tweets on a sliding window

Down



Built-in Libraries

SQL and DataF

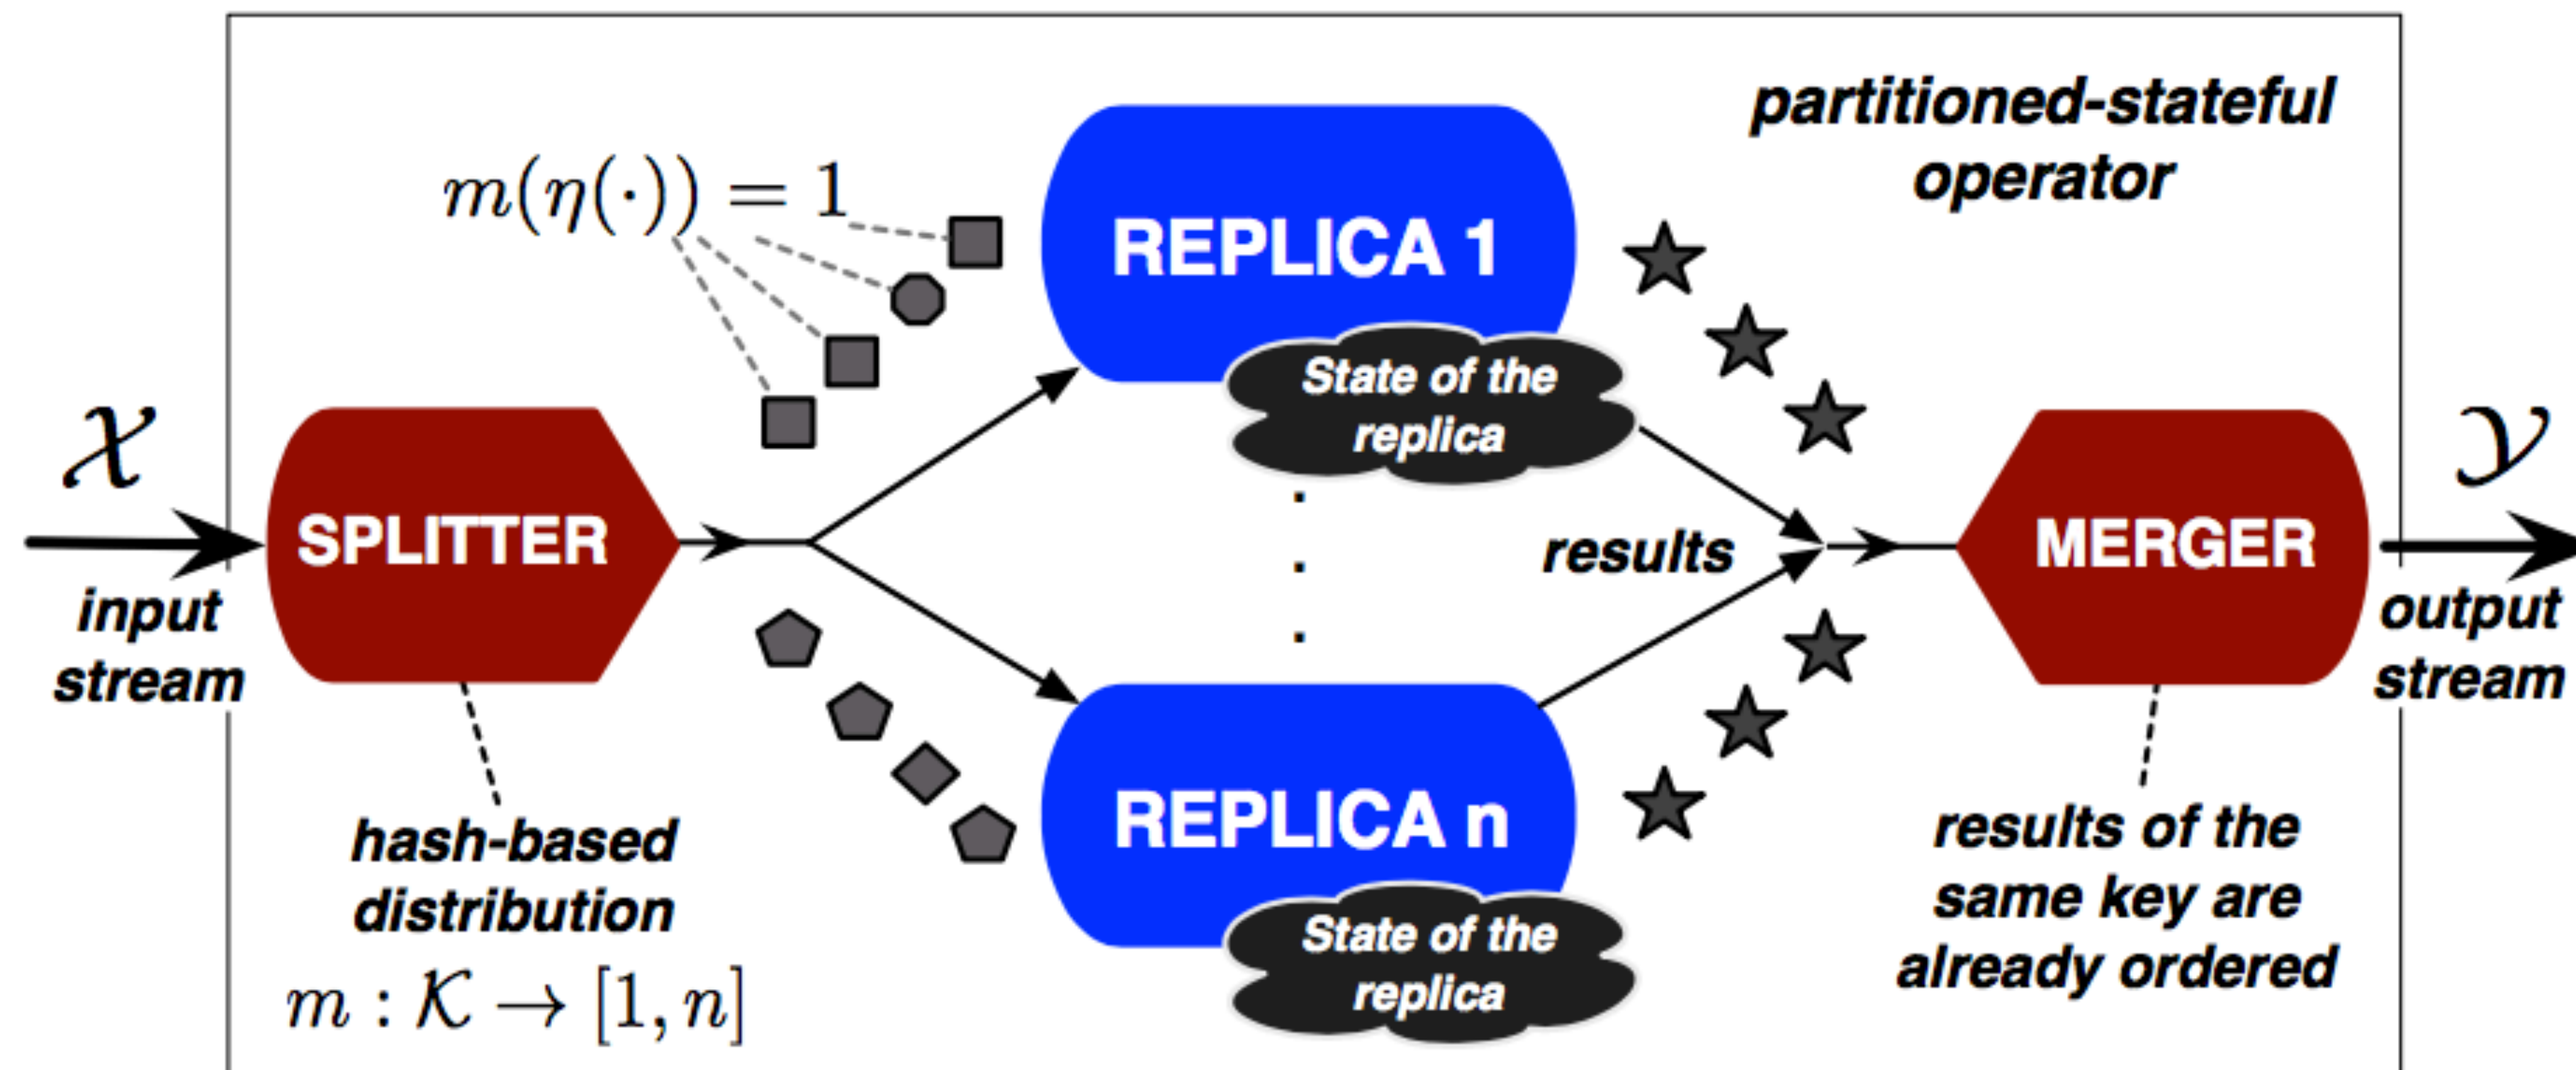
Spark Streaming

MLlib (machine

GraphX (graph)

Third-Party Pro


STREAMING: CORE PARADIGM




Tiziano De Matteis and Gabriele Mencagli. Keep Calm and React with Foresight: Strategies for Low- Latency and Energy-Efficient Elastic Data Stream Processing, PPoPP 2016.

operator



control actors



Boolean actors



decider



LABORATORY FOR
COMPUTER SCIENCE



MASSACHUSETTS
INSTITUTE OF
TECHNOLOGY

MIT/LCS/TM-61

FIRST VERSION OF A
DATA FLOW PROCEDURE
LANGUAGE

Jack B. Dennis

May 1975


545 TECHNOLOGY SQUARE, CAMBRIDGE, MASSACHUSETTS 02139

Figure 2. Node types for data flow programs.

NETWORK OF EXECUTORS VS TASK GRAPH


- Graph of tasks

- Dataflow — typically DAG
- Each node is a task
- E.g. a C++ object
- Problems: firing, scheduling, etc.



- Network of executors

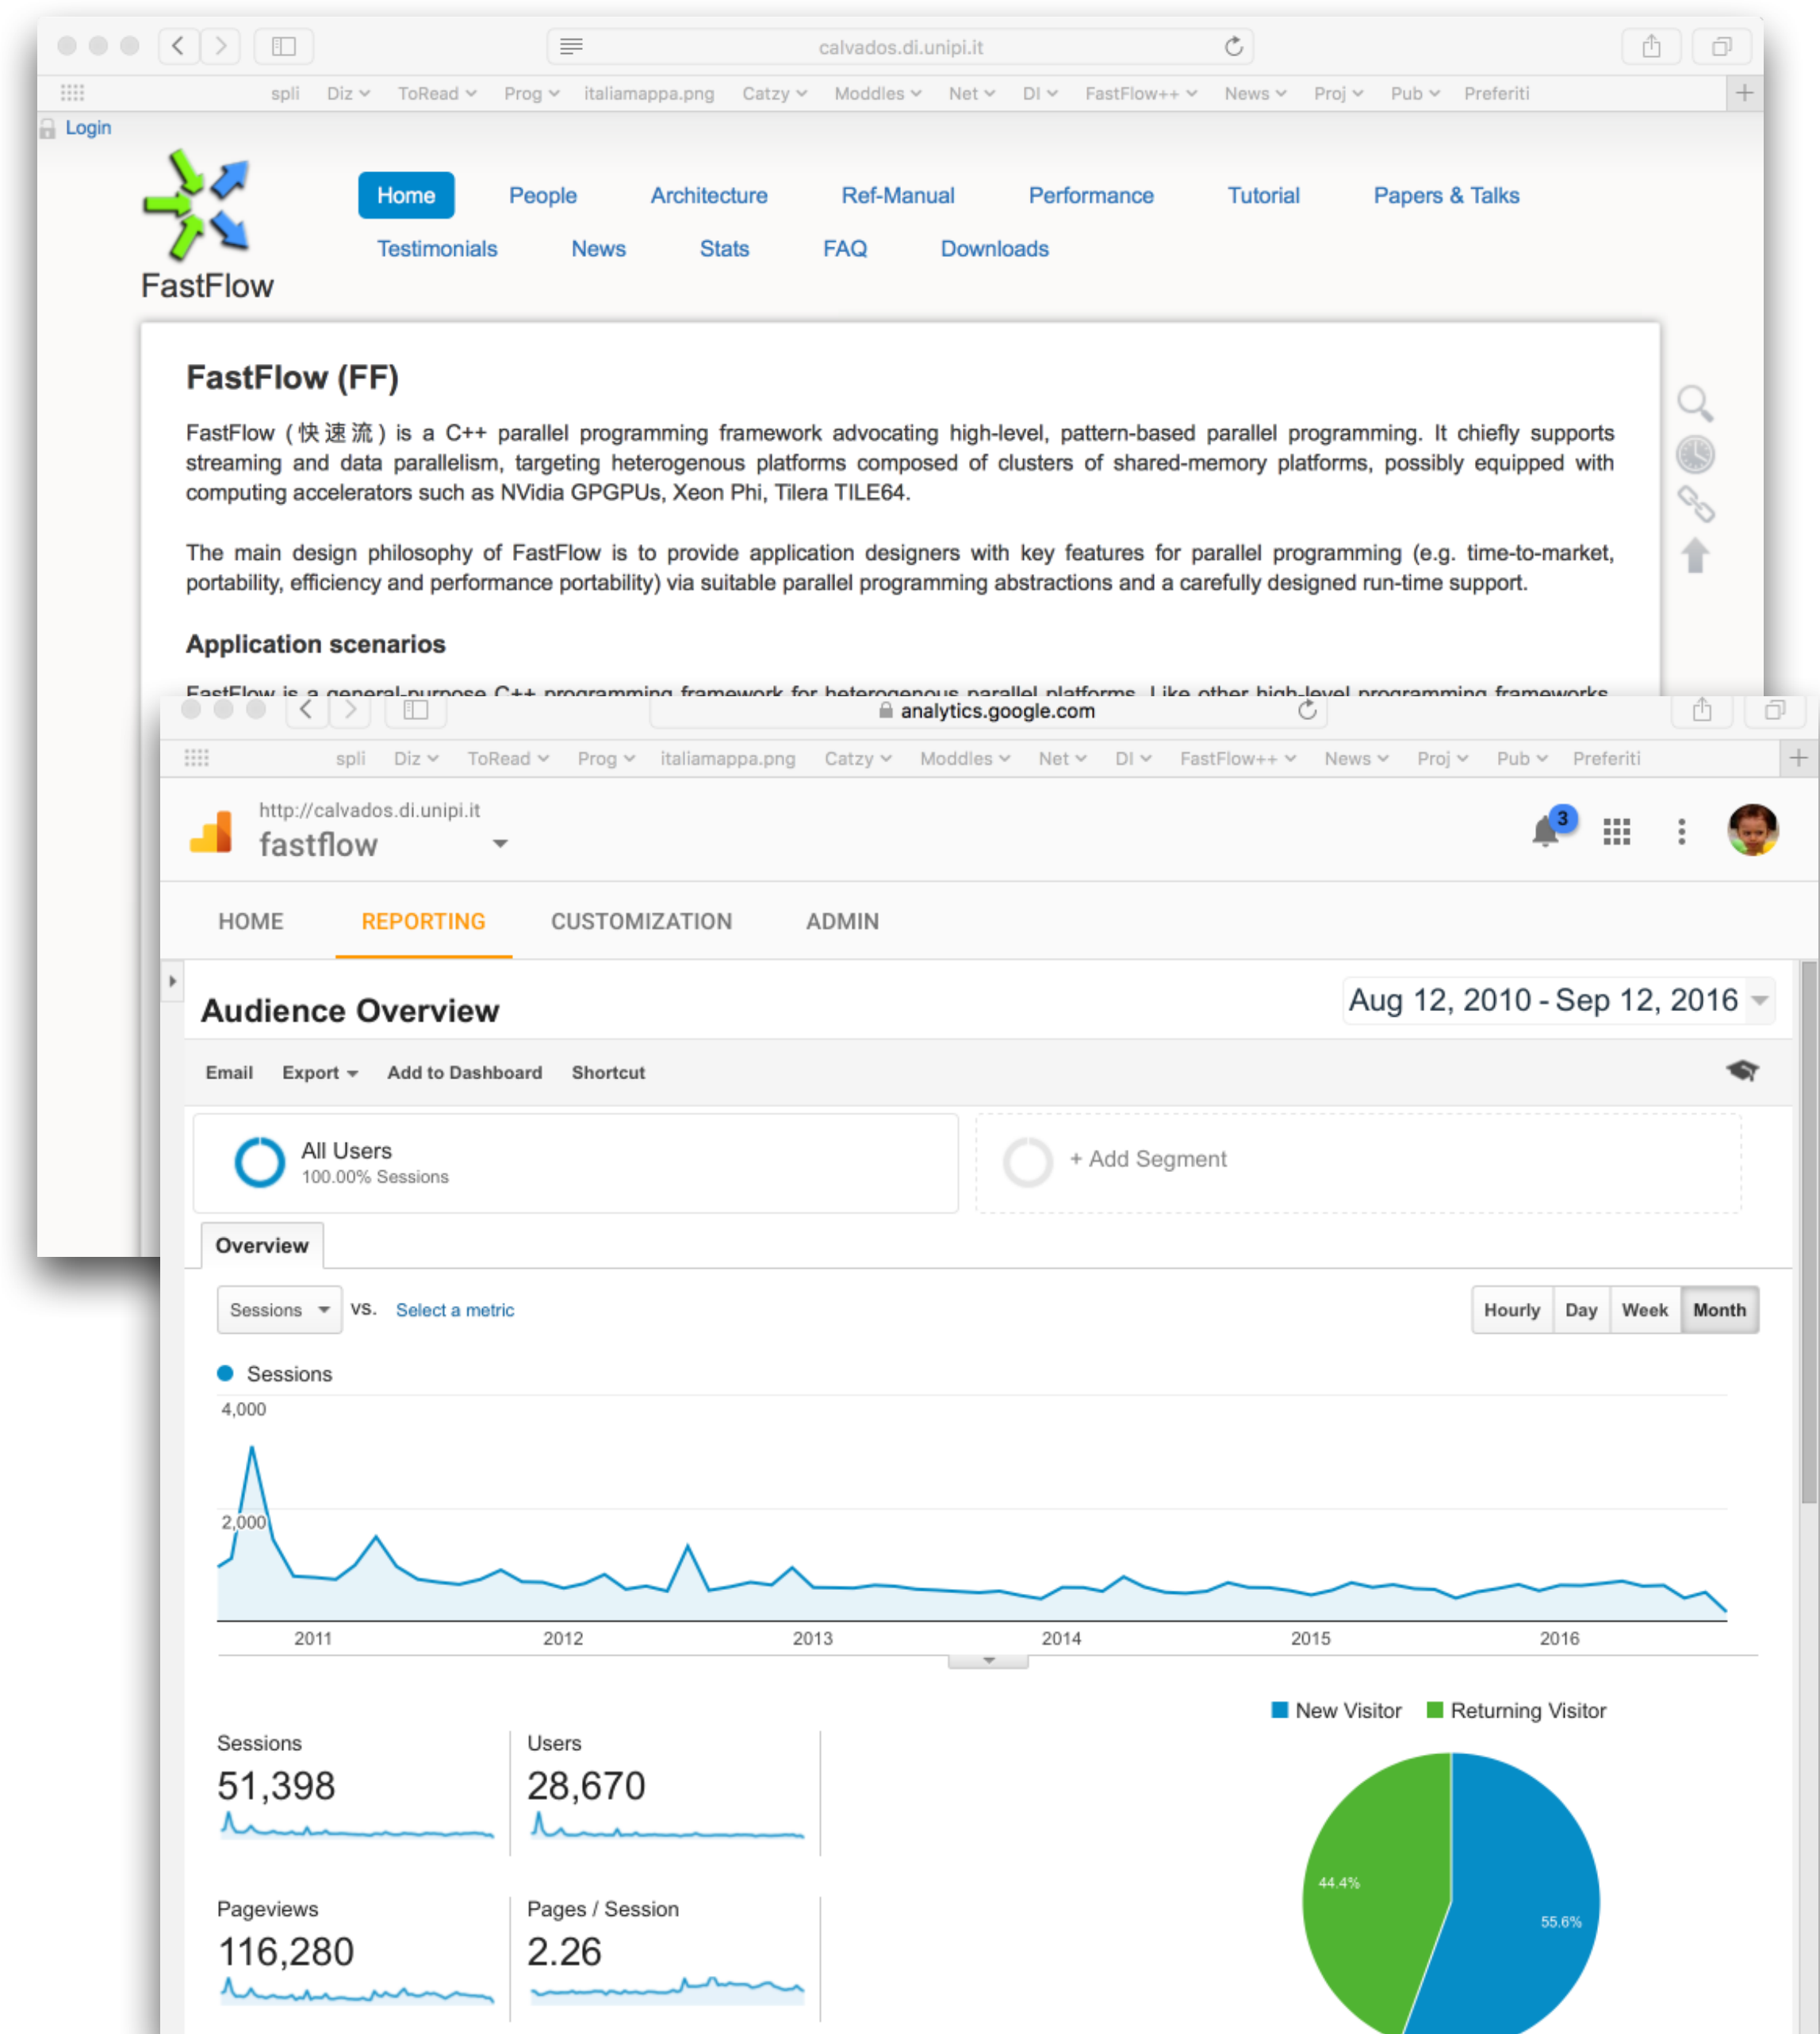
- "Controlflow" — typically cyclic graph
- E.g. threads or processes
- Problems: pinning, mapping, pooling, etc.




FastFlow

<http://mc-fastflow.sourceforge.net/>

- **Toreador** (EC-RIA, H2020, ICT-16-2015 big data): Trustworthy model-aware Analytics Data platform (2016, 36 months, total cost 6.5M €)
- **Rephrase** (EC-RIA, H2020, ICT-2014-1): Refactoring Parallel Heterogeneous Resource-Aware Applications – a Software Engineering Approach (2015, 36 months, total cost 3.5M €)
- **HyVar** (EC-RIA, H2020, ICT-2014-1): Scalable Hybrid Variability for Distributed Evolving Software Systems (2015, 36 months, total cost 2.8M €)
- **REPARA** (EC-STREP, 7th FP): Reengineering and Enabling Performance And power of Applications (2013, 36 months, total cost 3.5M €)
- **ParaPhrase** (EC-STREP, 7th FP): Parallel Patterns for Adaptive Heterogeneous Multicore Systems (2011, 42 months, total cost 4.2M €)
- **IBM Research** 3 faculty awards 2015 (50K \$)
- **Noesis Solutions**: Machine learning for engineering 2015 (75K €)
- **A3CUBE Inc.**: FastFlow/PGAS with in memory fabric 2014
- **Nvidia Corp**: CUDA Research Center at University of Torino 2013



SWSR QUEUES + MEDIATORS



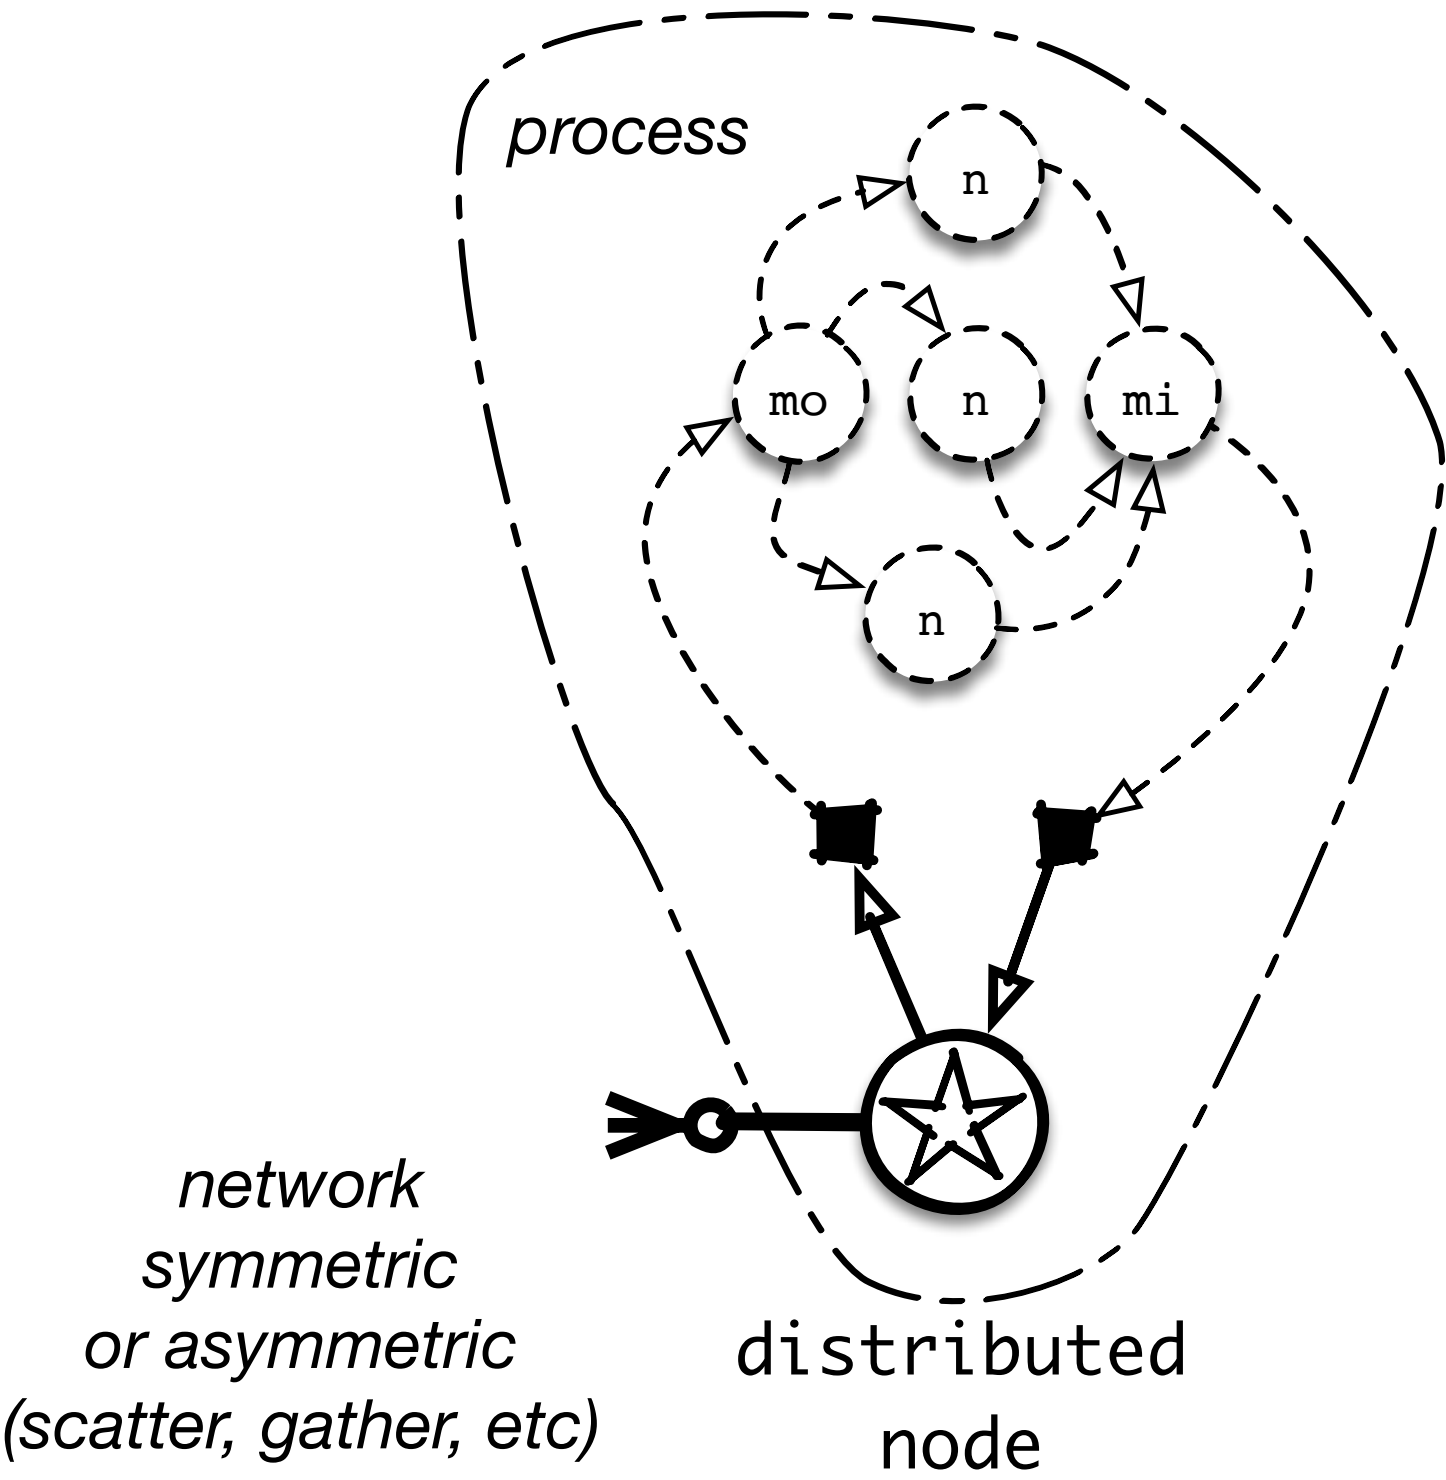




FF bound shmem FIFO channel
Single-Producer-Single-Consumer
lock-free fence-free queue



FF unbound shmem FIFO channel
Single-Producer-Single-Consumer
lock-free fence-free queue



FF (lock-free) distributed memory channel
RDMA or TCP



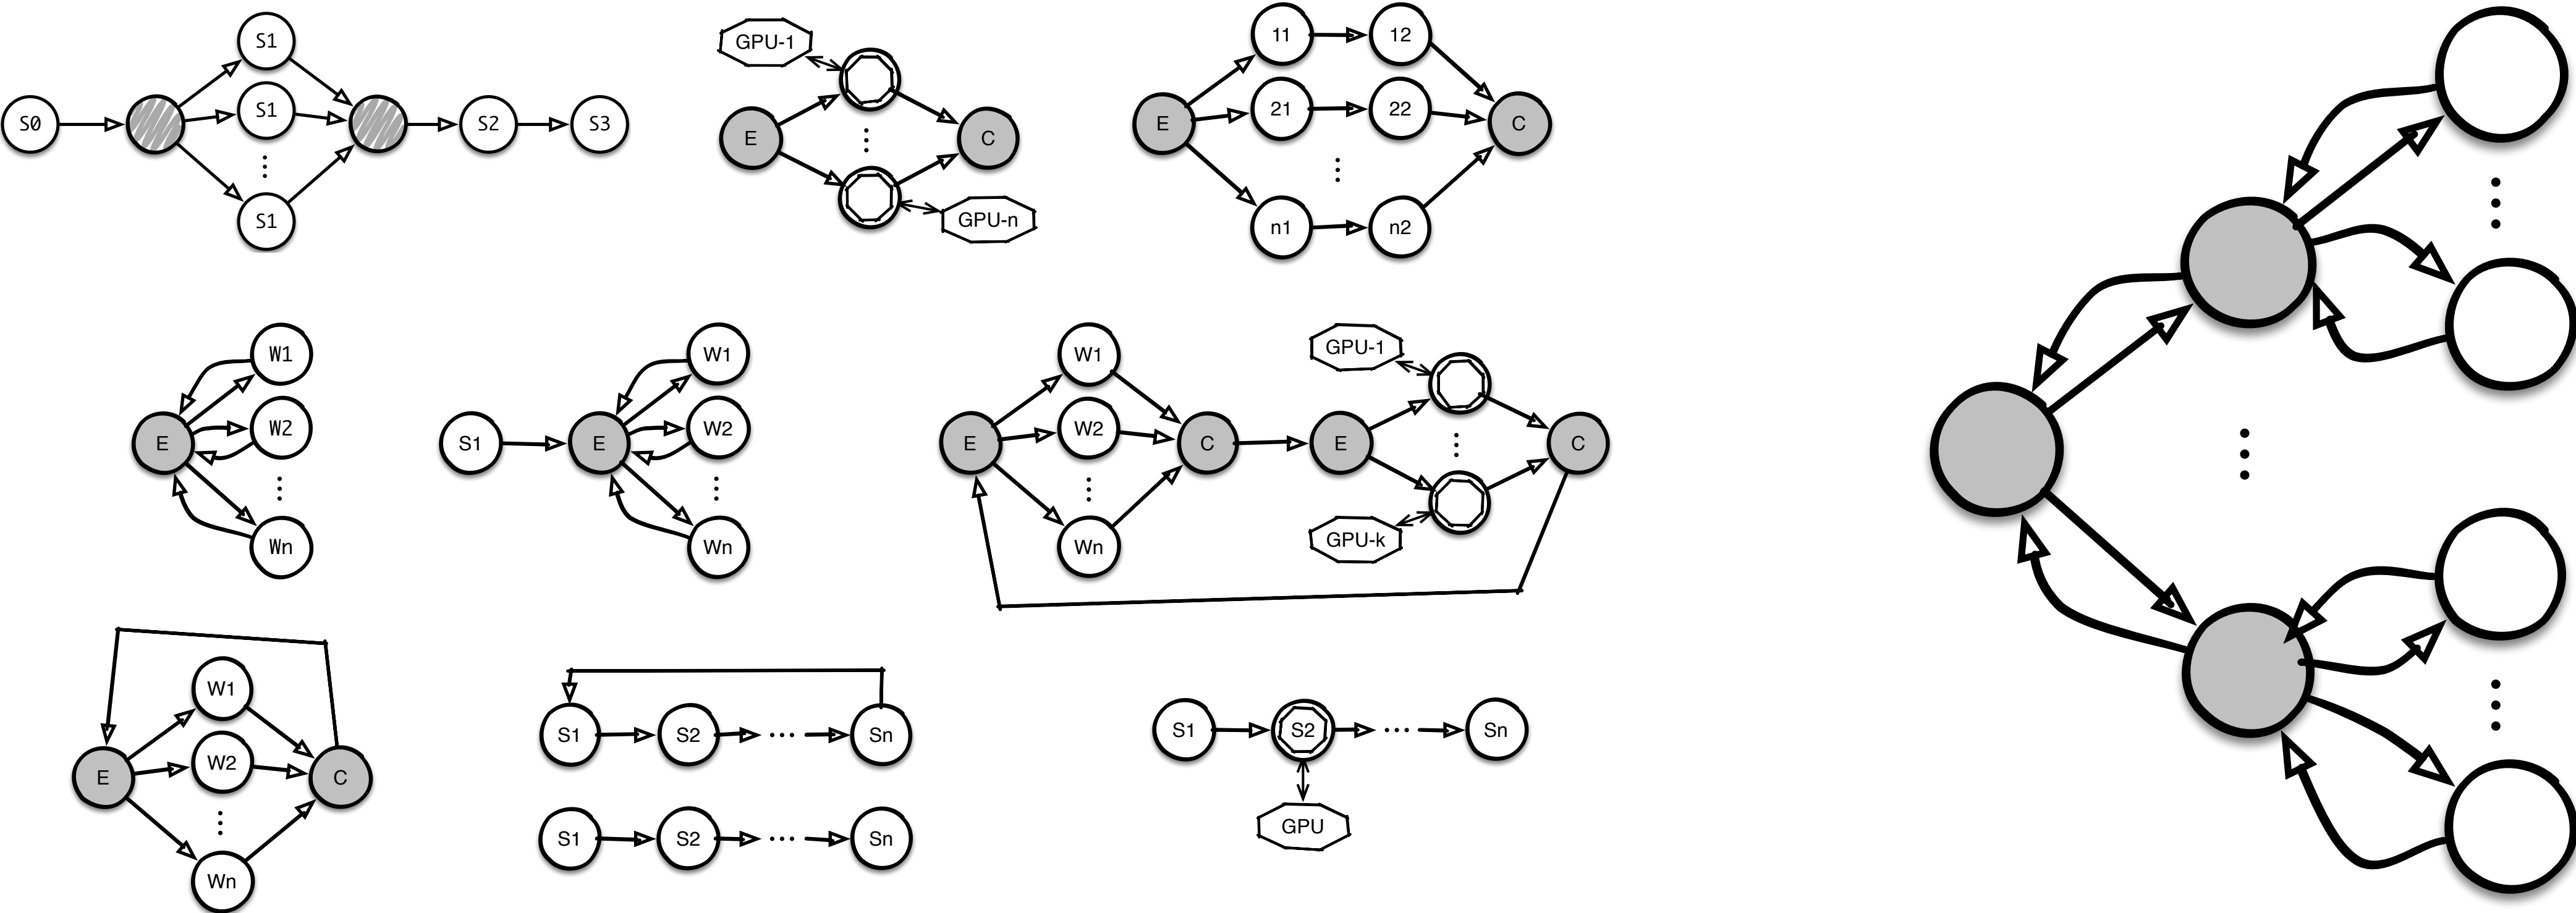
GENERATE THE NETWORK

TRUE DATA DEPENDENCIES MOVES ACROSS ARROWS

Composing via mediator guarantee
correctness (data races & deadlock freedom)

GENERATE THE NETWORK


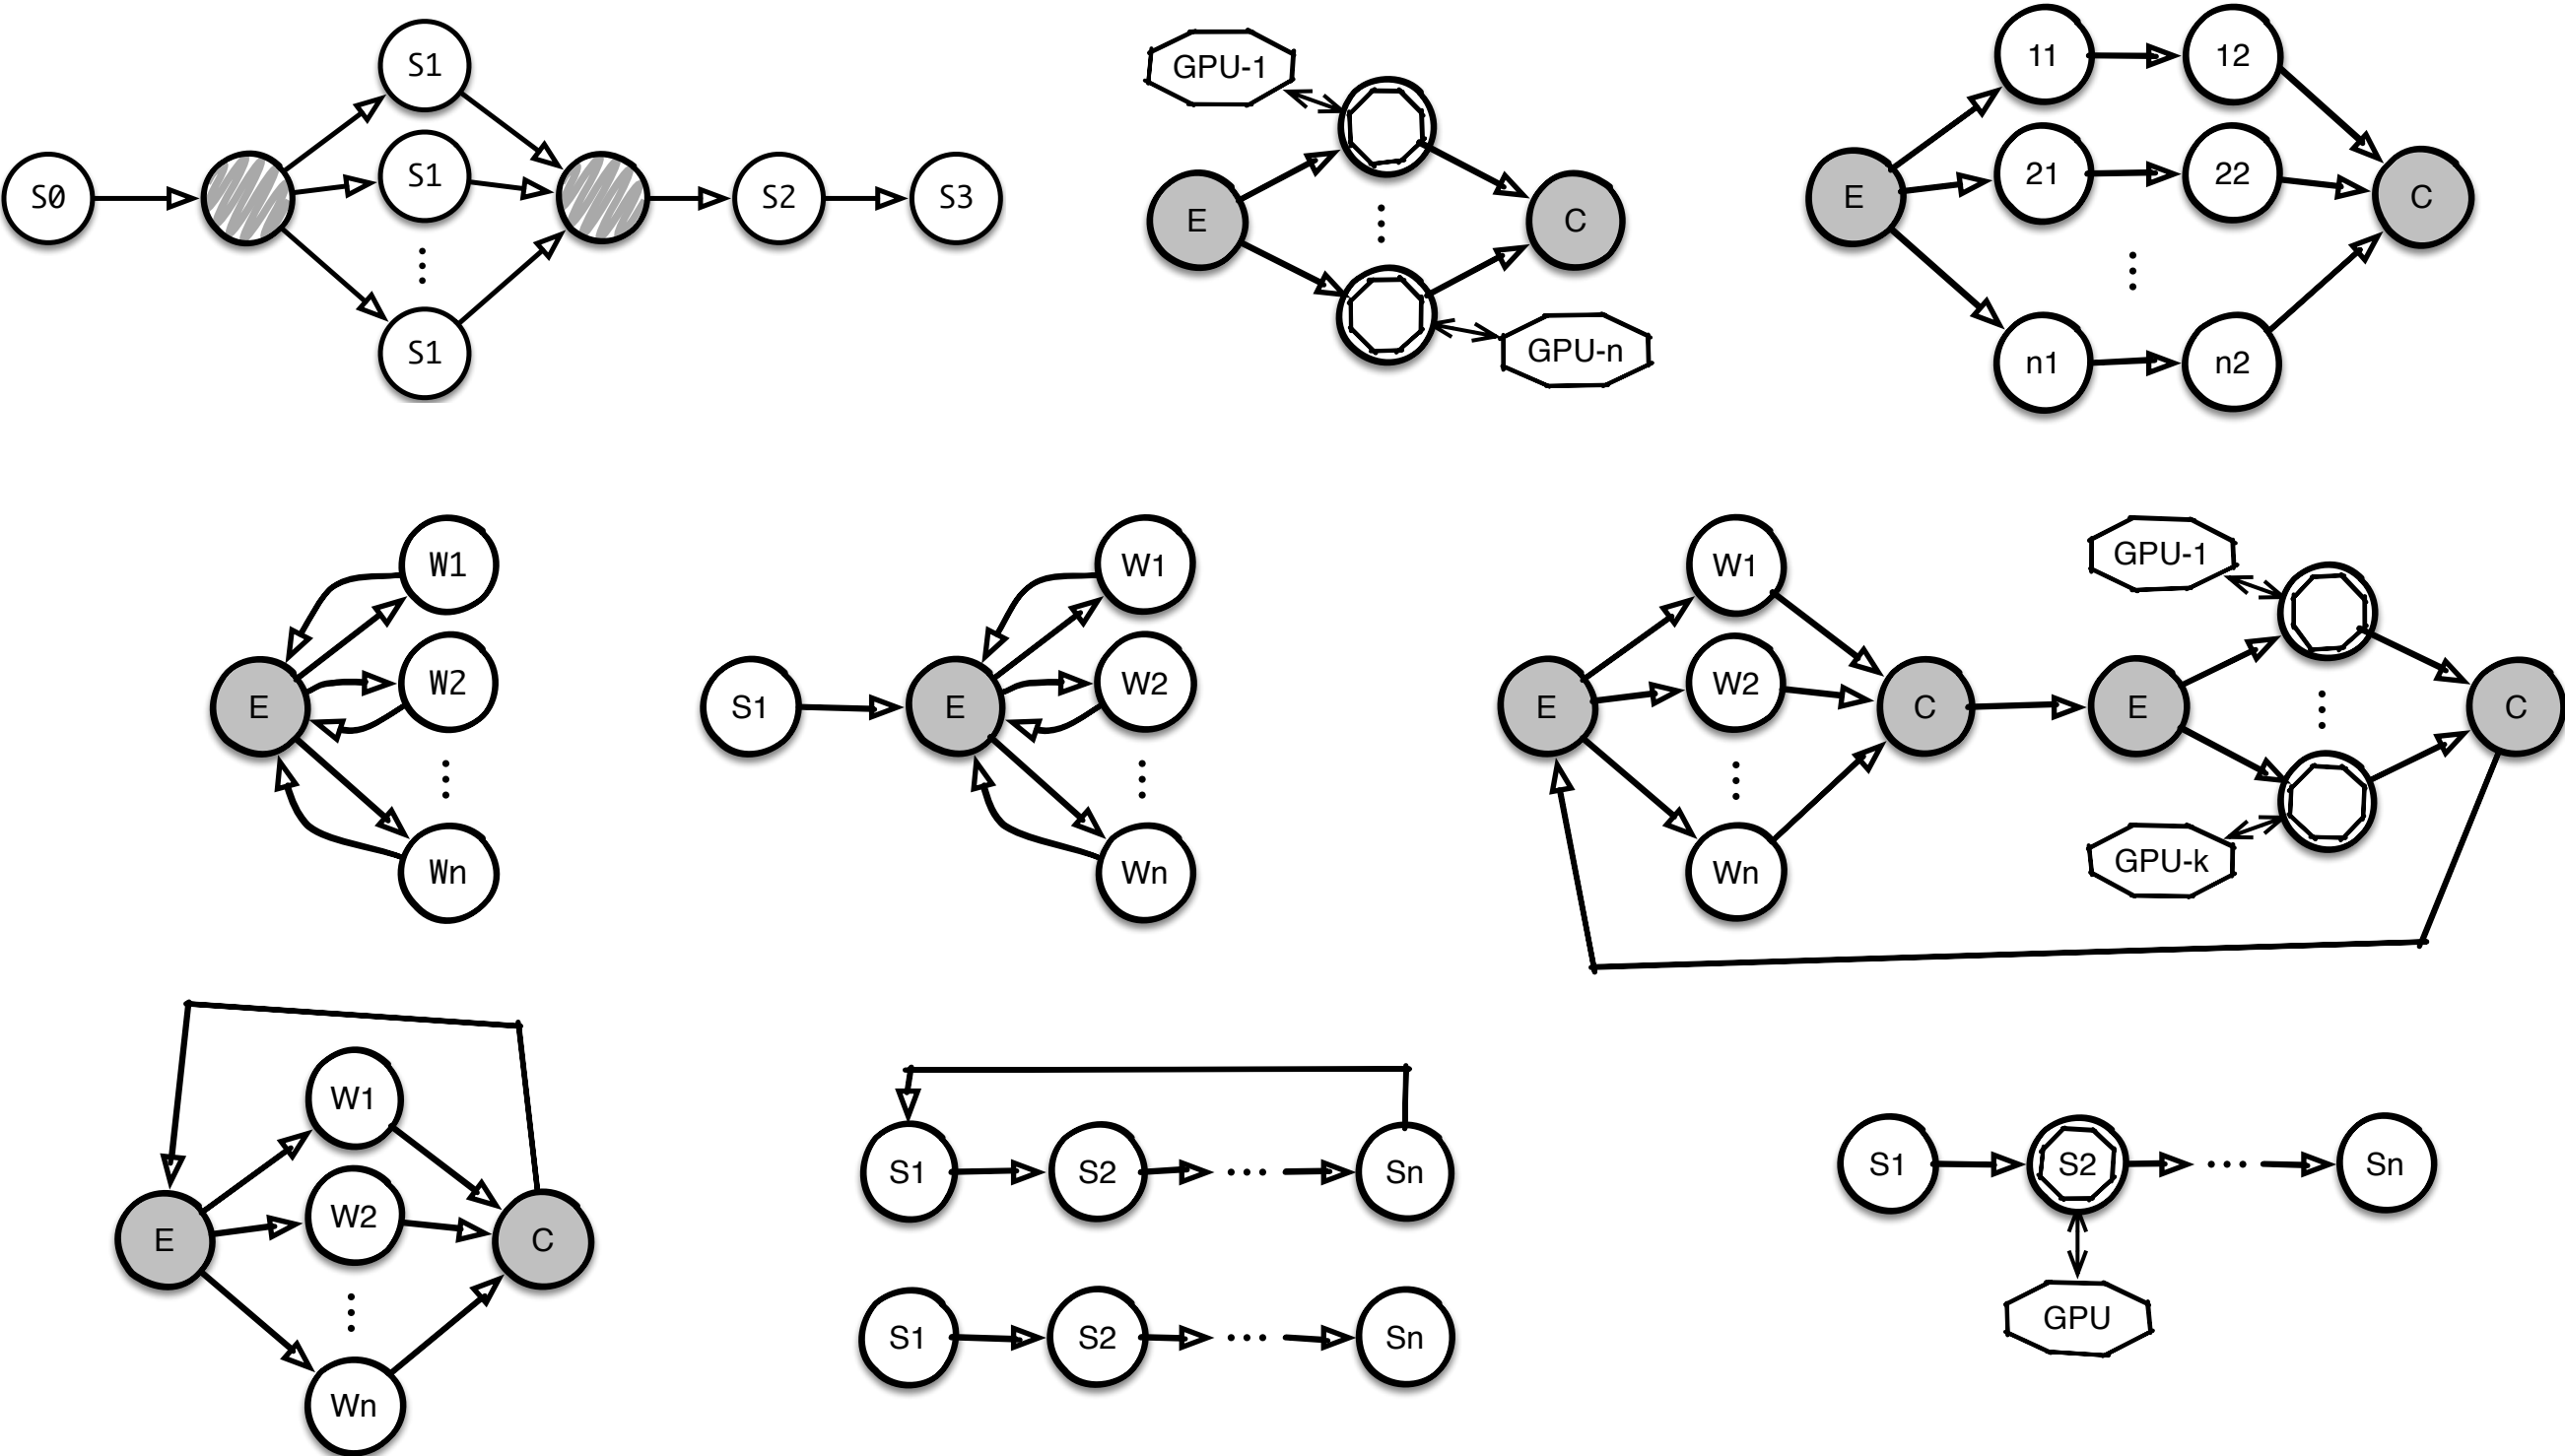
TRUE DATA DEPENDENCIES MOVES ACROSS ARROWS



Composing via mediator guarantee
correctness (data races & deadlock freedom)

GENERATE THE NETWORK

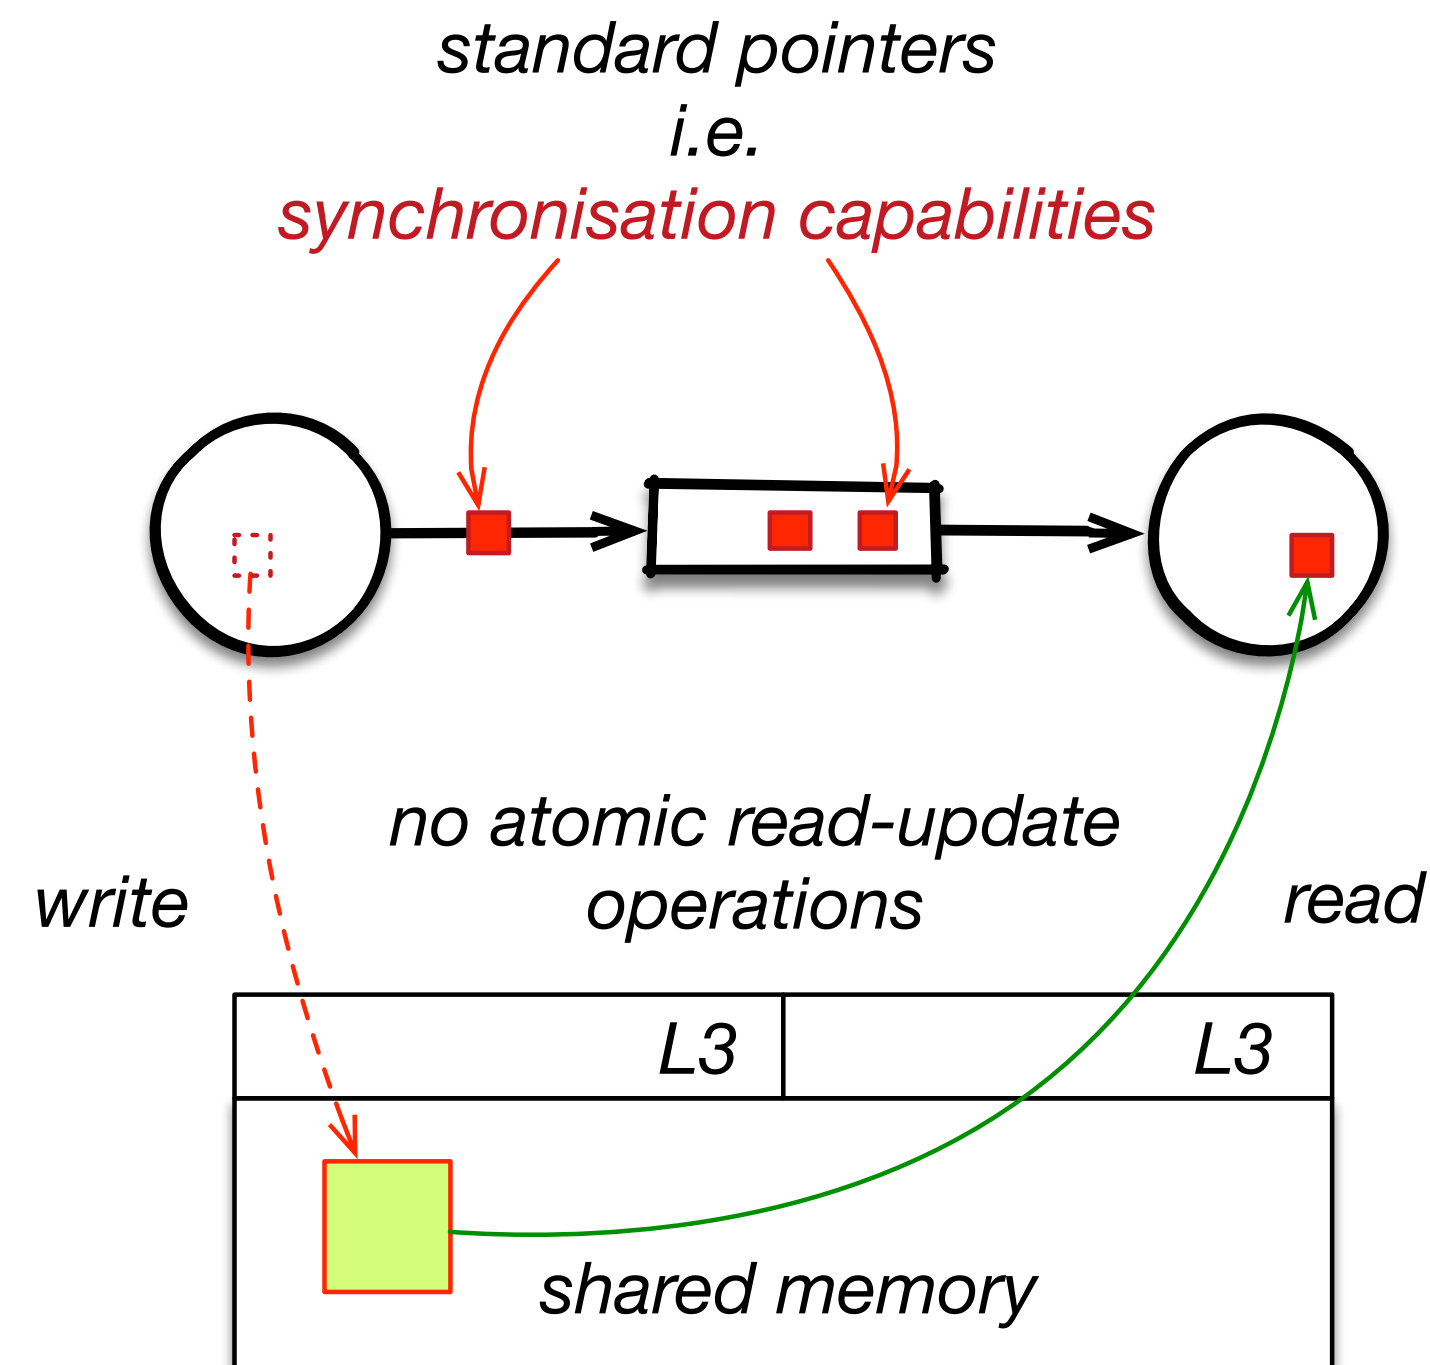
TRUE DATA DEPENDENCIES MOVES ACROSS ARROWS



Composing via mediator guarantee correctness (data races & deadlock freedom)

Possible data race

PROGRAMMING MODEL: SYNCHRONISATIONS HAPPEN BY WAY OF P2P DATA DEPENDENCIES (THUS NO ATOMICS ARE NEEDED)

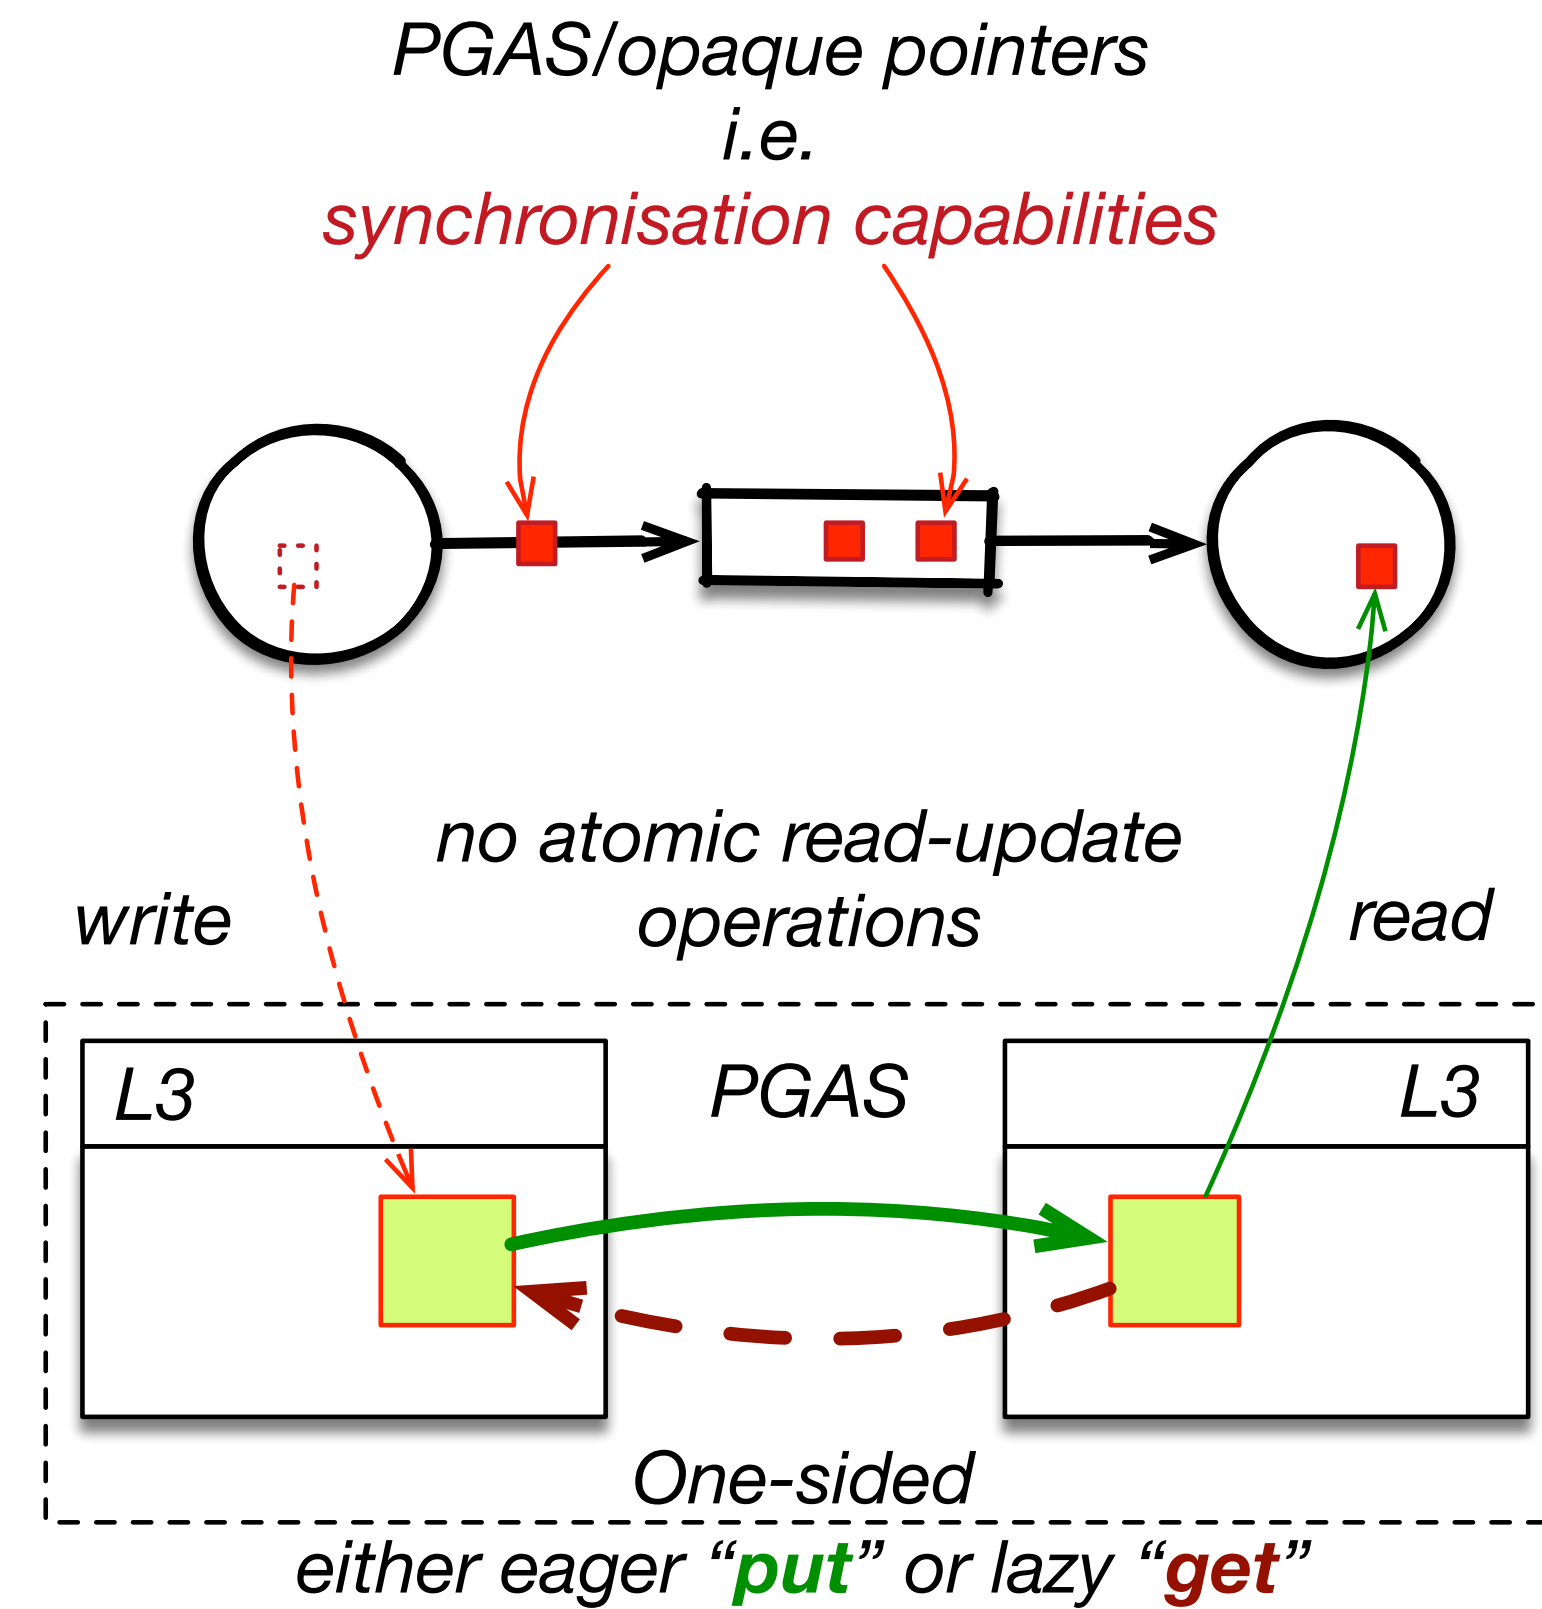
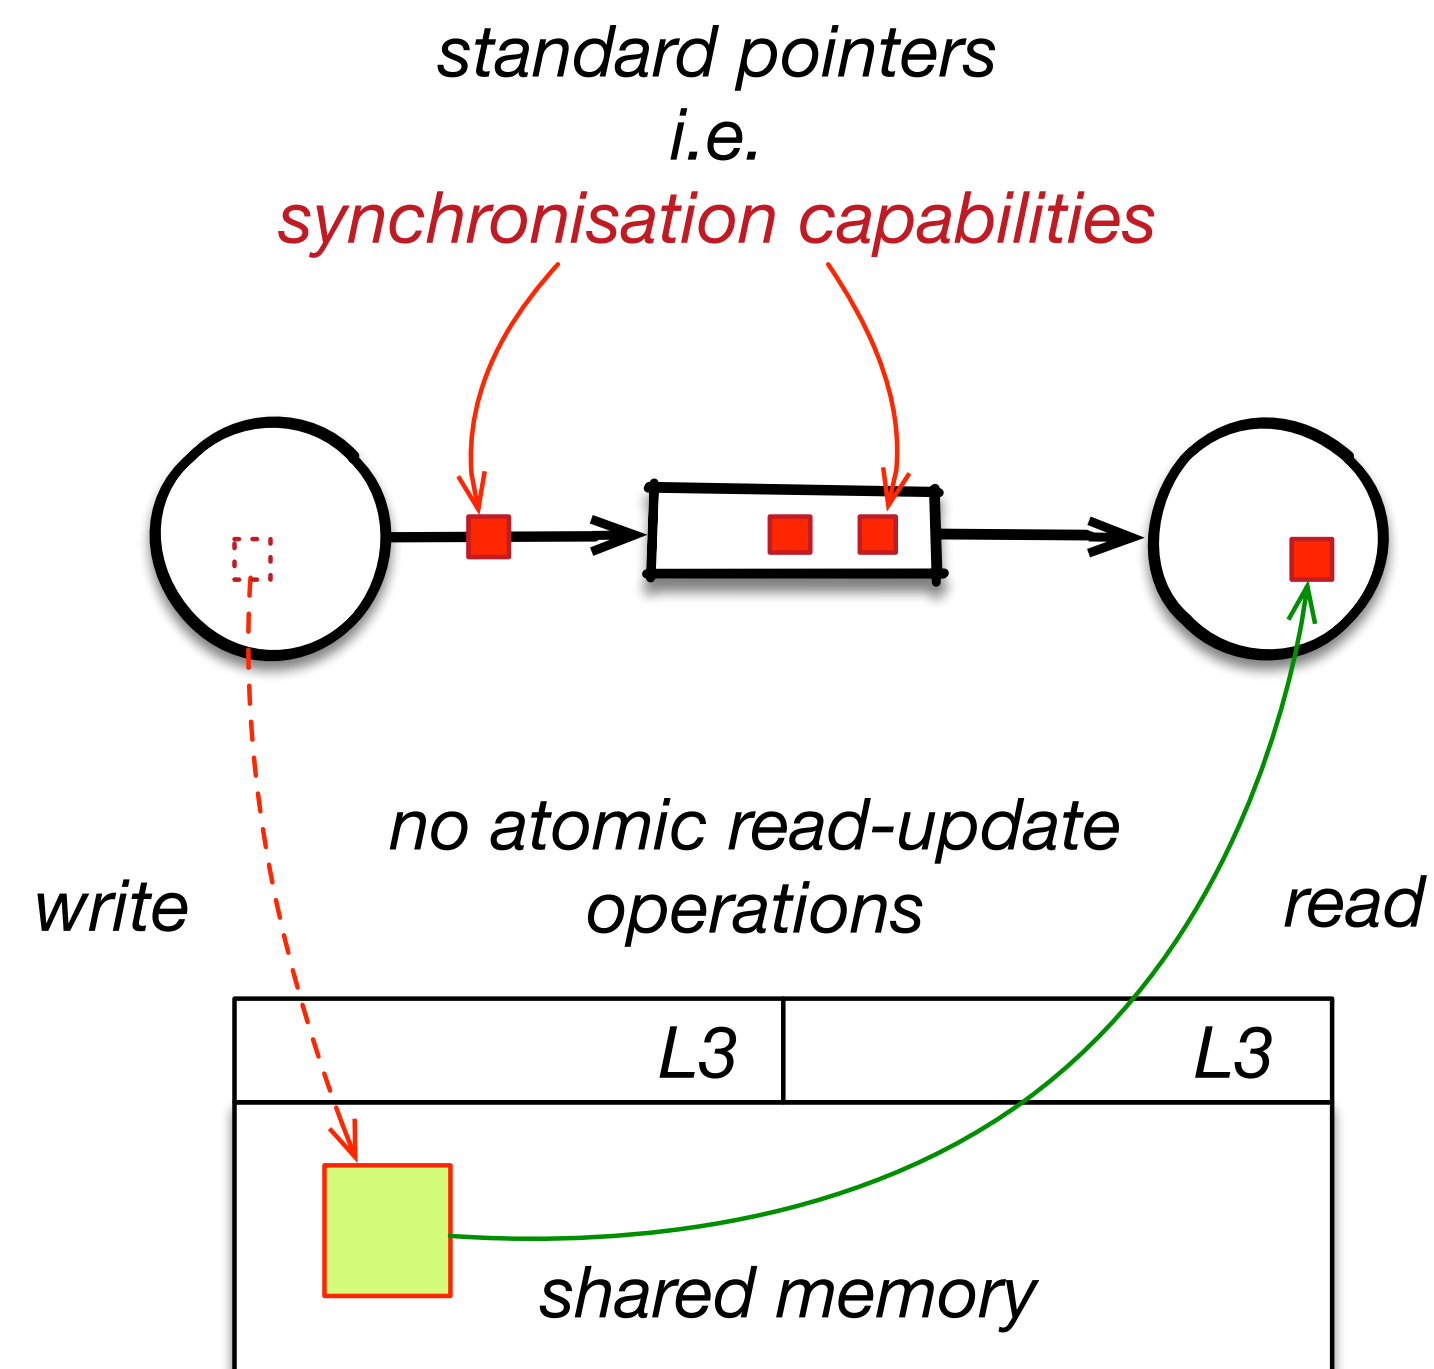


Synchronisations are in a message-passing style, but designers are not forced to think in a distributed way

No copies are needed, the memory fences are but asynchrony helps

Shared-memory cache-coherent or non-coherent multicore

PROGRAMMING MODEL: SYNCHRONISATIONS HAPPEN BY WAY OF P2P DATA DEPENDENCIES (THUS NO ATOMICS ARE NEEDED)




Shared-memory cache-coherent or
non-coherent multicore

Distributed GAM

BASED ON SINGLE-WRITER-SINGLE-READER FIFO

- Does not require atomic
 - No fence under TSO, WriteFence under WO
 - *J. Giacomoni et al. Fastforward for efficient pipeline parallelism: a cache-optimized concurrent lock-free queue. PPOPP 08*
 - *M. Aldinucci et al. An Efficient Unbounded Lock-Free Queue for Multi-core Systems. Euro-Par 2012*
- Enough to support Producer-Consumer
 - Inherently asynchronous
 - Powerful enough to build a general purpose parallel programming model


FastFlow unbound queue
core-to-core message latency
Xeon E7-4820 @2.0GHz Sandy Bridge



Bringing Parallel Patterns out of the corner: the P³ARSEC Benchmark Suite



A **Benchmark Suite** for
parallel patterns-based
applications



Parallel Pattern design
of 12 out of 13 **PARSEC**
benchmark applications




Implementation with
FastFlow and **SkePU2**
publicly available

FASTFLOW

<http://calvados.di.unipi.it/paragroup/>

- Compared in performance and features against:
openMP, TBB, Pthreads, OpenSs, MPI, ...
- M. Aldinucci, M. Meneghin, and M. Torquati,
“Efficient Smith-Waterman on multi-core with
FastFlow,” Euromicro PDP 2010.
- 80+ papers from 2010

Comparison over **3 different shared memory multicore architectures**
(Intel Xeon, Intel Xeon Phi, IBM Power 8) using different implementations




*Normalized with respect to Pthreads, averaged over all the benchmarks (the lower the better)

Detailed results and comparison with additional frameworks can be found in:

D. De Sensi, T. De Matteis, M. Torquati, G. Mencagli and M. Danelutto,
“Bringing Parallel Patterns out of the corner: the P³ARSEC Benchmark Suite”
Under Review in ACM Transactions on Architecture and Code Optimization

M. Danelutto, T. De Matteis, D. De Sensi, G. Mencagli, and M. Torquati,
“P³ARSEC: towards parallel patterns benchmarking”
in Proceedings of the 32nd annual ACM Symposium on Applied Computing (SAC 2017)

WINDOWED STREAM PROCESSING



WINDOWED STREAM PROCESSING

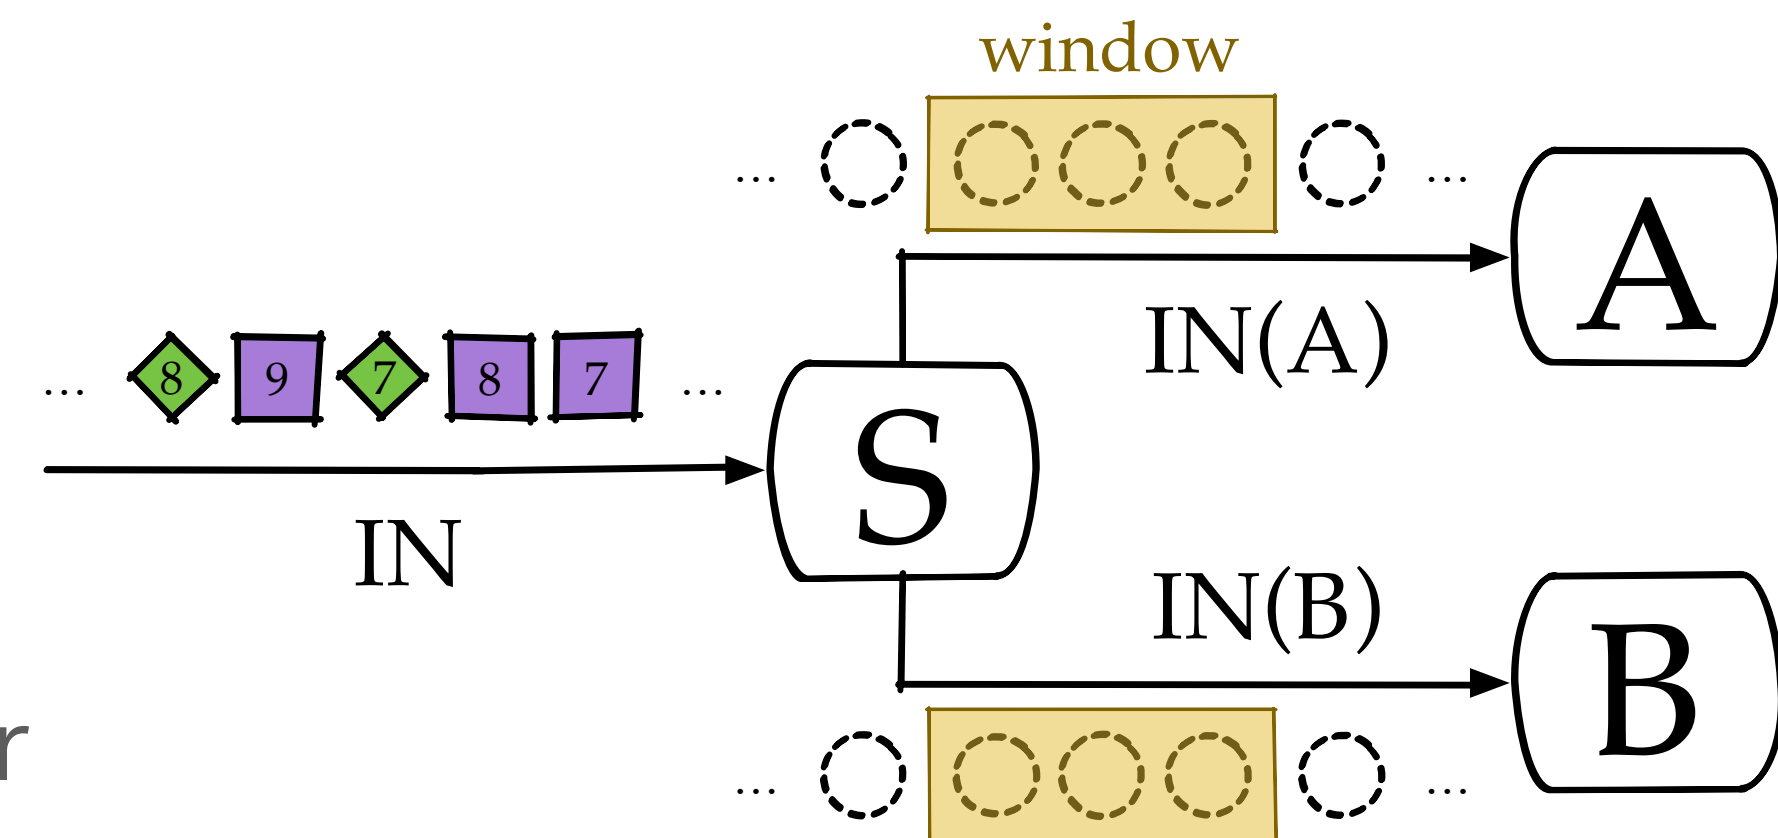
- Windows approximate infinite stream history

- tuple significance is often time-decaying
- only the most recent tuples are kept

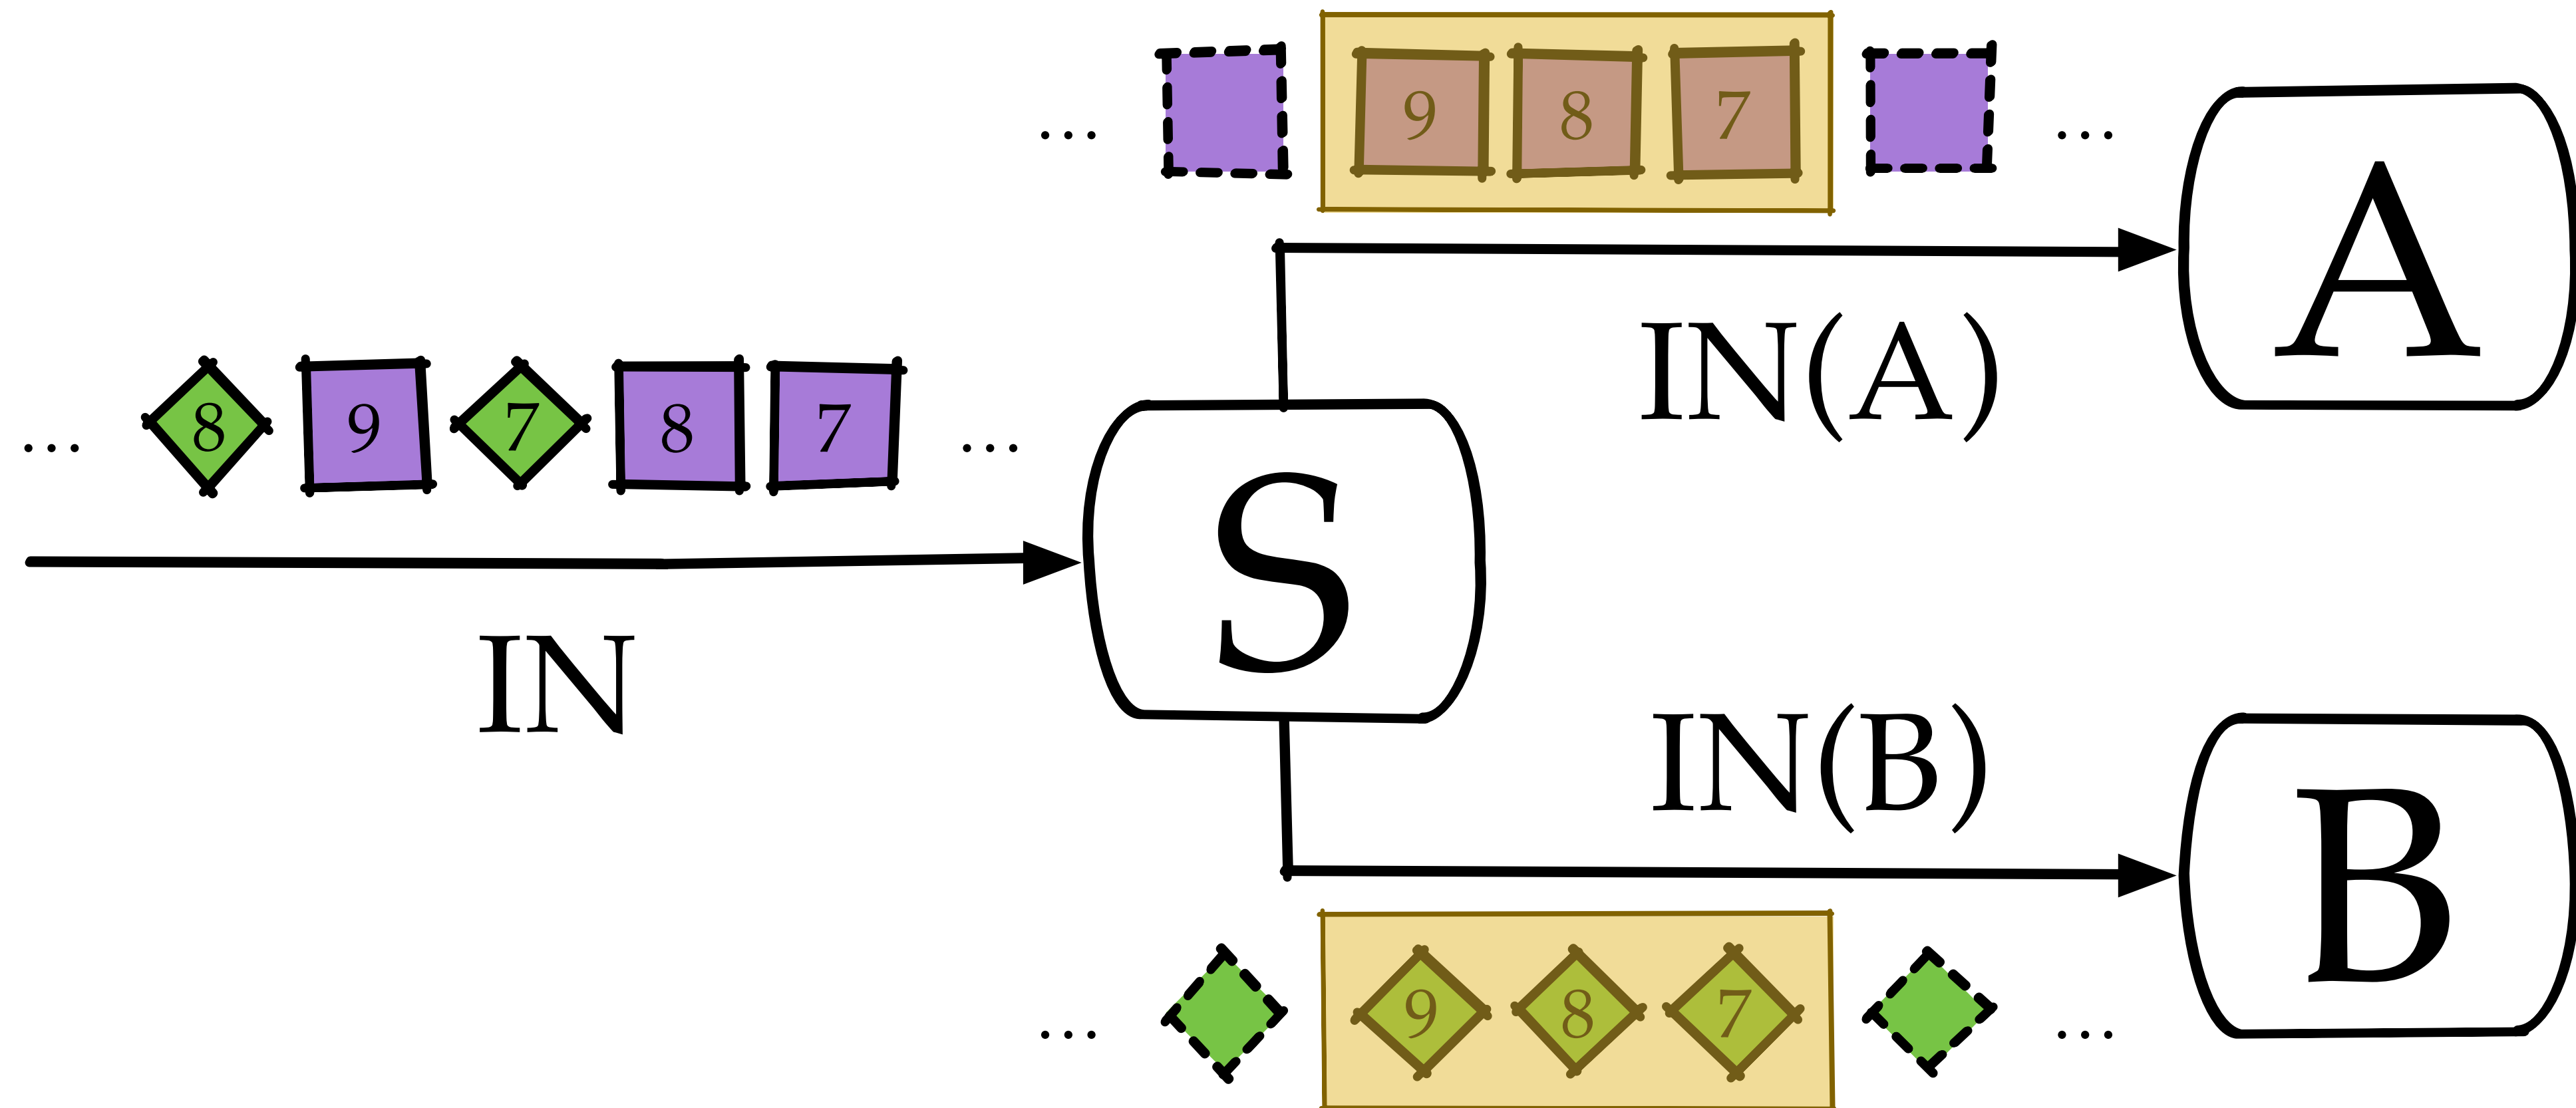
- Different windowing policies

- Sliding windows: window size + sliding factor
- Session windows [Apache Flink, Apache Beam...]

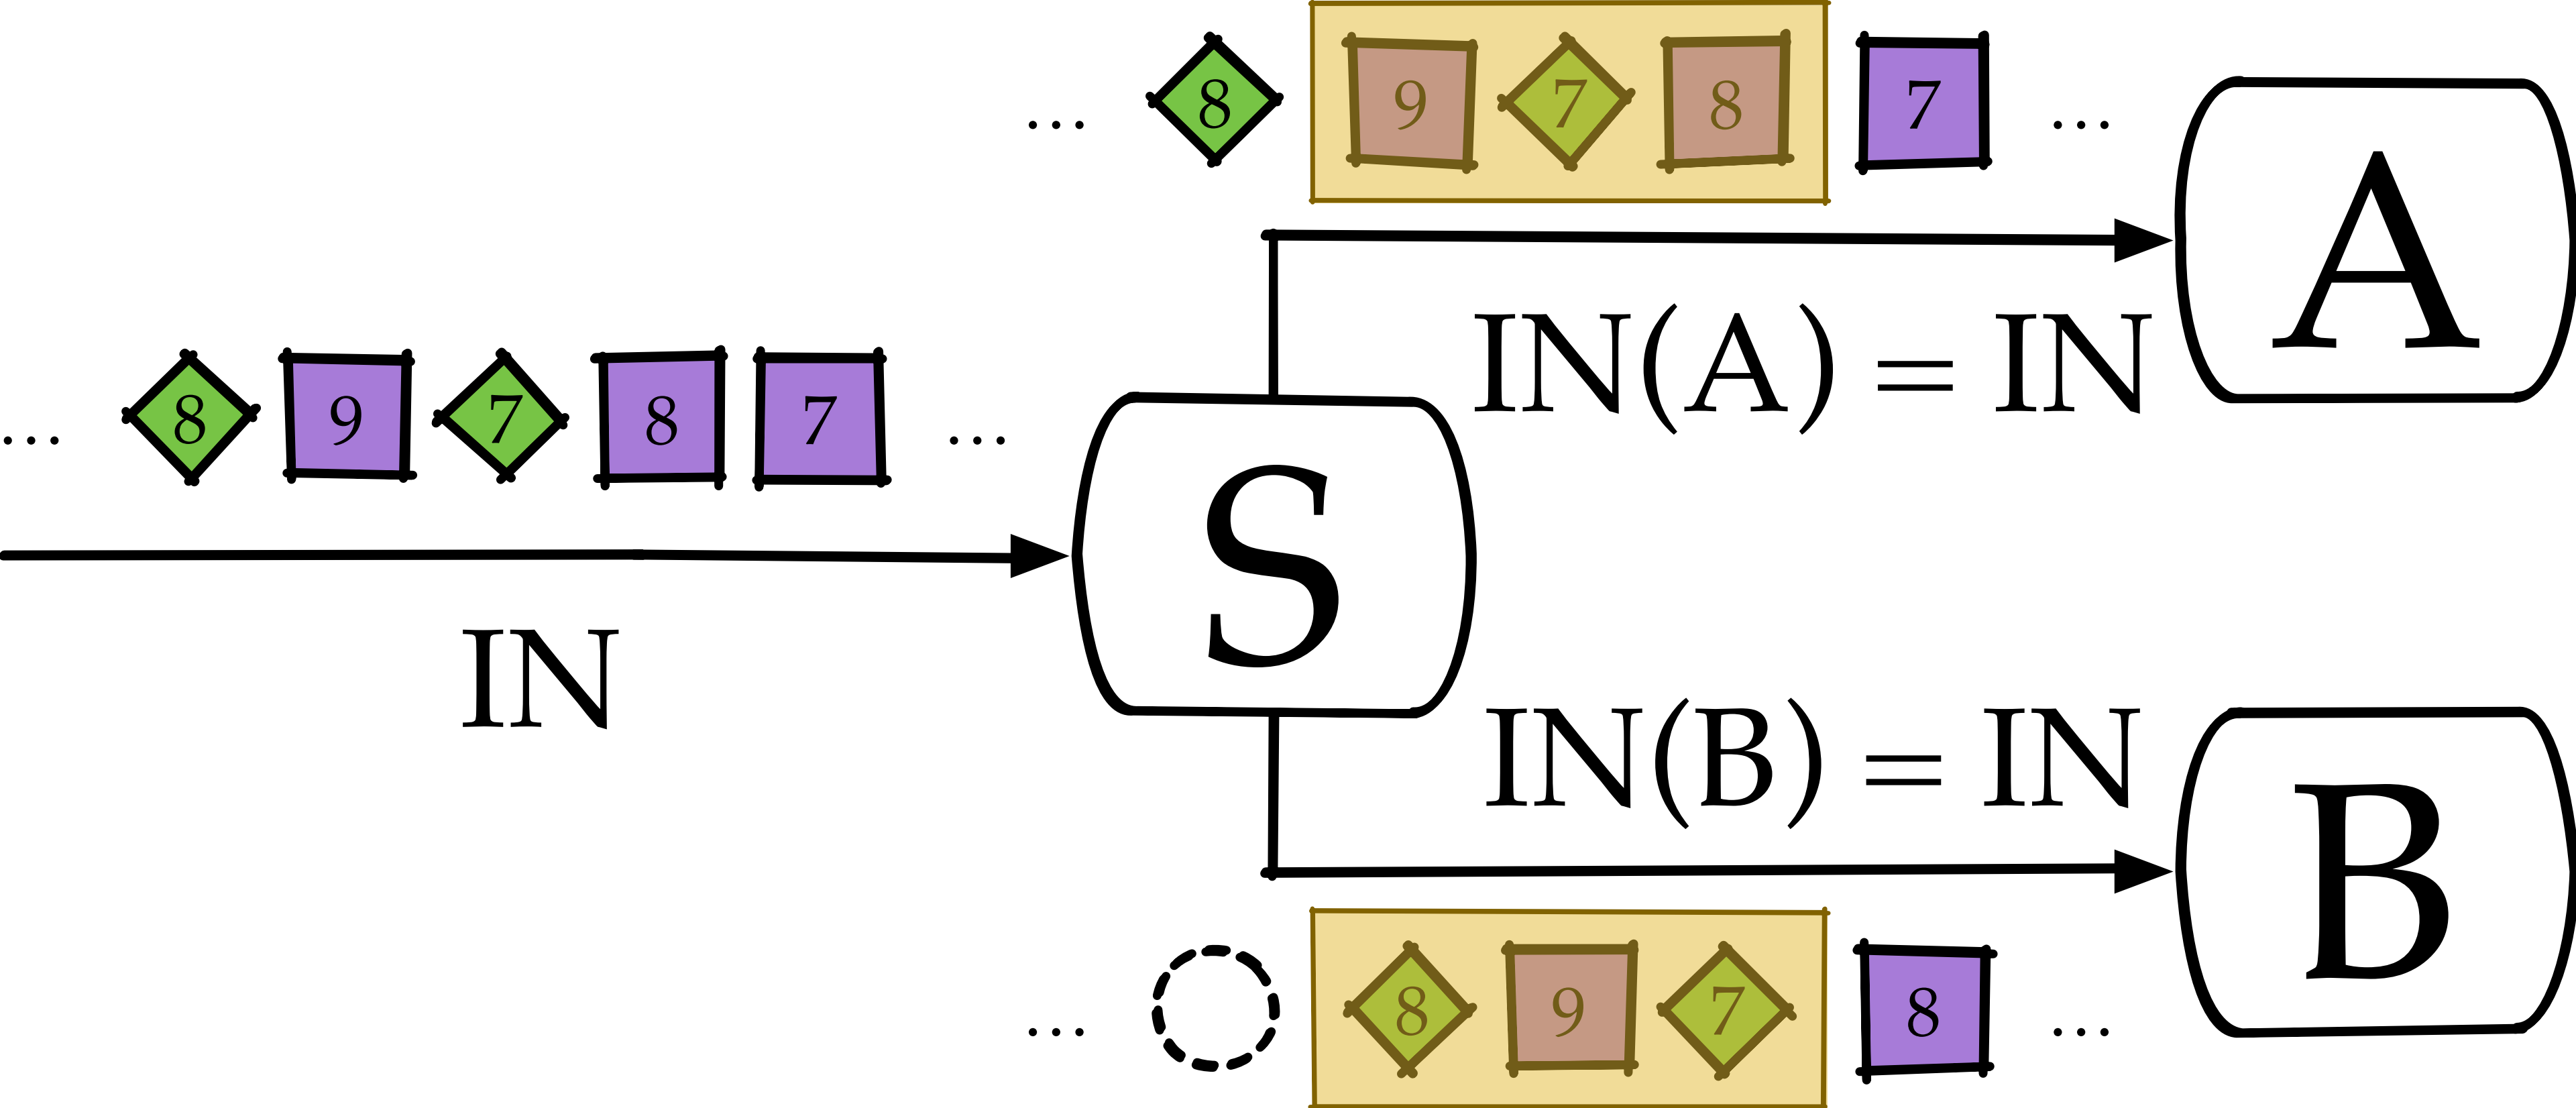
- Common implementation = Worker-side windowing (more parallelism)



KEY-PARTITIONING (KP)

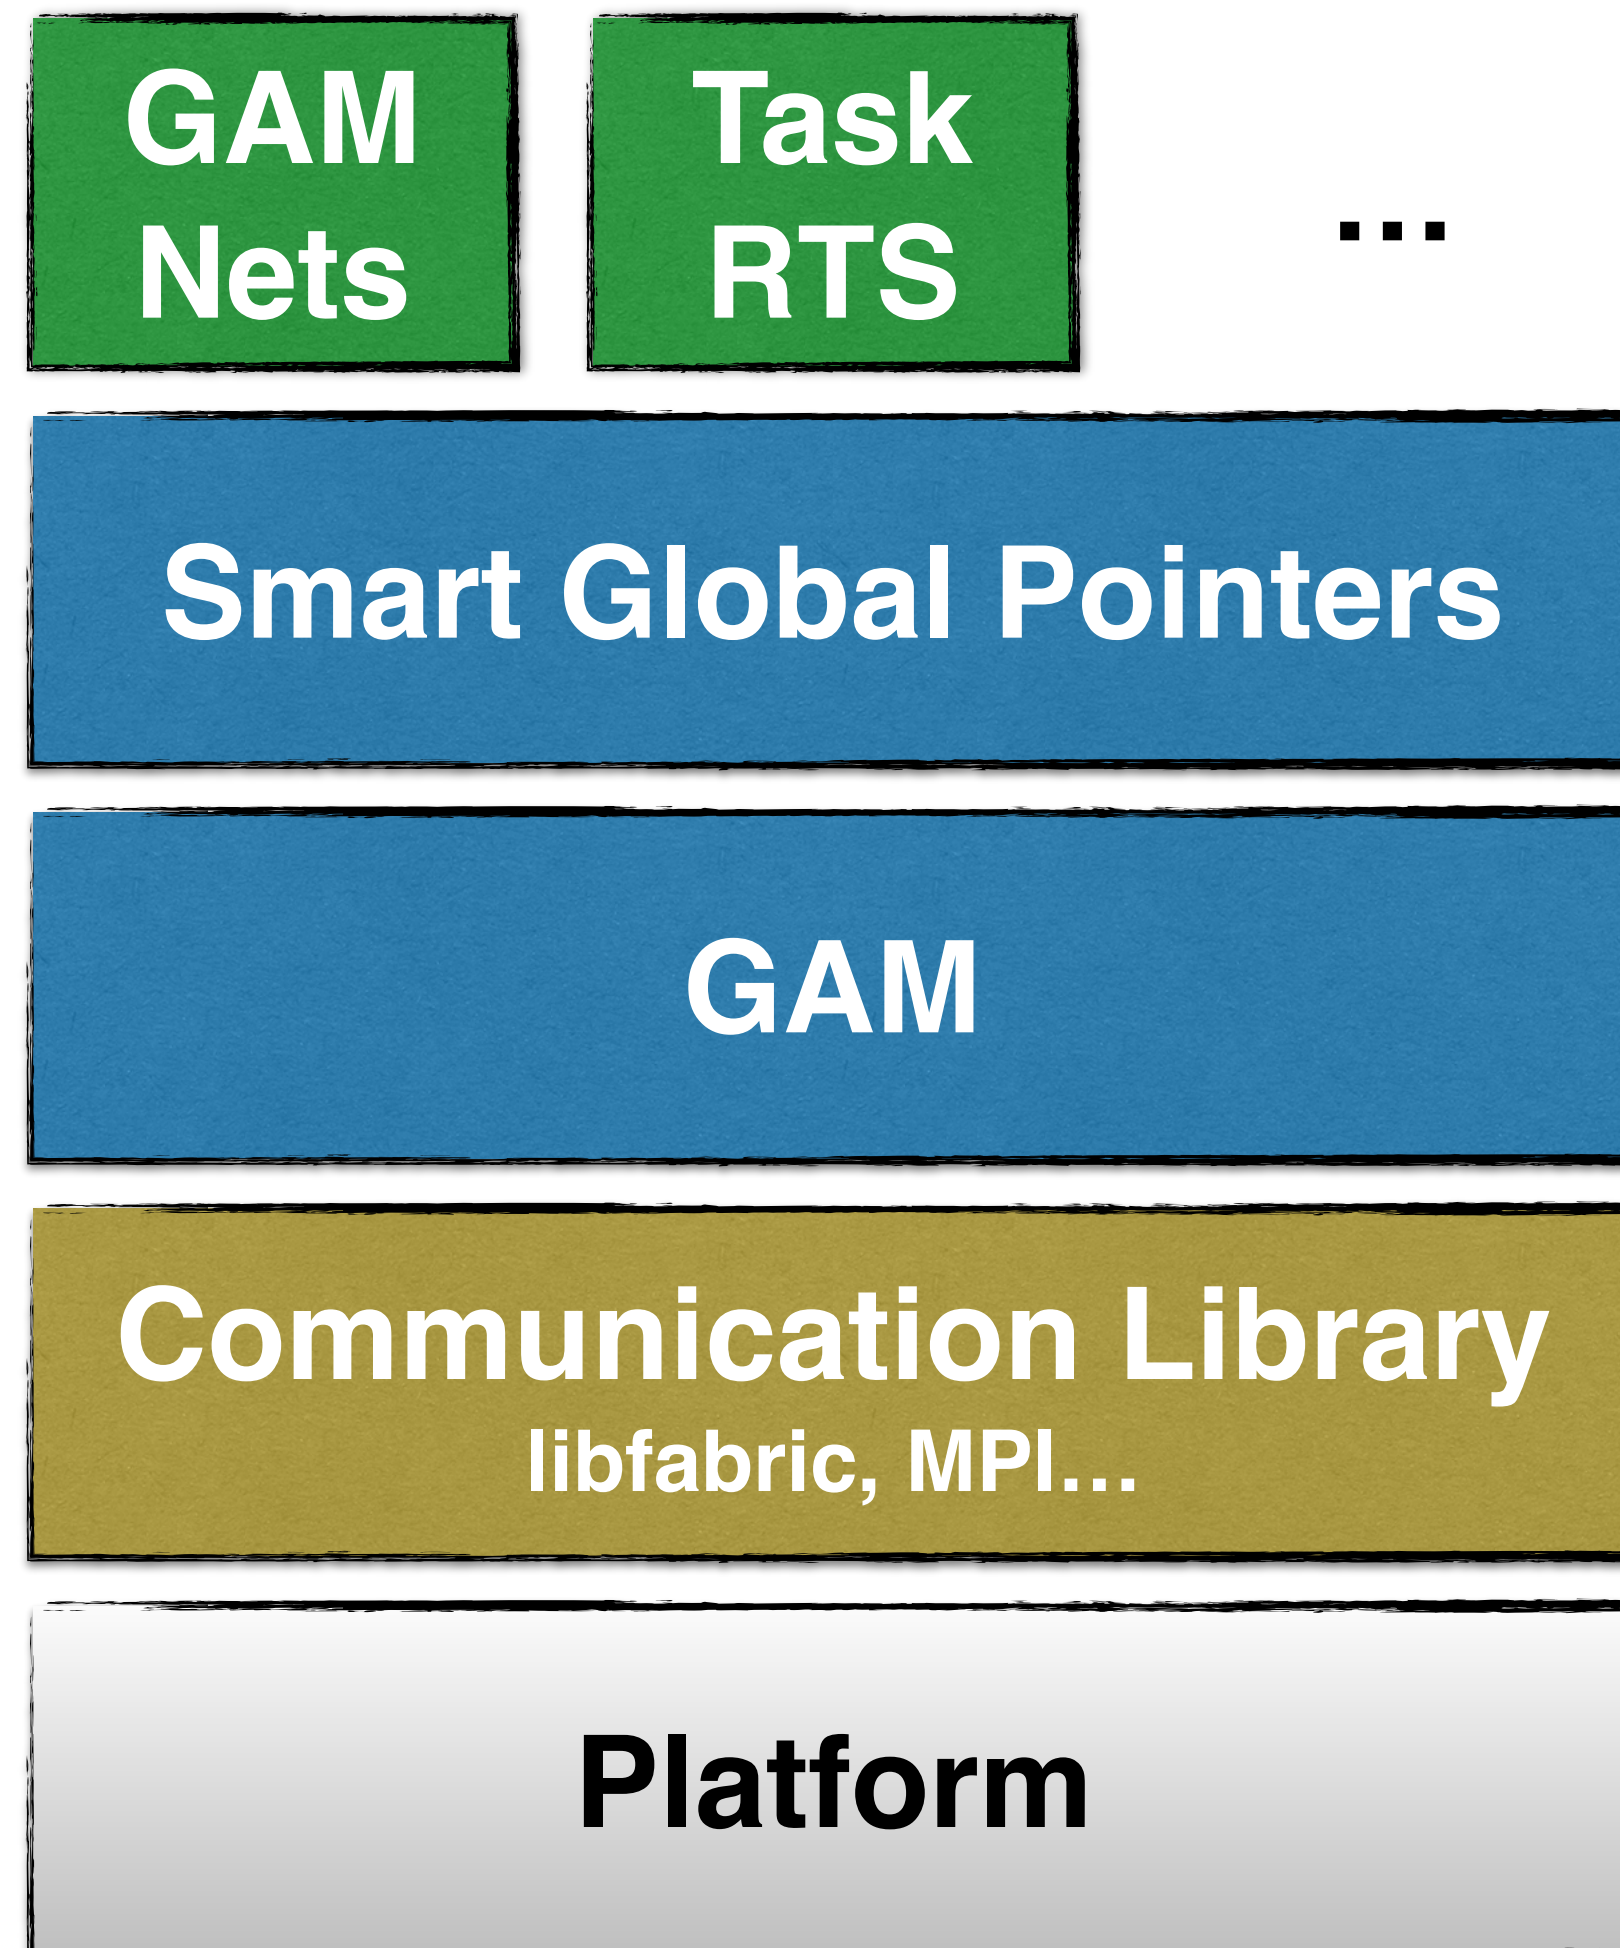


WINDOW-FARMING (WF)



A NEW PROPOSAL FOR A STREAM-ORIENTED PGAS

DISTRIBUTED FASTFLOW V.2 — C++ GLOBAL MEMORY STACK




M. Drocco. Parallel Programming with Global Asynchronous Memory: Models, C++ APIs and Implementations. PhD. Thesis proposal, University of Torino, 2017 (not yet defended).

Aldinucci, S. Campa, M. Danelutto, P. Kilpatrick, and M. Torquati, “Targeting Distributed Systems in FastFlow,” in *Euro-Par 2012 Workshops*,

GLOBAL ASYNCHRONOUS MEMORY

A NEW PROPOSAL FOR A STREAM-ORIENTED PGAS (BASED ON FF)

- From MPI Style:
communicate pointers
(a.k.a., capabilities)
 - From DSM Style:
shared address space
- ➔ Capability = both **data reference**
and **synchronization token**



PUBLIC CAPABILITIES

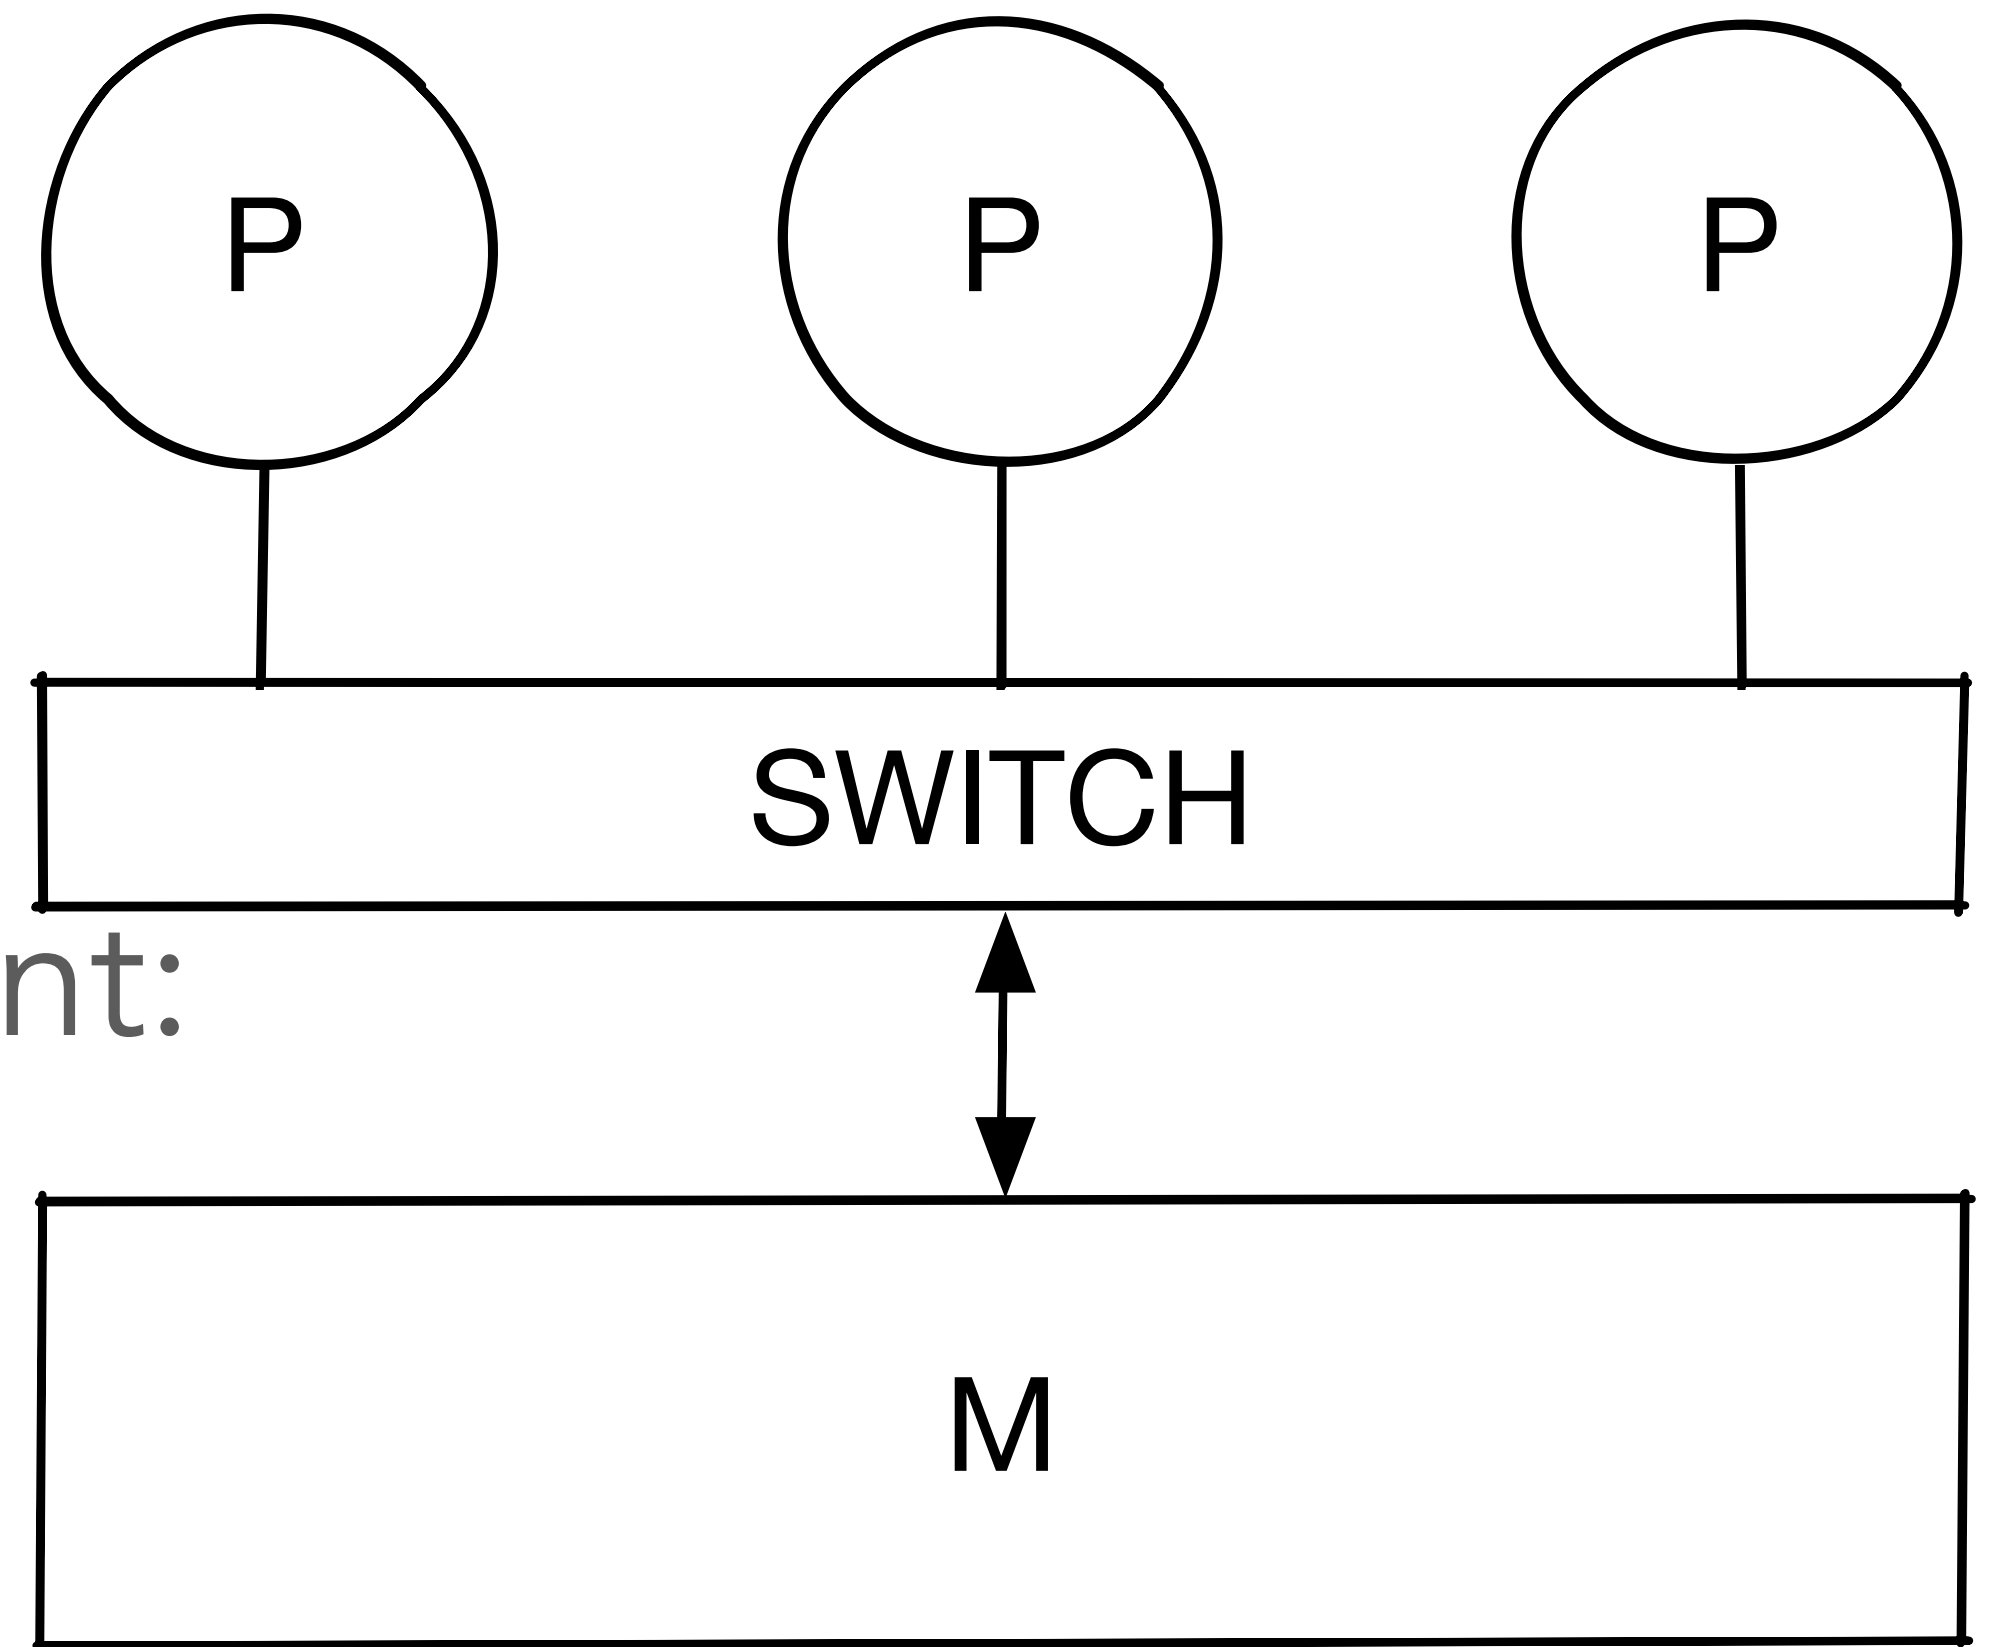
- Read-only (single assignment)
- Cacheable
- Can be copied

PRIVATE CAPABILITIES

- Exclusive read-write
- Not cacheable
- Can be moved

GAM MEMORY MODEL

- (Trivial) Sequential Consistency
- **Avoiding** consistency issues
(vs **solving** as in DSM/PGAS)
- SWMR cache-coherence invariant:
 - public \rightarrow (NW)MR
 - private \rightarrow SWSR



SMART GLOBAL POINTERS

- Rooted in modern C++
 - Intentional programming:
public → shared, private → unique
- Automatic Memory Management — the C++ way
 - Smartness = memory-reference lifetime binding
 - No memory leaks, no dangling pointers
 - No garbage collection (vs Java & friends)

PUBLIC POINTERS

```
public_ptr(T * const, Deleter);  
public_ptr<T> make_public(Args&&...);
```

```
//copy constructor/assignment...  
//move constructor/assignment...
```

```
public_ptr(private_ptr<T> &&);  
public_ptr& operator=(private_ptr<T> &&);
```

```
std::shared_ptr<T> local();
```



enables
plain C++ code

```
void push(executor_id to);  
public_ptr<T> pull_public(const exec_id from);  
public_ptr<T> pull_public();
```

PRIVATE POINTERS

```
private_ptr(T * const, Deleter);  
private_ptr<T> make_private(Args&&...);
```

```
//move constructor/assignment...
```

```
//NO copy constructor/assignment
```



```
gam_unique_ptr<T> local(); //unique_ptr + custom deleter
```

```
void push(executor_id to);  
private_ptr<T> pull_private(const exec_id from);  
private_ptr<T> pull_private();
```

enables
plain C++ code




SMARTNESS FOR PUBLIC POINTERS



- Distributed **reference counting** protocol
- Creation/copy/push trigger +1, destruction triggers -1
- C++ shared pointers: atomic-based reference counting


SMARTNESS FOR PRIVATE POINTERS



- Distributed **memory releasing** protocol
- Destruction triggers releasing
- C++ unique pointers: destruction-triggered release
- Inherently simpler than public pointers (as unique vs shared)


KP - INHERENTLY EFFICIENT

- Disjoint $IN(A)/IN(B)$
 - each tuple accessed "exclusively"
- GAM implementation
 - private pointers - exclusive capabilities
 - 1 RMA access per tuple



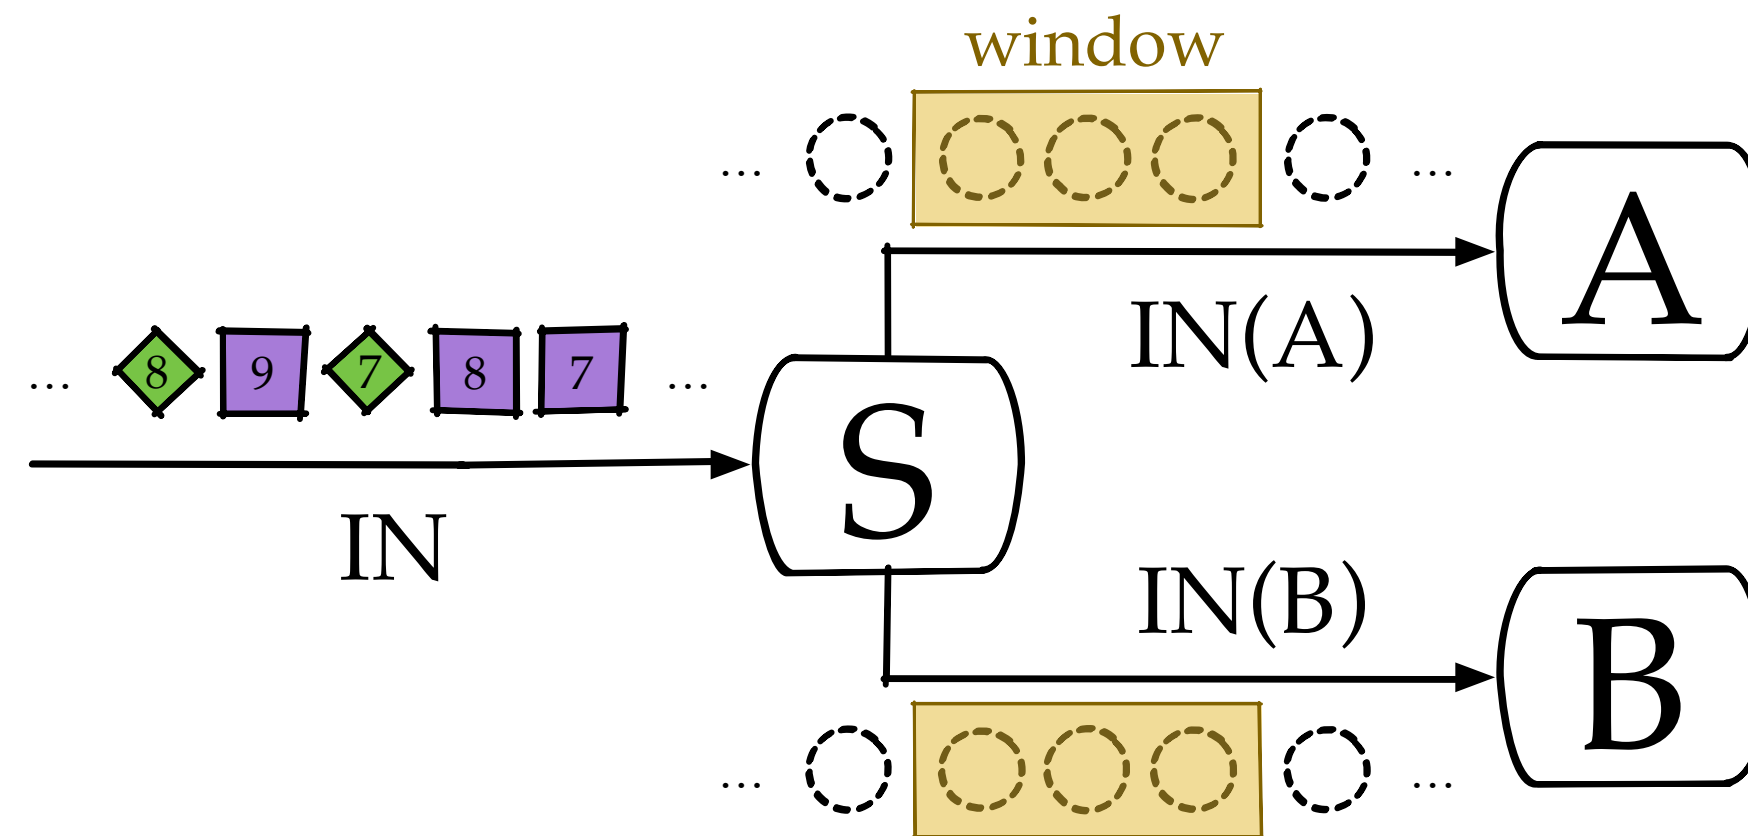
WF - INHERENTLY COMPLEX

- $IN(A) = IN(B) = IN$
 - each tuple accessed "by-any"
- GAM implementation
 - public pointers - read-only replicas
 - multiple RMA accesses per tuple




GAM ADVANTAGES


- Passing capabilities versus data
 - efficient worker-side windowing
 - extreme case: static dispatching not viable



HTP Video Restoration




Original Baboon
1024x1024



10% impulsive noise






50% impulsive noise



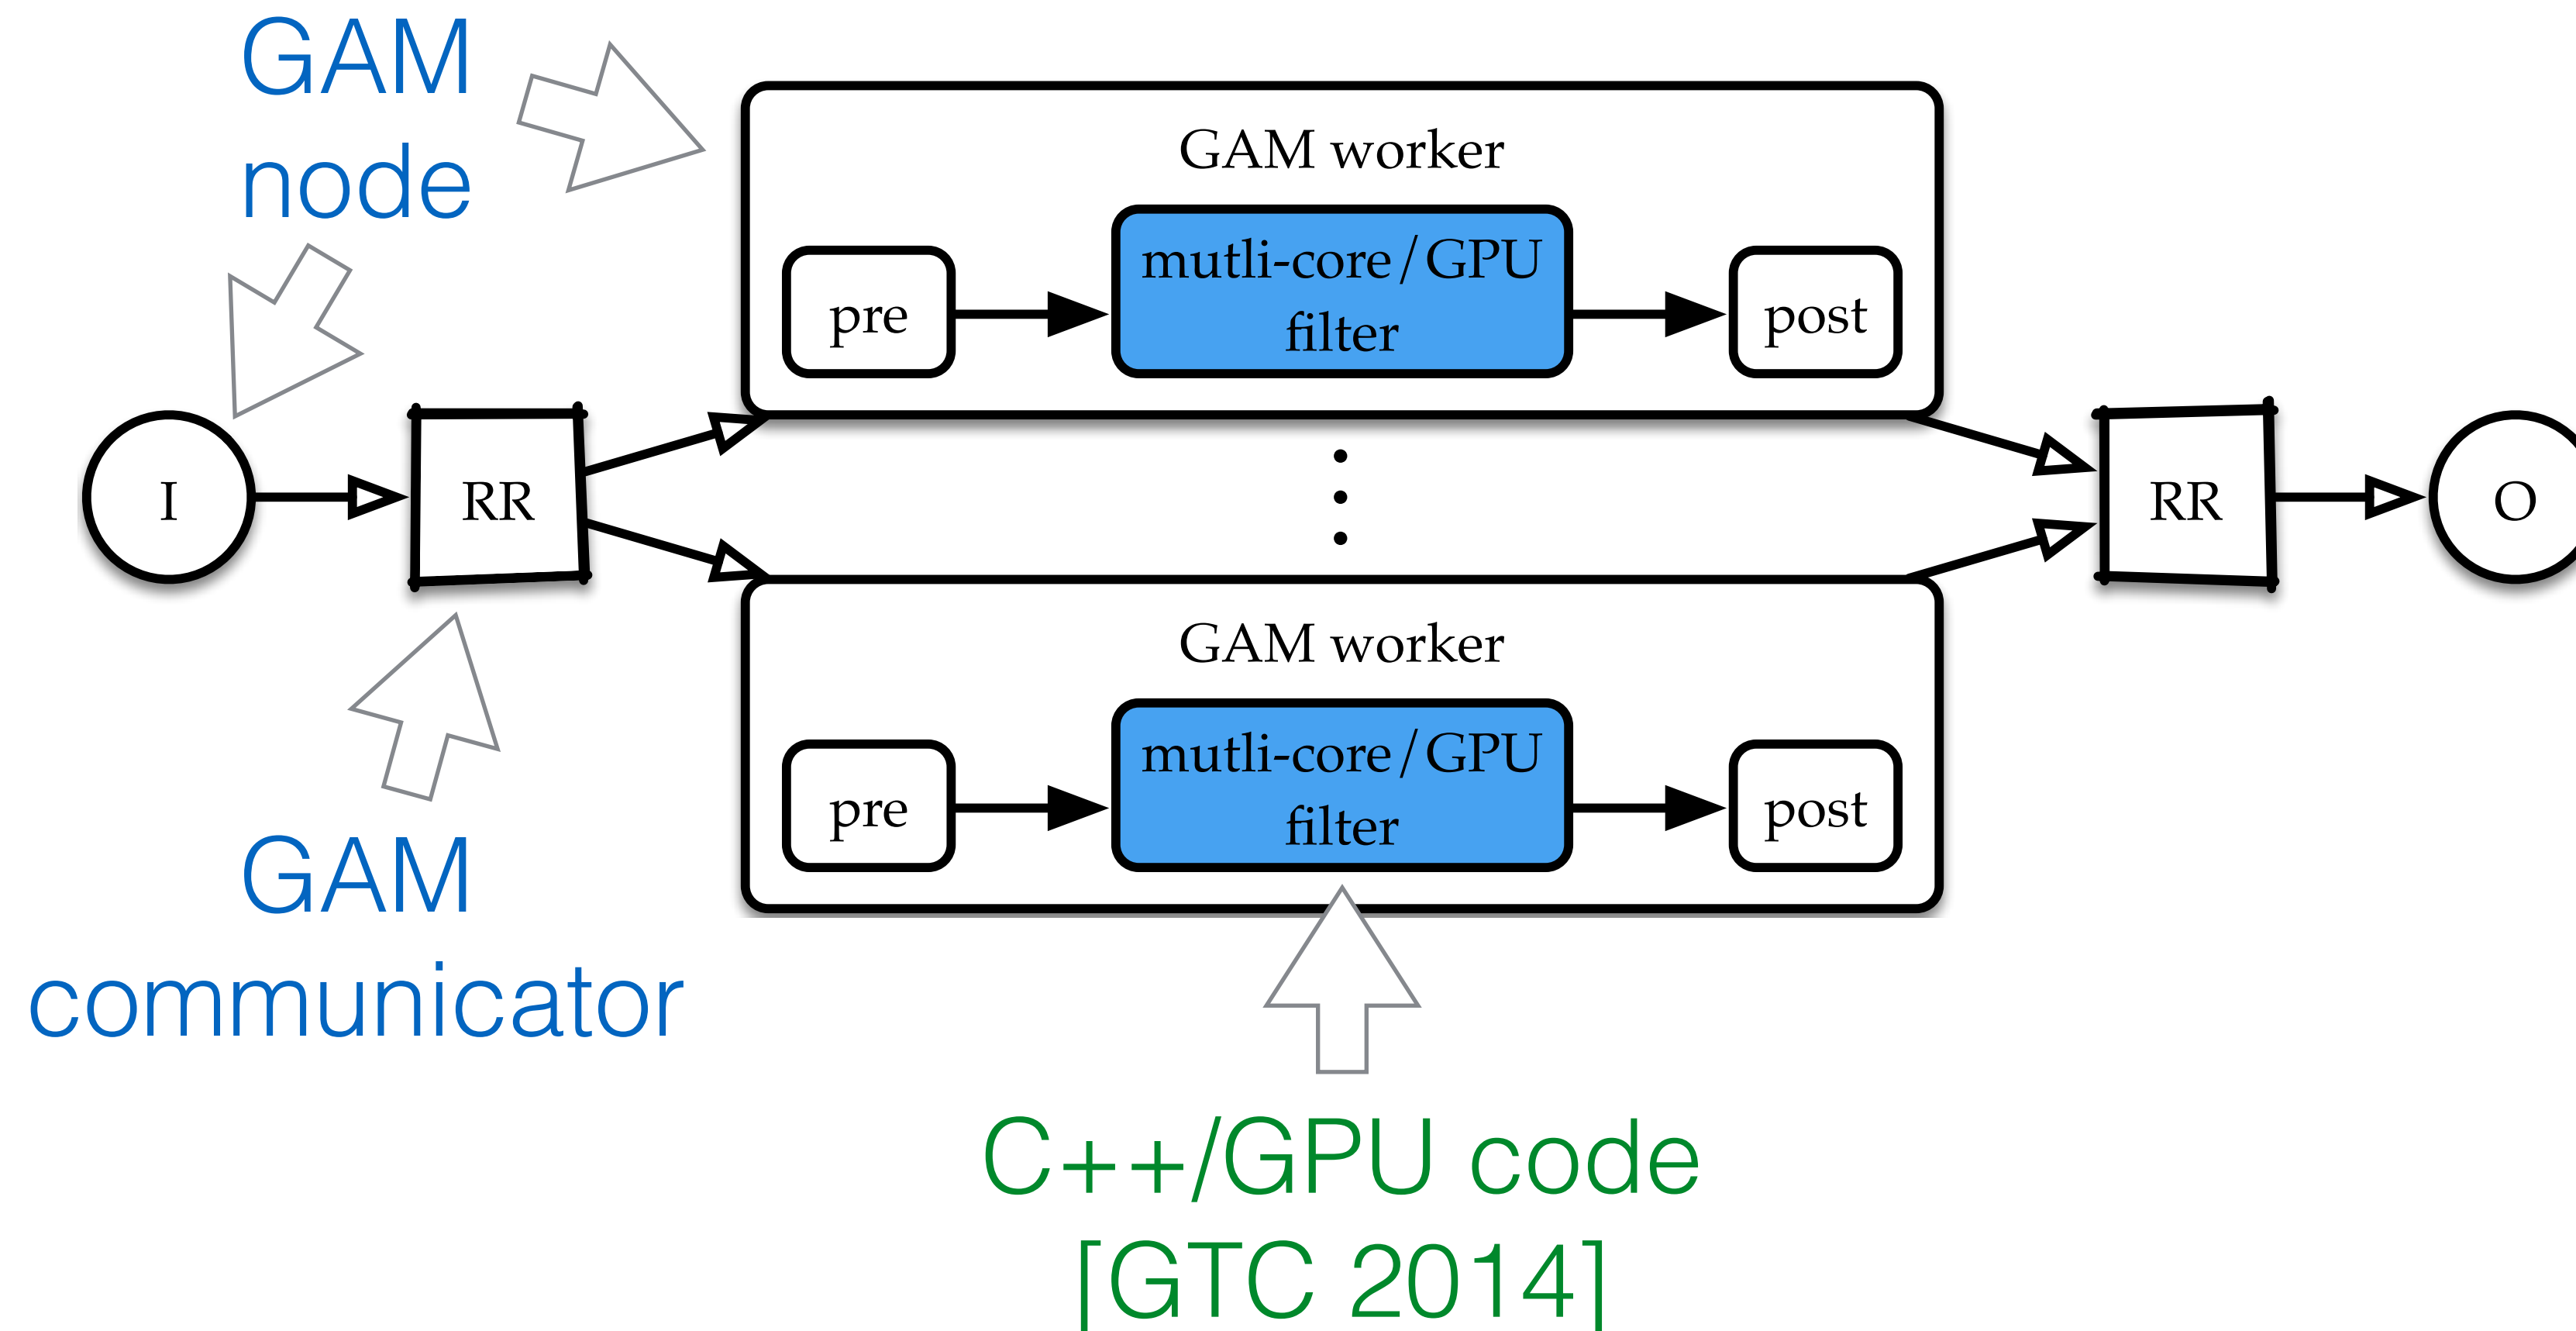
90% impulsive noise

Restored




HTP Video Restoration

- A GAM implementation of **ordering farm**



HTP Video Restoration





EuroPar 2018
Torino, Italy — 27-31 August 2018

Co-chairs: M. Aldinucci, L. Padovani, M. Torquati