



# Kent Academic Repository

**Marr, Stefan (2021) *Actors! And now? An Implementer's Perspective on High-level Concurrency Models, Debugging Tools, and the Future of Automatic Bug Mitigation*. In: 11th ACM SIGPLAN International Workshop on Programming Based on Actors, Agents, and Decentralized Control, AGERE!'21, 17 Oct 2021. (Unpublished)**

## Downloaded from

<https://kar.kent.ac.uk/94848/> The University of Kent's Academic Repository KAR

## The version of record is available from

<https://2021.splashcon.org/details/agere-2021-papers/8/Actors-And-now-An-Implementer-s-Perspective-o>

## This document version

Presentation

## DOI for this version

## Licence for this version

UNSPECIFIED

## Additional information

## Versions of research works

### Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

### Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in *Title of Journal*, Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

## Enquiries

If you have questions about this document contact [ResearchSupport@kent.ac.uk](mailto:ResearchSupport@kent.ac.uk). Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).

# Actors! And now?

An Implementer's Perspective on  
High-level Concurrency Models,  
Debugging Tools,  
and the Future of Automatic Bug Mitigation

THE  
ROYAL  
SOCIETY



Engineering and  
Physical Sciences  
Research Council

Stefan Marr  
17 October 2021

**Got a Question?  
Feel free to interrupt me!**



We're Looking for a Postdoc!

**Job Ad**

# Project CaMELot: Catch and Mitigate Event-Loop Concurrency Issues



**Please get  
in touch!**

**University of  
Kent**

<https://stefan-marr.de/2021/02/open-postdoc-position-on-language-implementation-and-concurrency/>

# Outcomes of Project MetaConc and work by



C. Torres Lopez



D. Aumayr



E. Gonzalez Boix



H. Mössenböck



# Actors! What are Actors?

## **43 Years of Actors: A Taxonomy of Actor Models and Their Key Properties**

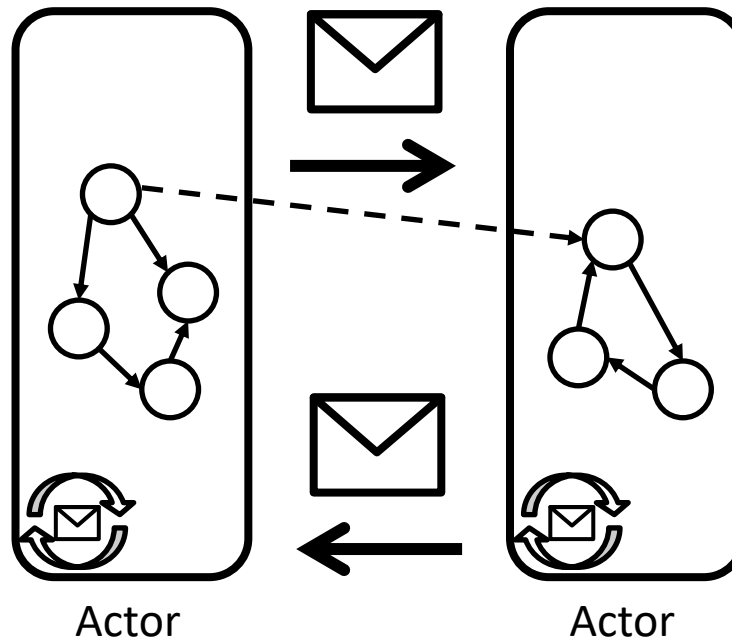
Joeri De Koster  
Vrije Universiteit Brussel

Tom Van Cutsem  
Nokia Bell Labs

Wolfgang De Meuter  
Vrije Universiteit Brussel

- Many different variants
- For the 50 Years' Edition:
  - Which model is good for what?
    - Suitable problems/applications
    - Unsuitable problems per model
  - ...

# Communicating Event Loops



# Concurrency Bugs are Common in Event Loop Systems



53 projects, 57 issues

2 studies

12 projects, 1000 potential issues



12 projects

1 study

53 concurrency issues



Websites in top 500

≈1-10 concurrency issues per website



**Tip of the Iceberg**



8-27 apps

≈2-20 concurrency issues per app



6 projects

1 study

35 known event races



**How to get rid  
of all these bugs?**

Perhaps not a way to get rid of them all, but at least to make it easier

# **DEBUGGING ACTORS WITH SUITABLE BREAKPOINTS/STEPPING**

# Actor Breakpoints/Stepping

```
prom := aResult <-: get.
```

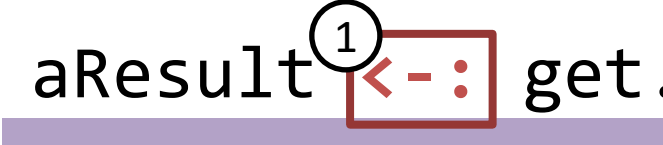
```
prom whenResolved: [:r |  
  r println  
].
```

Actor A

# Actor Breakpoints/Stepping

msg send  
msg receive  
promise resolver  
promise resolution

```
prom := aResult1 <- : get.
```



```
prom whenResolved: [:r |  
  r println  
].
```

Actor A

# Actor Breakpoints/Stepping

msg send  
msg receive  
promise resolver  
promise resolution

```
prom := aResult <-: get.
```

```
prom whenResolved: [:r |  
  r println  
].
```

Actor A

```
class Result = ()(  
  
  public get = (  
    ② | result |  
      result := 42.  
      ^ result  
  )  
  
)
```

Actor B

# Actor Breakpoints/Stepping

msg send  
msg receive  
promise resolver  
promise resolution

```
prom := aResult <-: get.
```

```
prom whenResolved: [:r |  
  r println  
].
```

Actor A

```
class Result = ()(  
  
  public get = (  
    | result |  
    result := 42.  
    ③ ^ result  
  )  
  
)
```

Actor B

# Actor Breakpoints/Stepping

msg send  
msg receive  
promise resolver  
promise resolution

```
prom := aResult <-: get.
```

```
prom whenResolved: [:r |  
  r println  
].
```

Actor A

```
class Result = ()(  
  public get = (  
    | result |  
    result := 42.  
    ^ result  
  )  
)
```

Actor B

# Actor Breakpoints/Stepping

```
prom := aResult <-: get.  
prom whenResolved: [:r |  
  r println  
].
```

Actor A

```
class Result = ()(  
  before async  
  after async  
  public get = (  
    ① | result |  
      result := 42.  
    ② ^ result  
  )  
)
```

Actor B



# Actor Breakpoints/Stepping

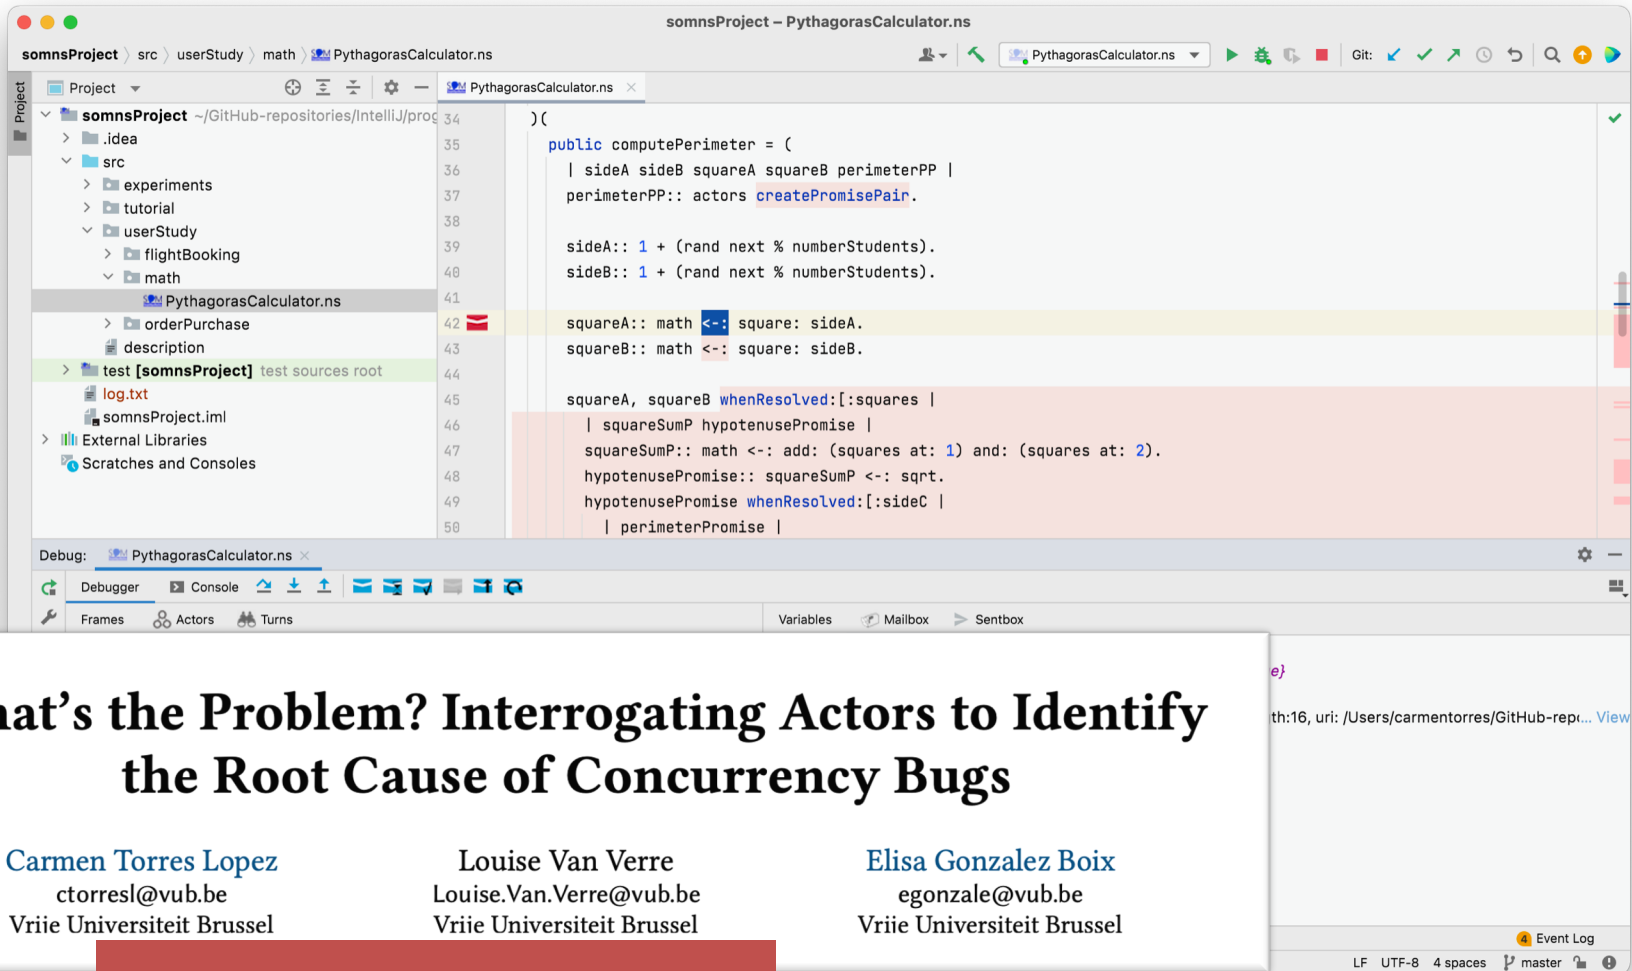
```
prom := aResult <-: get.  
  promise resolver  
  promise resolution  
prom whenResolved: [:r |  
① r println  
].
```

Actor A

```
class Result = ()(  
  public get = (  
    | result |  
    result := 42.  
    ^ result  
  )  
)
```

Actor B

# Apgar: A Debugger Made for Actor Programs



**What's the Problem? Interrogating Actors to Identify the Root Cause of Concurrency Bugs**

**Carmen Torres Lopez**  
ctorresl@vub.be  
Vrije Universiteit Brussel

**Louise Van Verre**  
Louise.Van.Verre@vub.be  
Vrije Universiteit Brussel

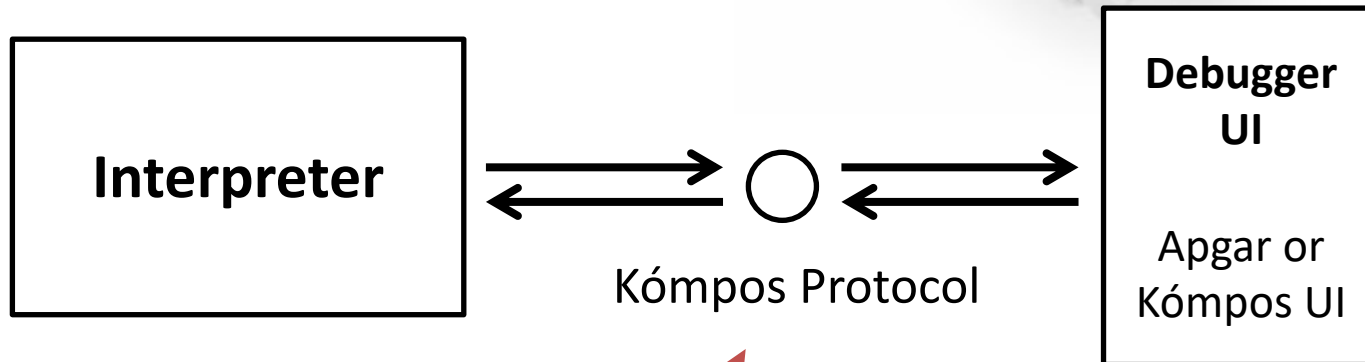
**Elisa Gonzalez Boix**  
egonzale@vub.be  
Vrije Universiteit Brussel

Event Log  
LF UTF-8 4 spaces master

Carmen's presentation is in  
about 5.5h here at AGERE

# Kómpos Architecture

Κόμπος  
kompos



**The "Magic" Bit**

# The Kómpos Debugger

The screenshot displays the Kómpos Debugger interface. At the top, the browser window title is "Kompos Debugger" and the address bar shows "localhost:8888/index.html". The program being debugged is "core-lib/KomposDemo.com".

The main area shows a diagram of the program's structure. It includes several actors: "InputGeneratorActor", "DataActor", "Platform", "JsonInputActor", and a highlighted actor "λcalculateSumOfwithfrominto@386@12@387@52 (9)". Arrows indicate dependencies and data flow between these actors.

On the right side, there is a control panel with a "Reconnect" button and several status icons. Below this, a log window displays the following messages:

```
2017-03-28T21:42:17.243: [WS] close  
2017-03-28T21:42:00.669: Send breakpoints: 0  
2017-03-28T21:42:00.668: [WS] open
```

At the bottom, there are panels for "Platform" and "ReportActor", each with its own control icons. The "ReportActor" panel shows the same actor ID as the highlighted one in the diagram.

A red box in the foreground contains the following text:

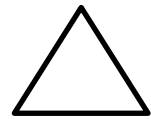
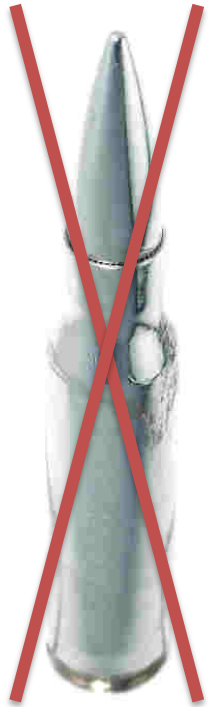
**Demo:**  
<https://stefan-marr.de/2017/10/multi-paradigm-concurrent-debugging/>

**Even with better debuggers,  
we'll still have concurrency bugs  
in our actor systems...**

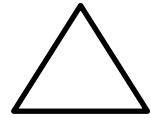
**Maybe, just maybe!**

**Maybe Actors aren't the best  
choice for every problem?**

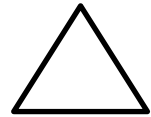
# Maybe there are no Silver Bullets?



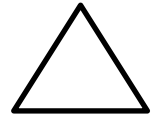
**Actors**



**CSP**



**Locks, Monitors, ...**



**Fork/Join**



**Transactional Memory**

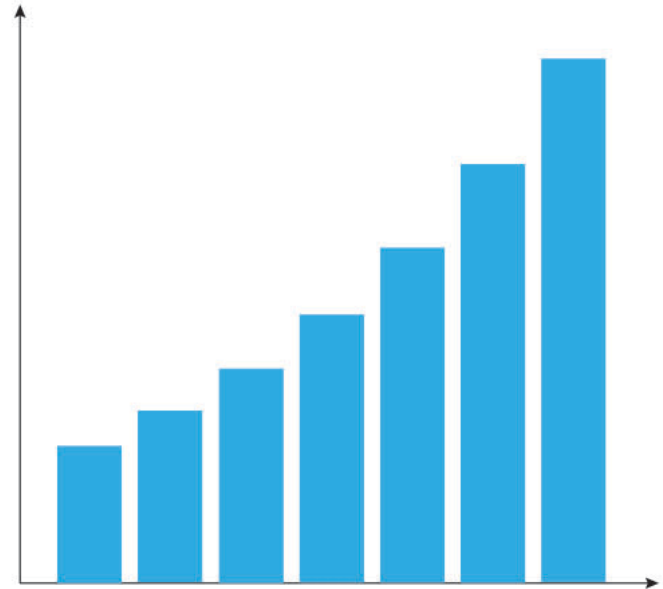


**Data Flow**

...

# Building an Online Sales-Data Processor

```
{"item": "beer",  
  "price": 5.5,  
  "quantity": 344,  
  "customer": "<Prog>",  
  "address": "Pleinlaan 2"}
```



## Stream of Sales Events

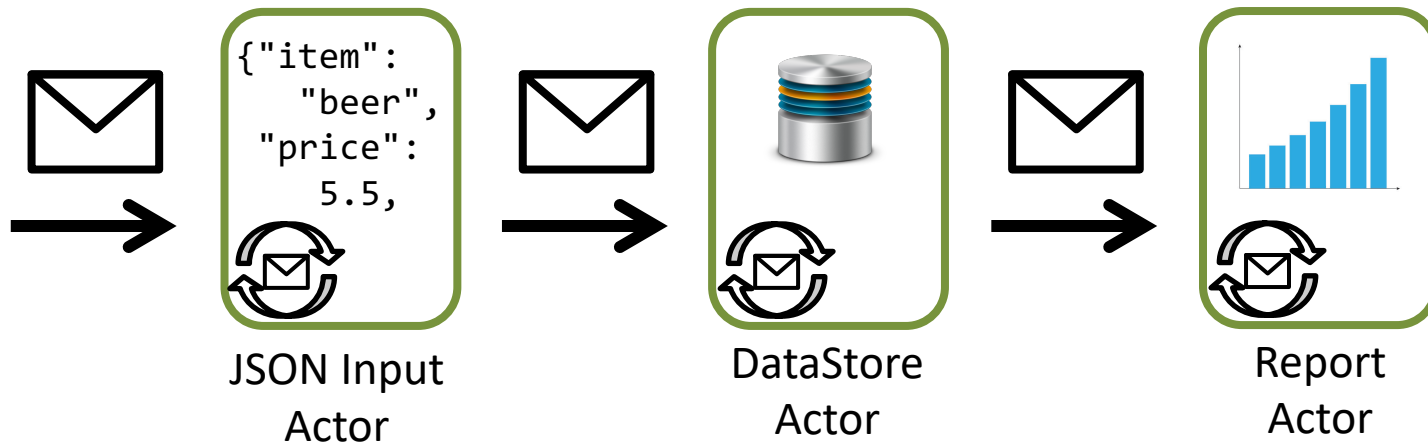
- Track revenue
- Report sales revenue over time



# Subsystems as Asynchronous Activities

## Use Actors as Main Abstraction

Event-Loop Model fits UI and System Paradigms

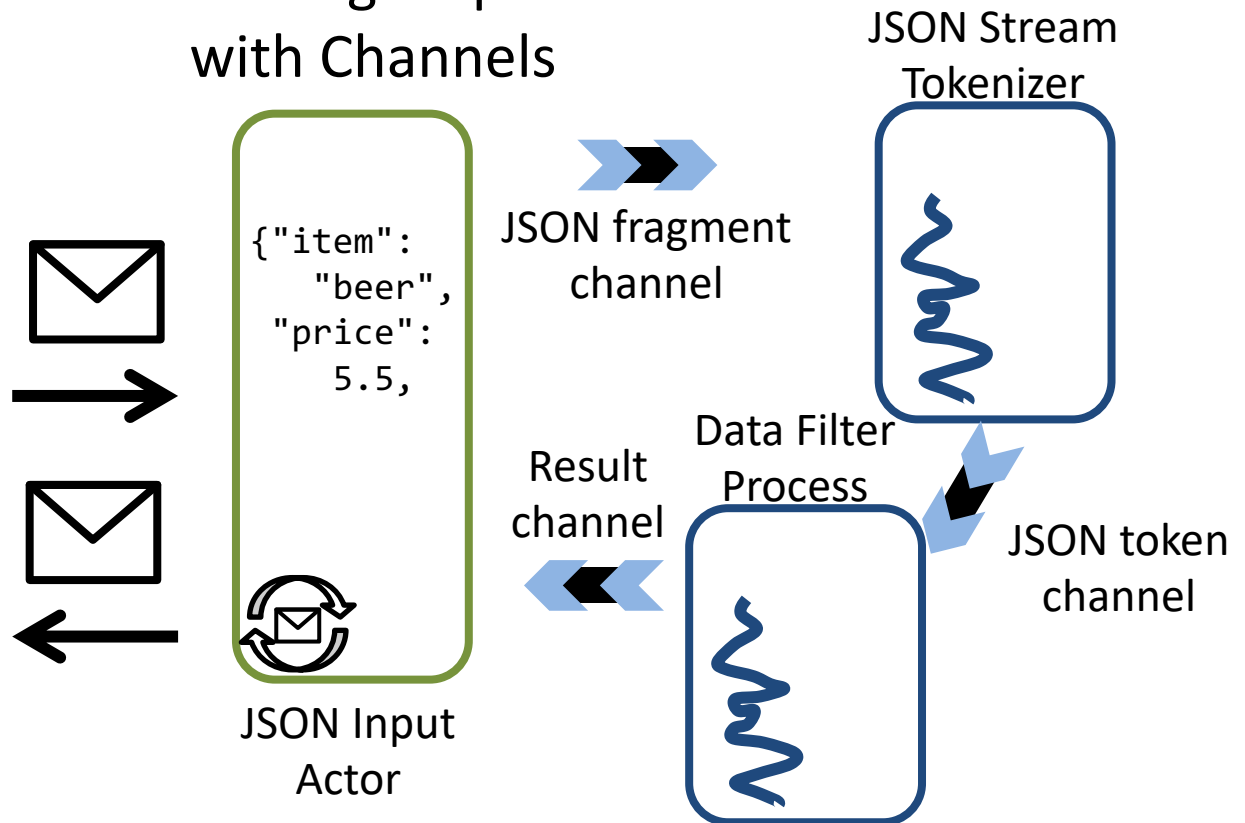


# Parallelize JSON Processing

Using Communicating Sequential Processes

with Channels

- Strict consumer/producer relationship
- Allow for pipeline parallelism



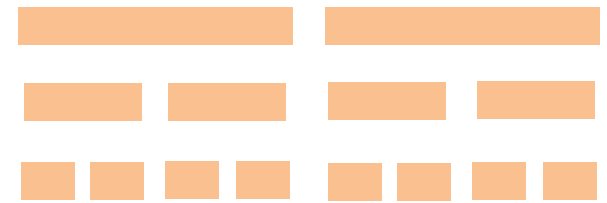
# Sales Revenue Over Time based on Large Data Array



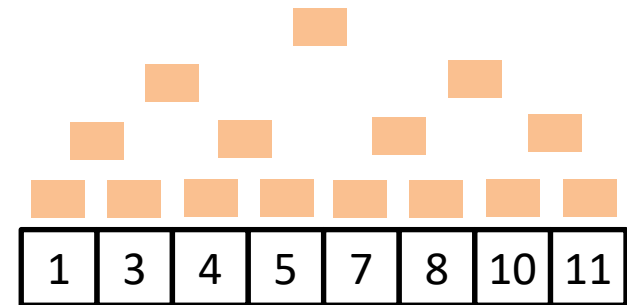
Report  
Actor

1	2	1	1	2	1	2	1
---	---	---	---	---	---	---	---

Construct Sum Tree  
*in parallel*



Calculate Prefix Sum  
*in parallel*



1	3	4	5	7	8	10	11
---	---	---	---	---	---	----	----

Parallel Prefix Sum Calculation  
with fork/join parallelism

# How to build debuggers to support all the Concurrency Models?

κόμπος  
kompos



# Κόμπος: A PLATFORM FOR DEBUGGING COMPLEX CONCURRENT APPLICATIONS

# The Kómpos Debugger

## **A Concurrency-Agnostic Protocol for Multi-Paradigm Concurrent Debugging Tools**

Stefan Marr  
Johannes Kepler University  
Linz, Austria  
stefan.marr@jku.at

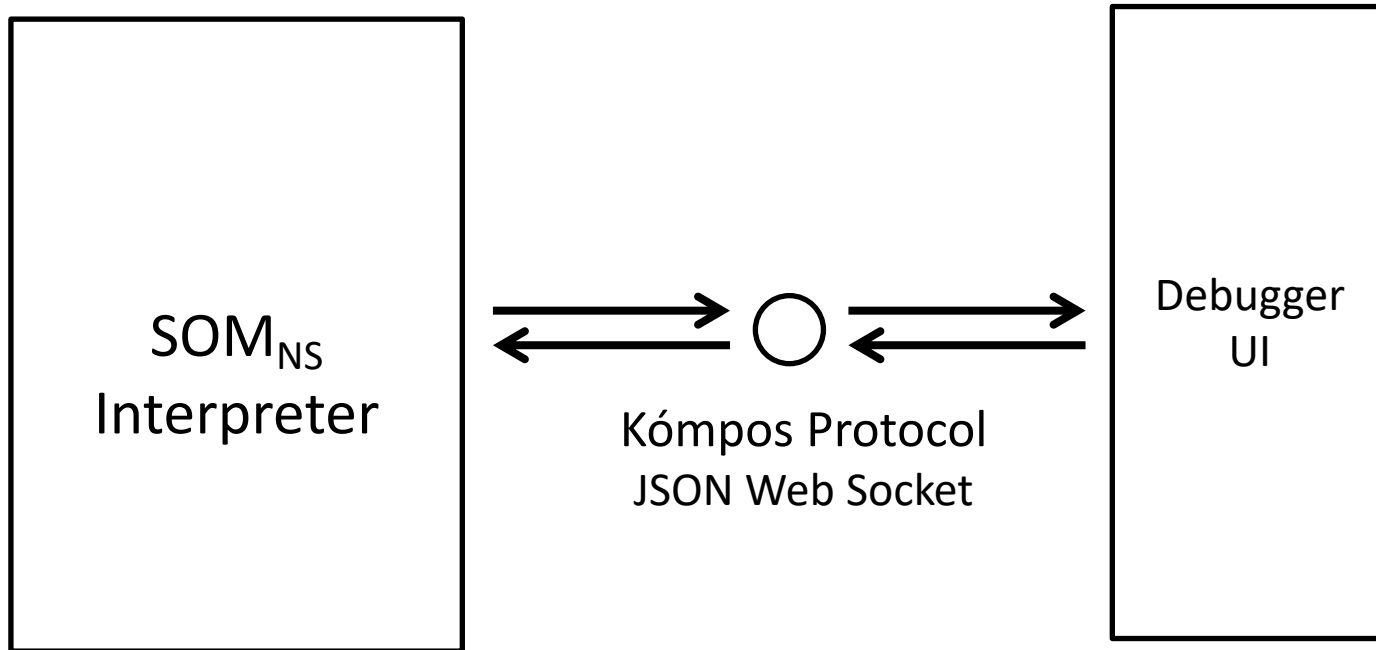
Carmen Torres Lopez  
Vrije Universiteit Brussel  
Brussels, Belgium  
ctorresl@vub.be

Dominik Aumayr  
Johannes Kepler University  
Linz, Austria  
dominik.aumayr@jku.at

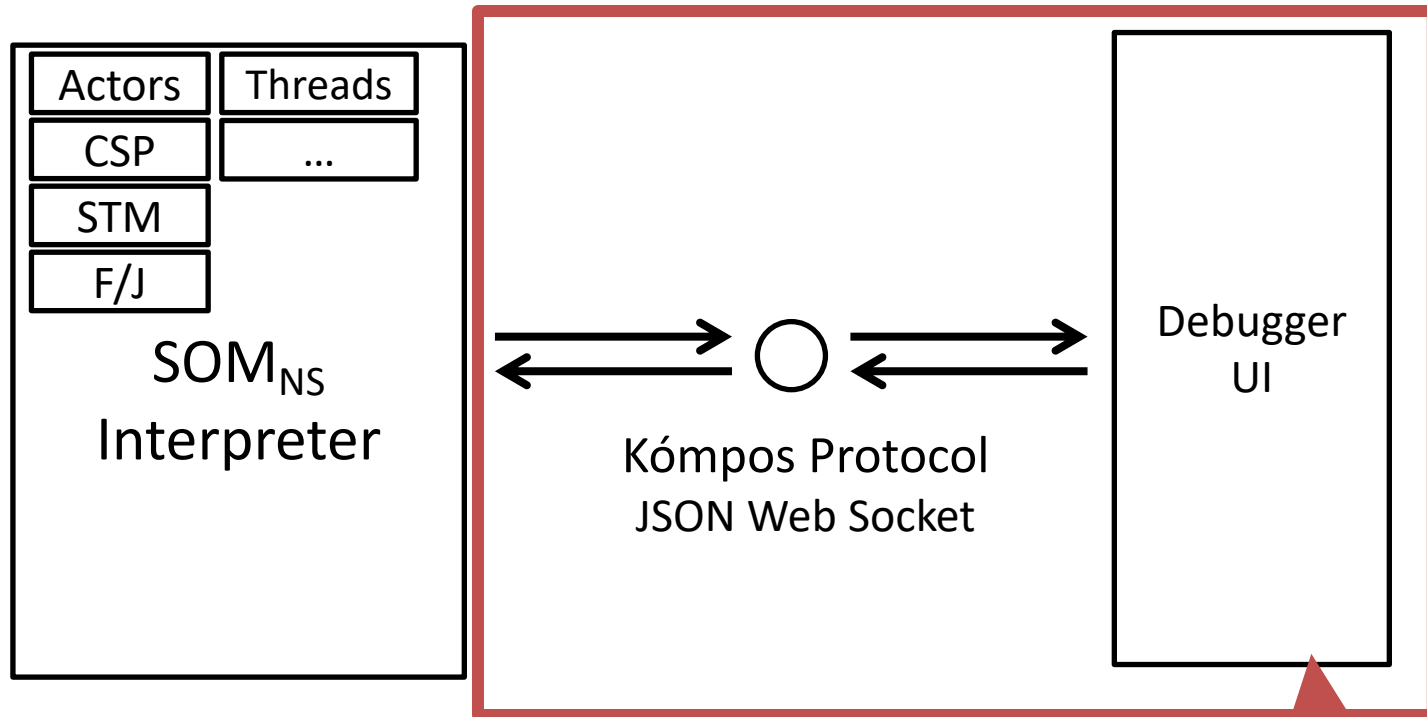
Elisa Gonzalez Boix  
Vrije Universiteit Brussel  
Brussels, Belgium  
egonzale@vub.be

Hanspeter Mössenböck  
Johannes Kepler University  
Linz, Austria  
hanspeter.moessenboeck@jku.at

# Kómpos Architecture



# Kómpos Architecture



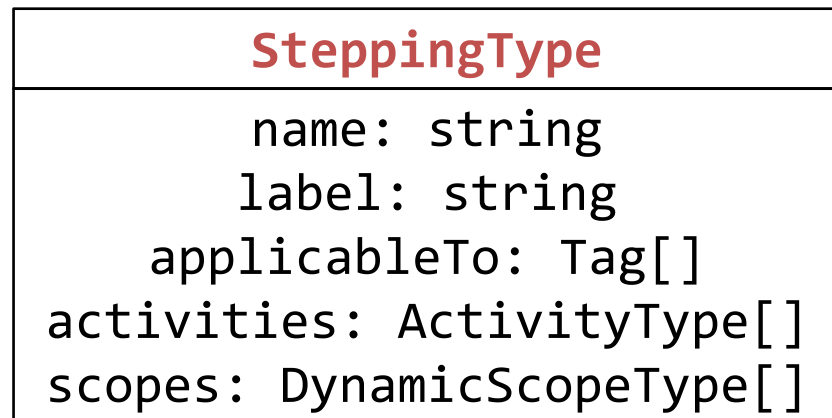
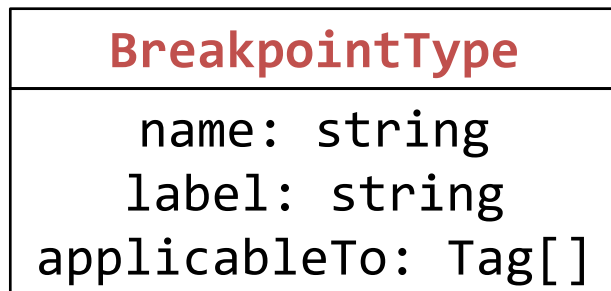
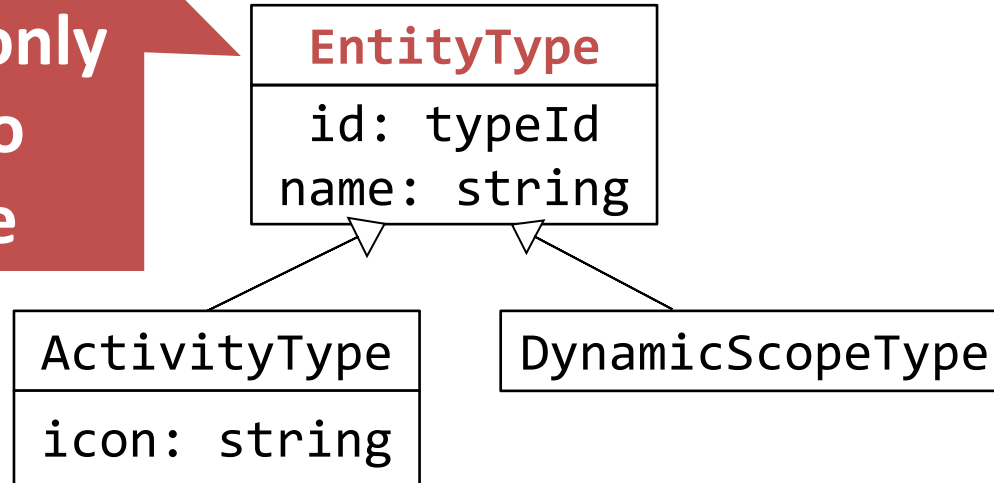
**Agnostic of  
Concurrency  
Models**

**And we have  
two UIs! Apgar  
& Kómpos UI**



# Kómpos Protocol Metadata

Concurrency semantics only known to language



# Kómpos Protocol Messages

SetBreakpoint
location: Coord <b>type: BreakpointType</b>

DoStep
activityId: id <b>type: SteppingType</b>

**Debugger UI just  
“lists” available  
types**

Stopped
activityId: id location: Coord <b>actType: ActivityType</b> scopes: DynamicScopeType[]

# A Model-Agnostic Debugger: Example Channel Breakpoints

channel out

① write: 42 ④

Process A

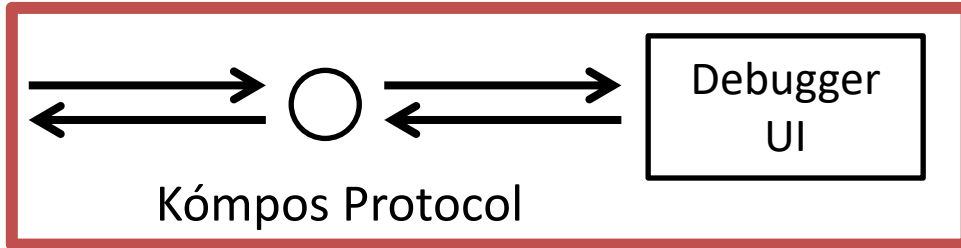
channel in

③ read ②

Process B

**“just” source locations and ids!  
UI doesn't need to know these  
concepts!**

# Debuggers can be Great for High-level Concurrency Models!



```
promise resolver  
promise resolution  
prom whenResolved: [:r |  
r println ].
```

**Make tools agnostic**

**Offer the Key Features  
as Breakpoints/Steps**



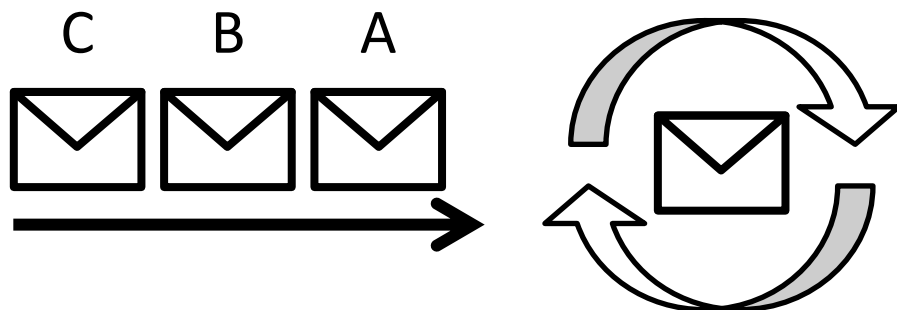


Reproduces only 1 in 10? How can I fix such a bug???

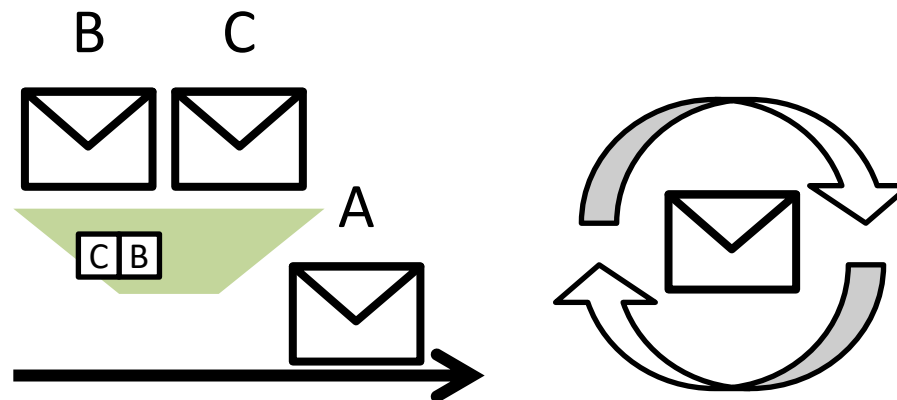
# **NON-DETERMINISM MAKES FOR UNHAPPY DEBUGGERS**

# One Solution: Record & Replay

- Record event order



- Replay reorder to fit

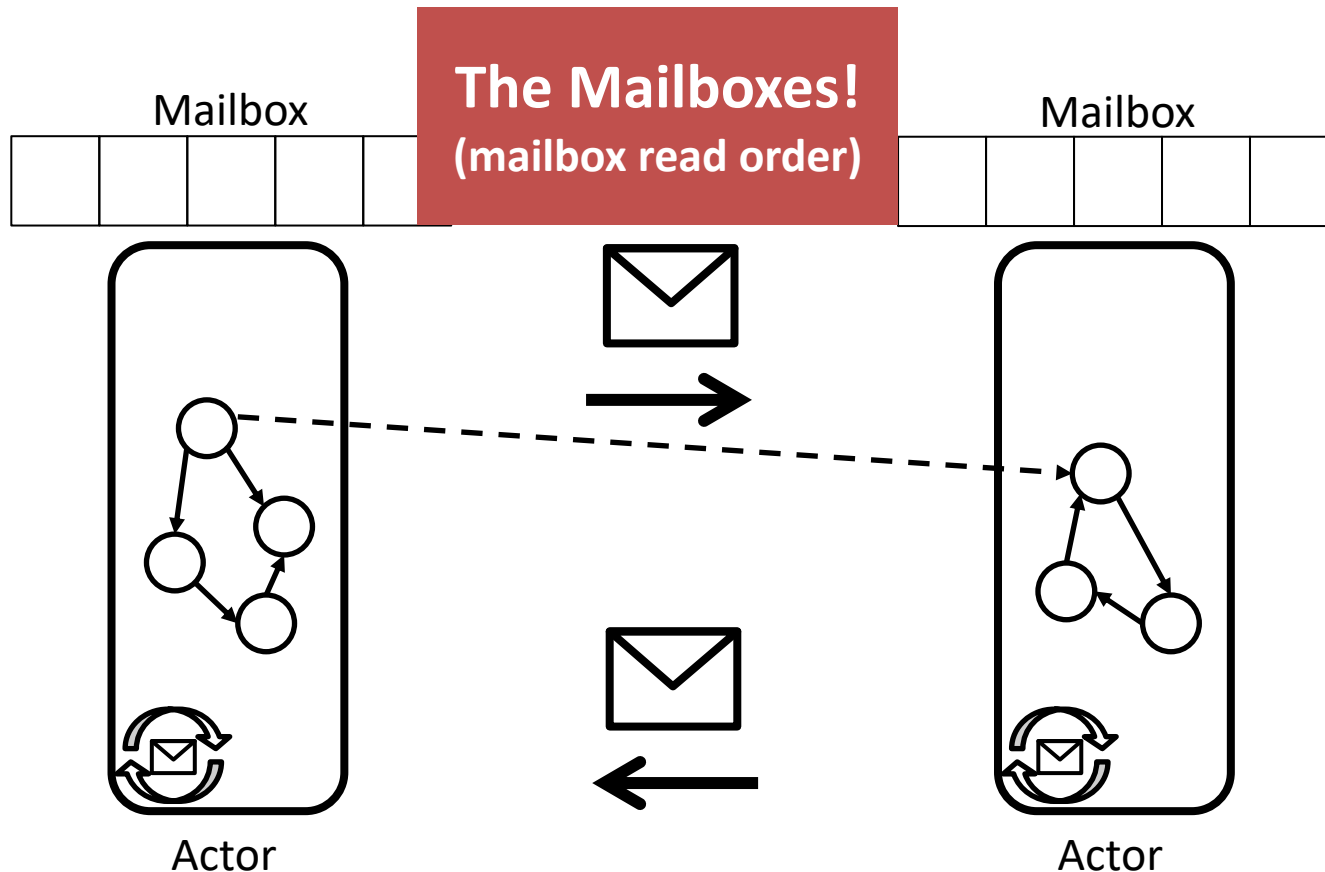


**Capturing High-level Nondeterminism in Concurrent Programs for Practical Concurrency Model Agnostic Record & Replay** D. Aumayr et al. The Art, Science, and Engineering of Programming, **Programming**, 2021.

**Efficient and Deterministic Record & Replay for Actor Languages** D. Aumayr et al. Proceedings of the 15th International Conference on Managed Languages and Runtimes, **ManLang'18**.

**How is that going to work  
agnostic to concurrency models?**

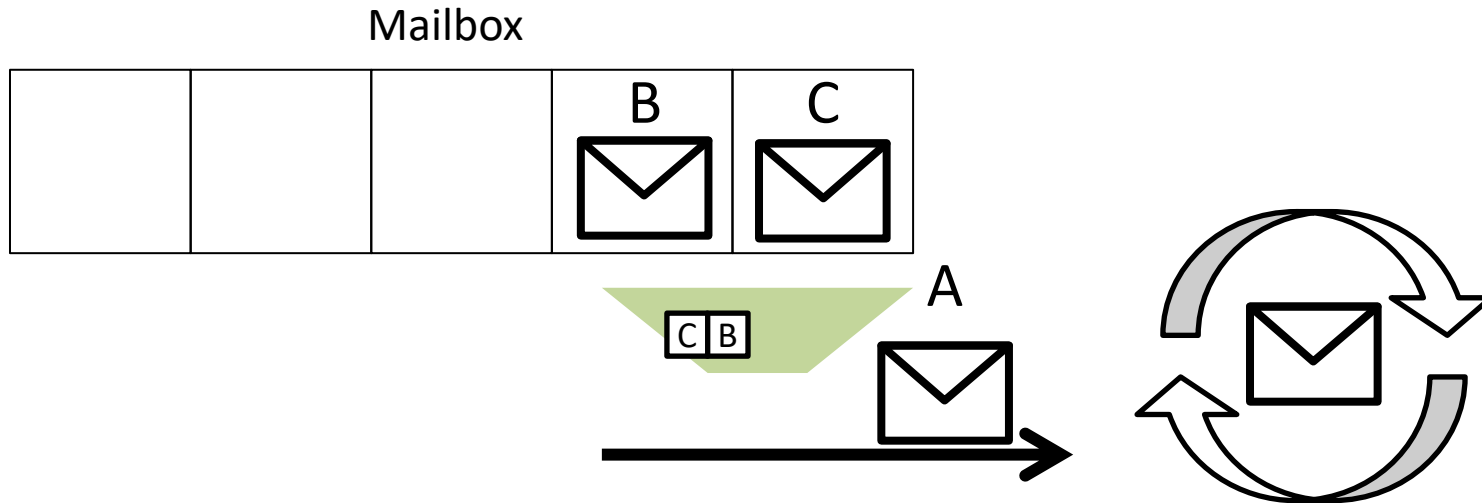
# Looking at Communicating Event Loops



**What are the Points of Non-determinism?**

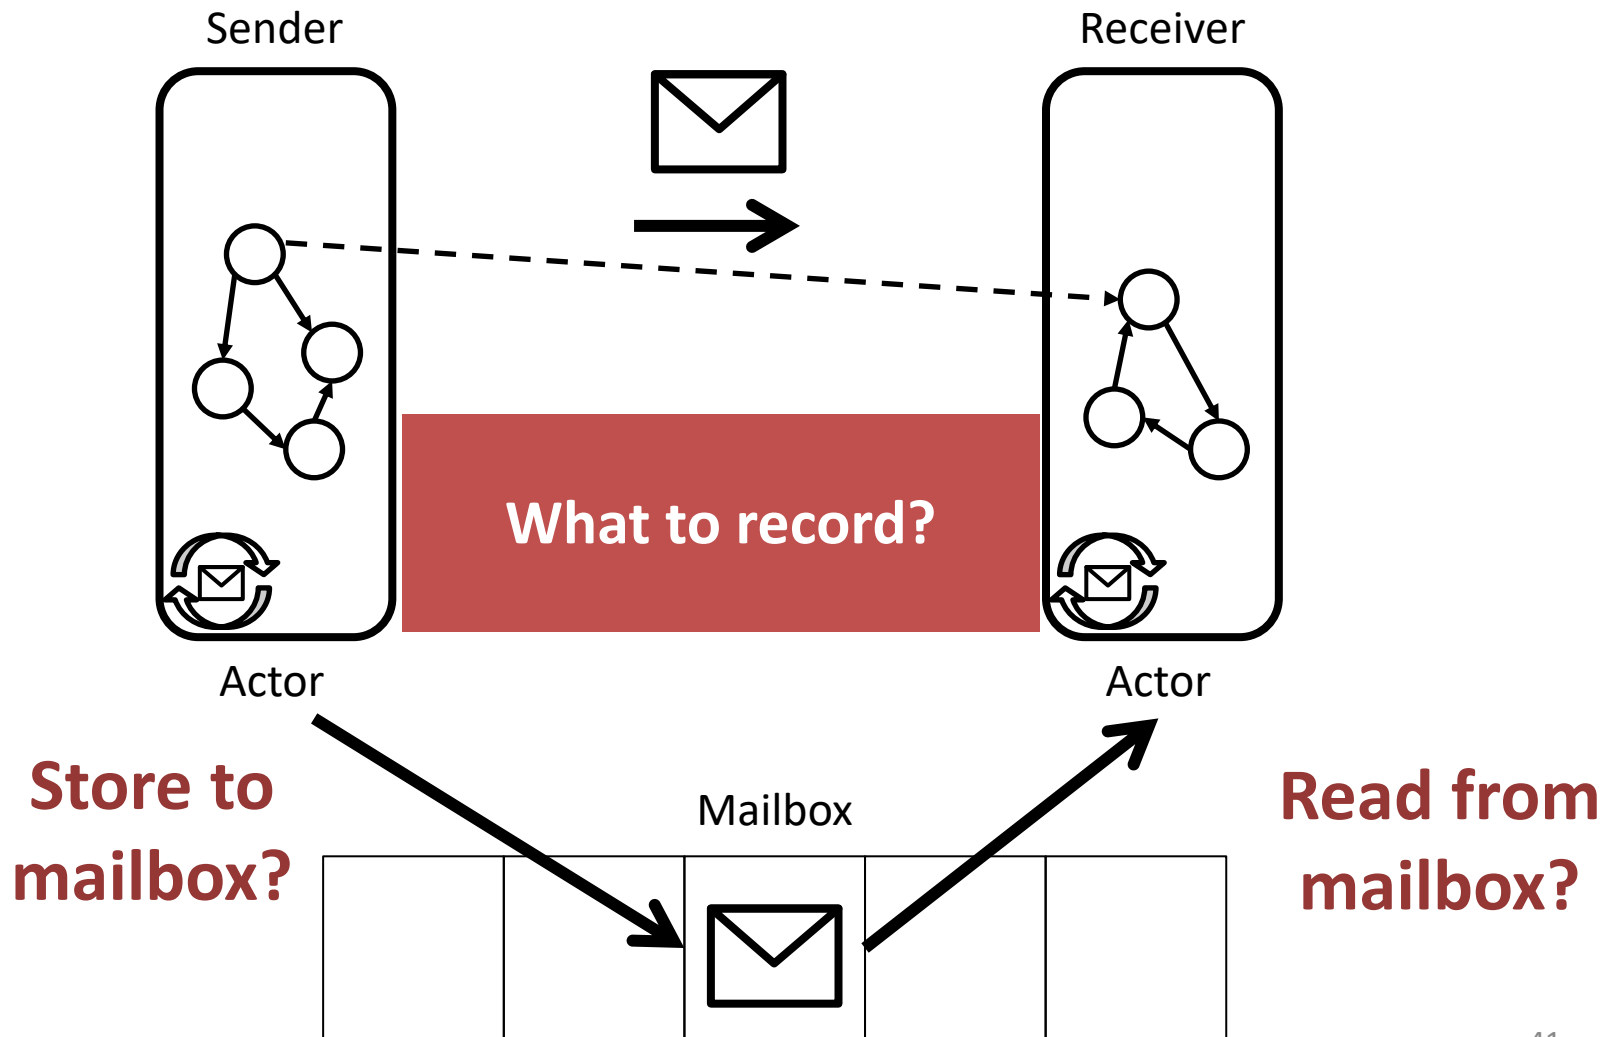


# Communicating Event Loops



**Replay messages in same  
order as originally**

# Recording Non-determinism in Communicating Event Loops



For Communicating Event Loops

**Sender-side and Receiver-Side**

**Recording are**

most interesting bit

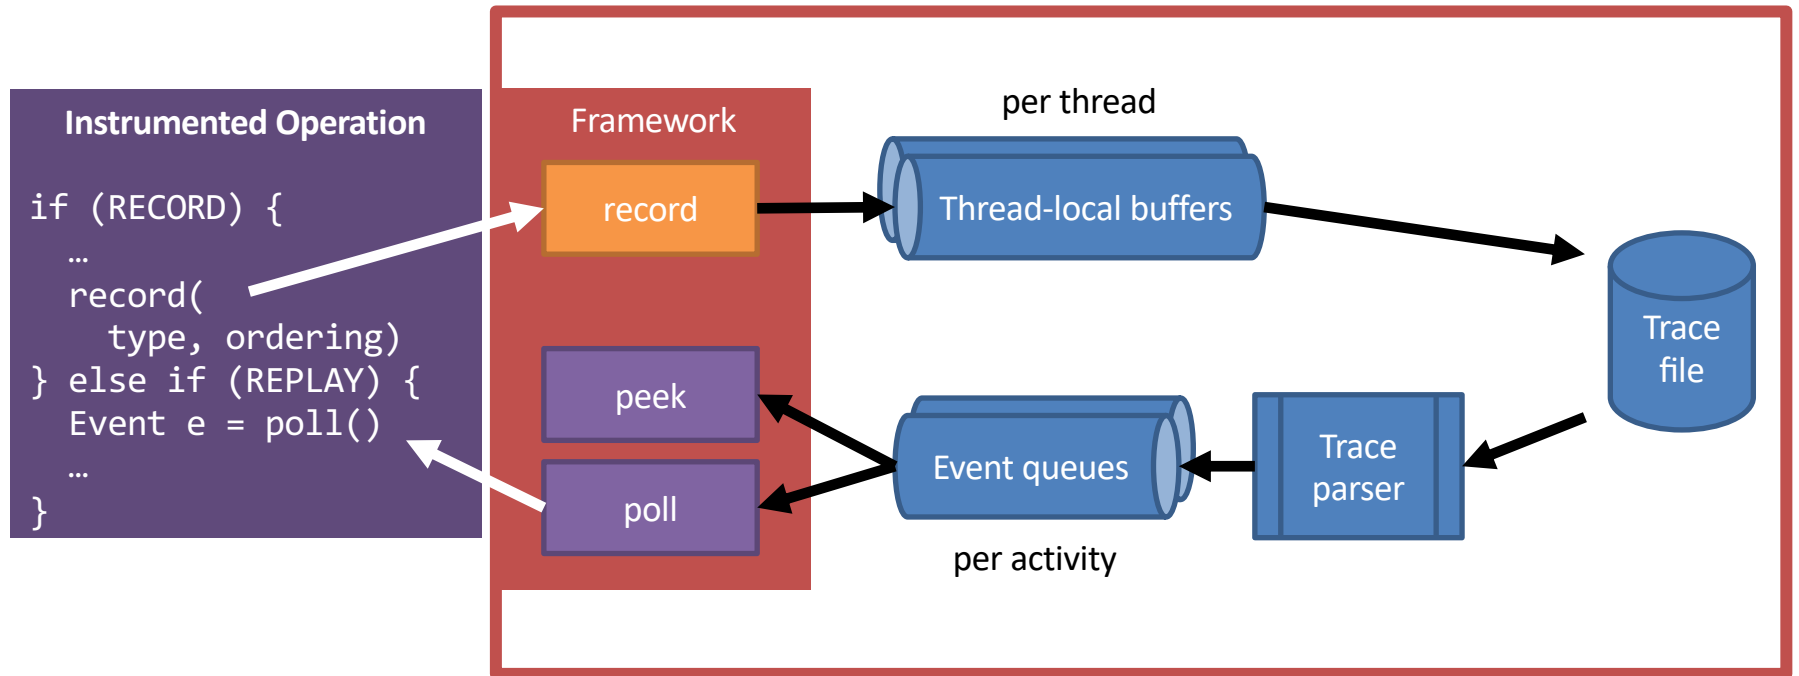
**“Functionally Equivalent”**

with complexity  
and performance trade-offs

# Overview for Concurrency Models

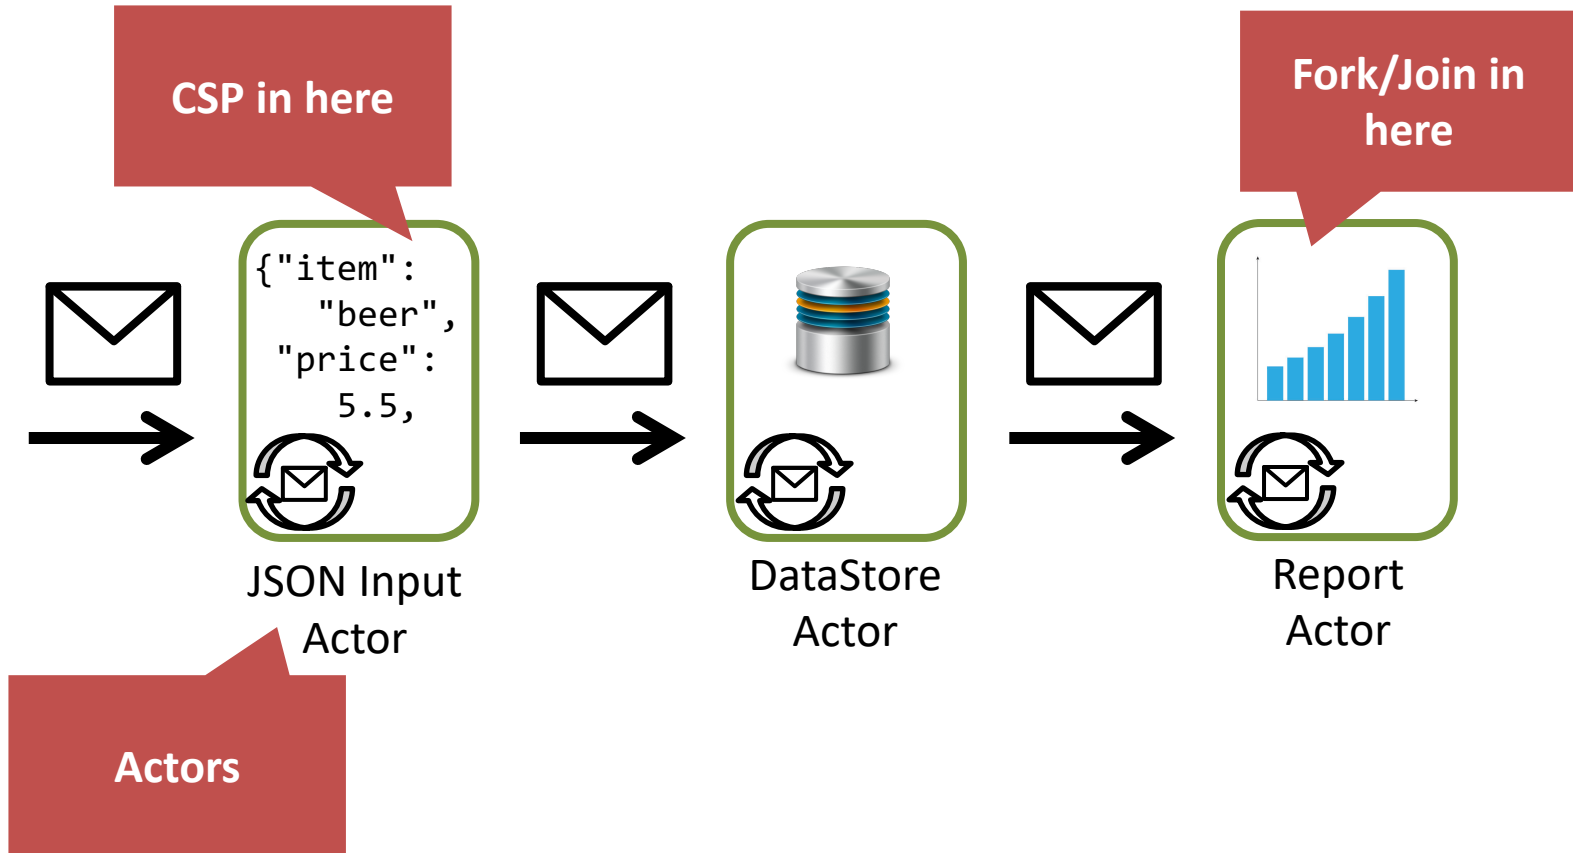
Model	Activities	Passive Entities	Non-determinism
Communicating Event Loops	Actor	Promise, Message	Message order per actor
Threads & Locks	Thread	Lock, Condition	Order of lock acquisitions
Communicating Sequential Processes	Process	Channel	Order of channel reads/writes
Software Transactional Memory	Transaction	-	Commit order

# Model Agnostic Framework



**Agnostic of  
Concurrency Models**

# Allows us to Record&Replay a Multi-Paradigm Application

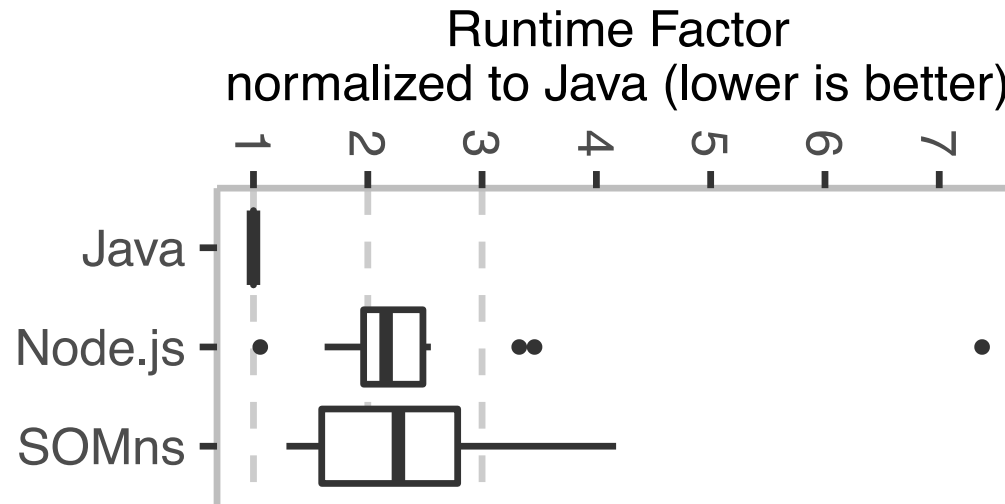




# **SOM<sub>NS</sub>: A NEWSPEAK FOR CONCURRENCY RESEARCH**

Newspeak: [newspeaklanguage.org](http://newspeaklanguage.org)  
SOM<sub>NS</sub>: [github.com/smarr/SOMns](https://github.com/smarr/SOMns)

# Performance: Baselines



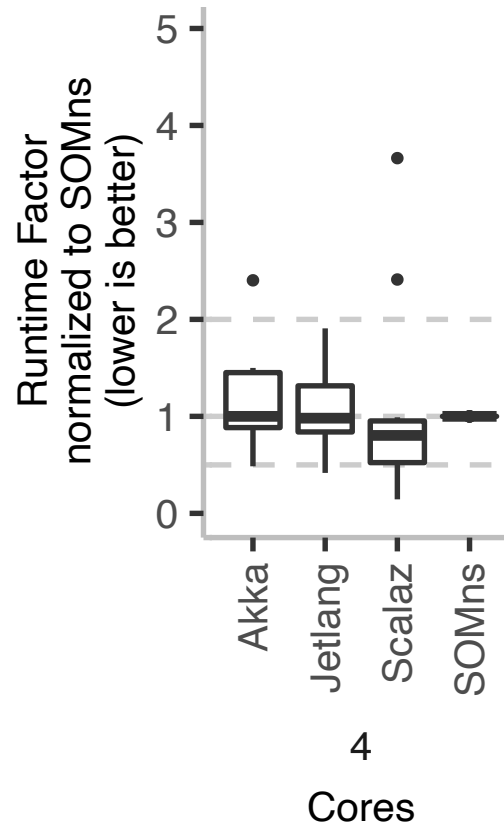
**SOM<sub>NS</sub> is on level of  
optimized dynamic  
languages!**

Are We Fast Yet: Cross-Language Comparison

<https://github.com/smarr/are-we-fast-yet#readme>



# Performance: Baselines

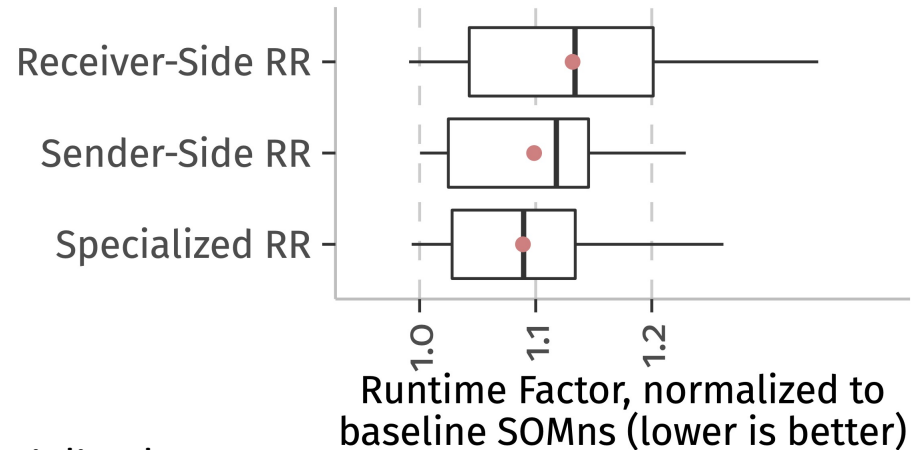


**Competitive  
with JVM actor  
frameworks!**

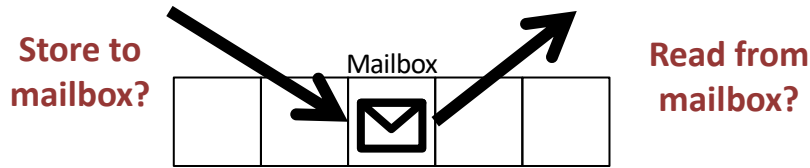
# Overhead of Recording Actors for Replay

Overhead on Savina benchmarks  
over execution without recording (geometric)

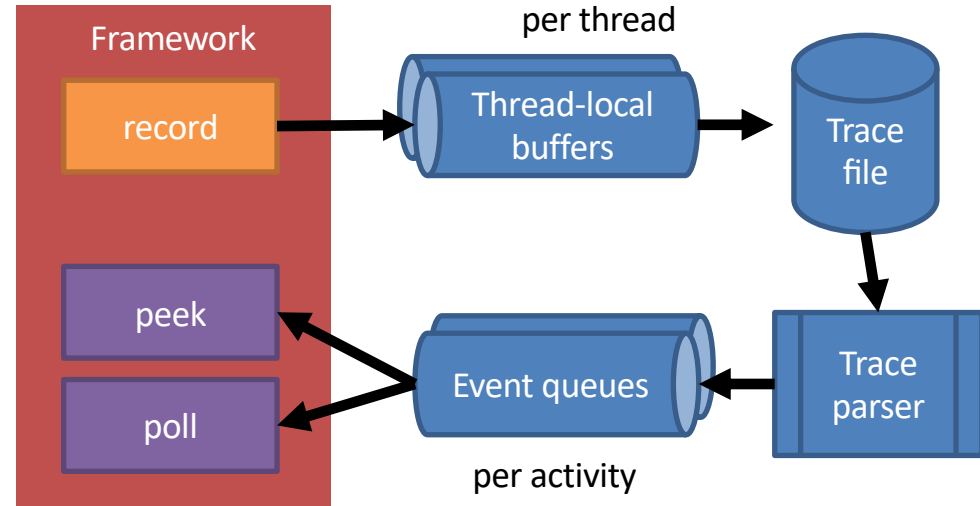
- Specialized: 7.89%  
min. -21.42%, max. 36.29%  
(specialized to actors,  
without support for  
other concurrency models)
- Sender-side: 7.82%  
min. -17.84%, max. 41.23%
  - Performance is competitive with specialized  
implementation
- Receiver-side: 13.23%  
min. -19.33%, max. 53.1%
  - Not as optimized as specialized



# Agnostic Record&Replay is Practical!

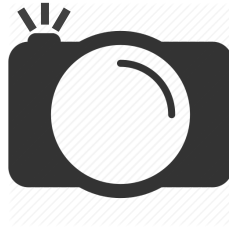


**Capture Non-determinism  
Per Concurrency Model**



**Keep Framework  
Agnostic**



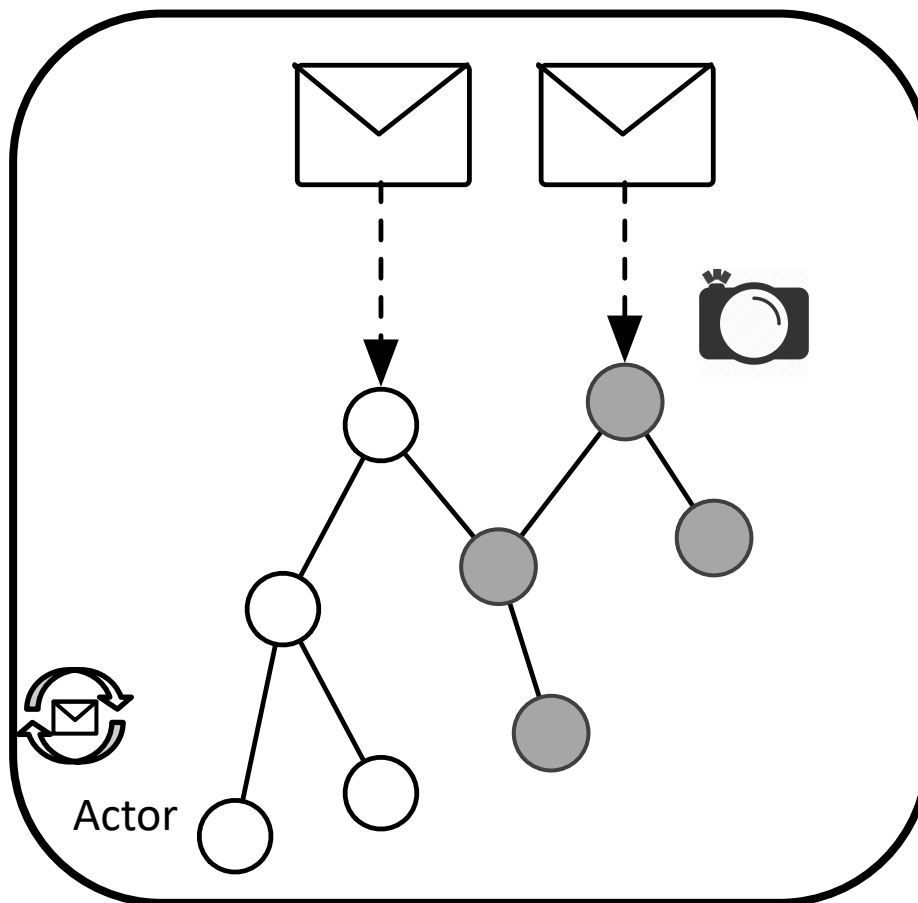


Snapshotting Actor Systems without Stopping Them

**LONG AND HUGE TRACES MAKE  
REPLAY IMPRACTICAL**

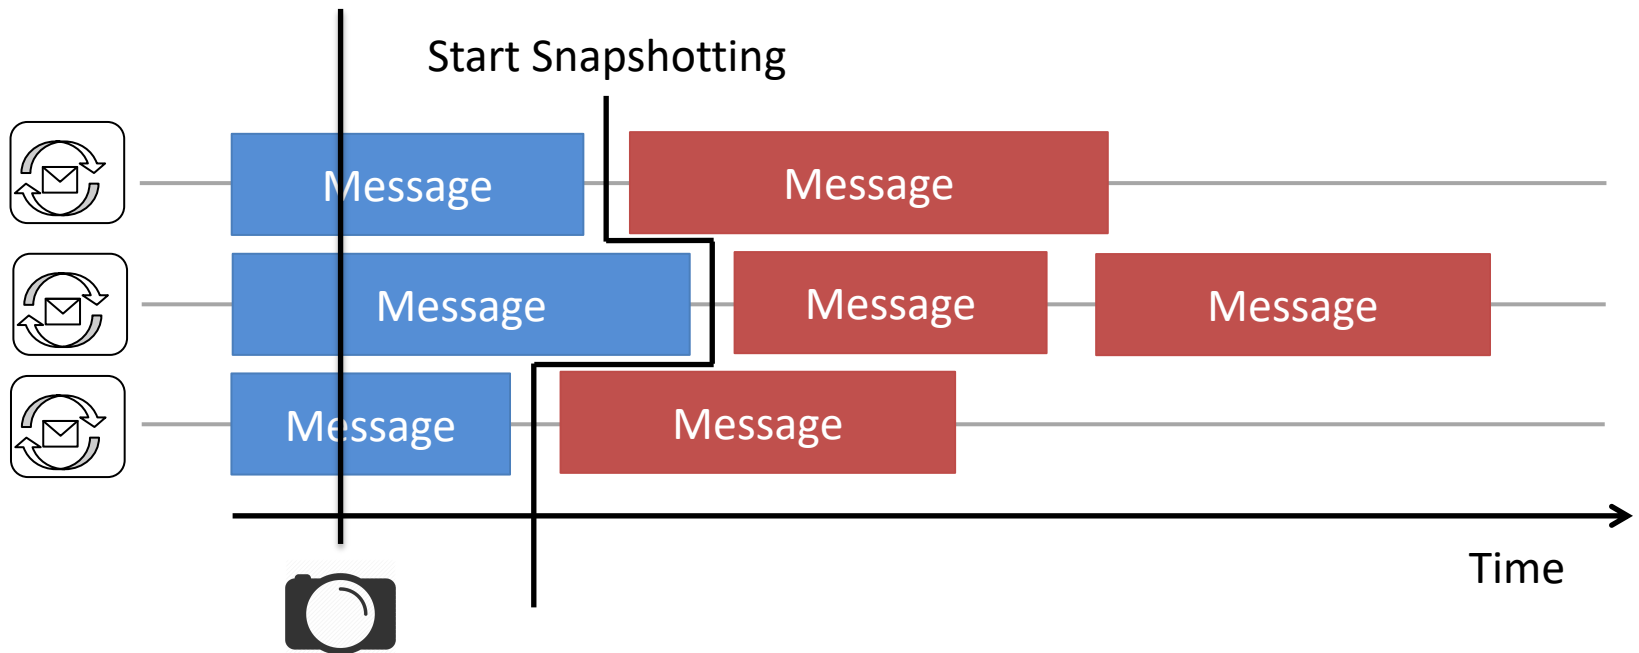


# Asynchronous and Partial Heap Snapshots



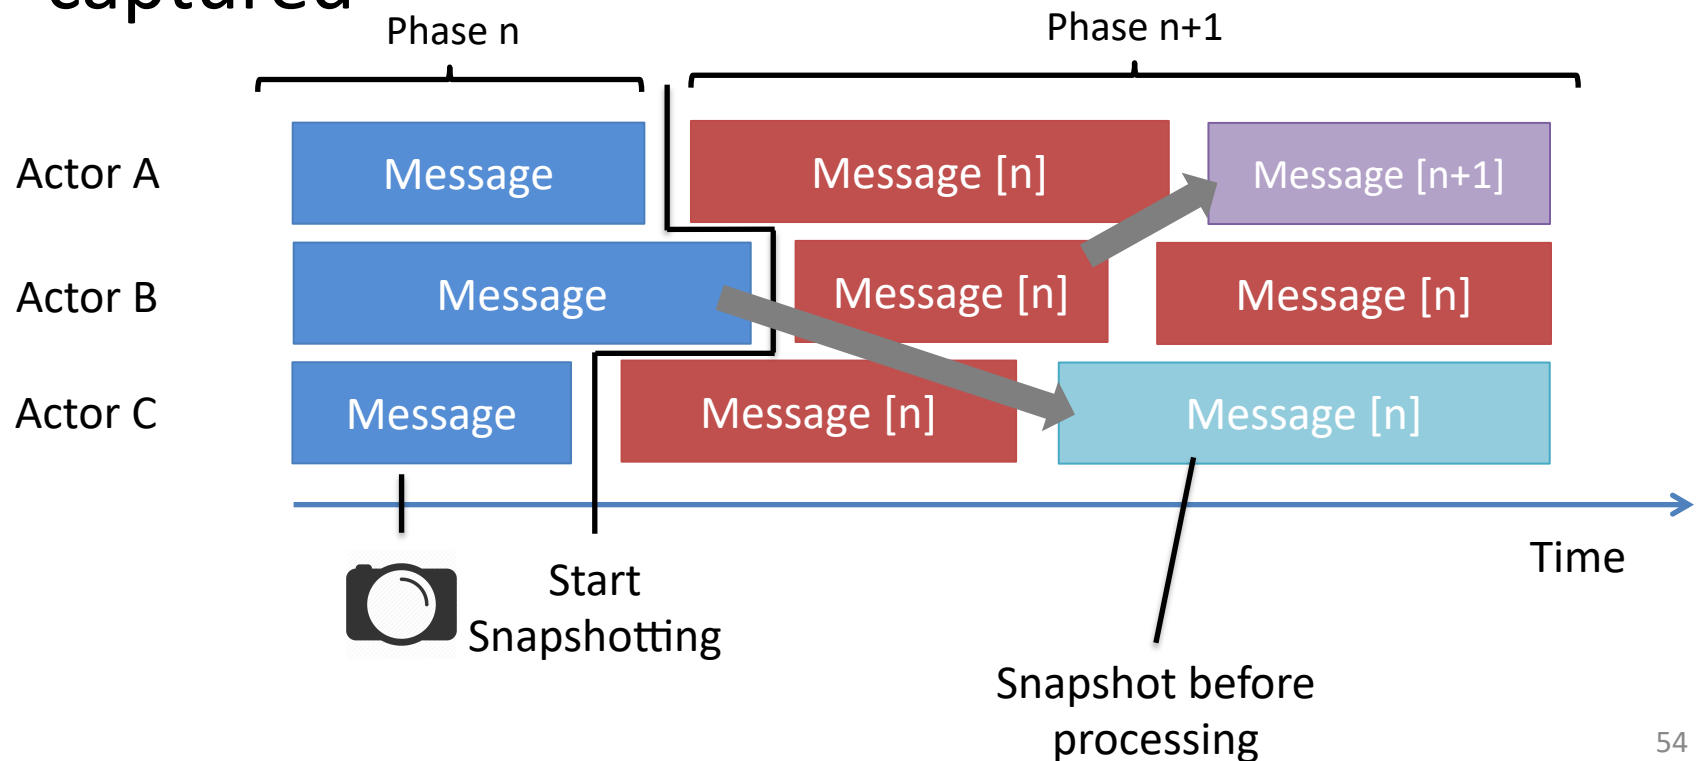
snapshot on message receive  
but only objects reachable from a message

# Snapshotting without Global Synchronization

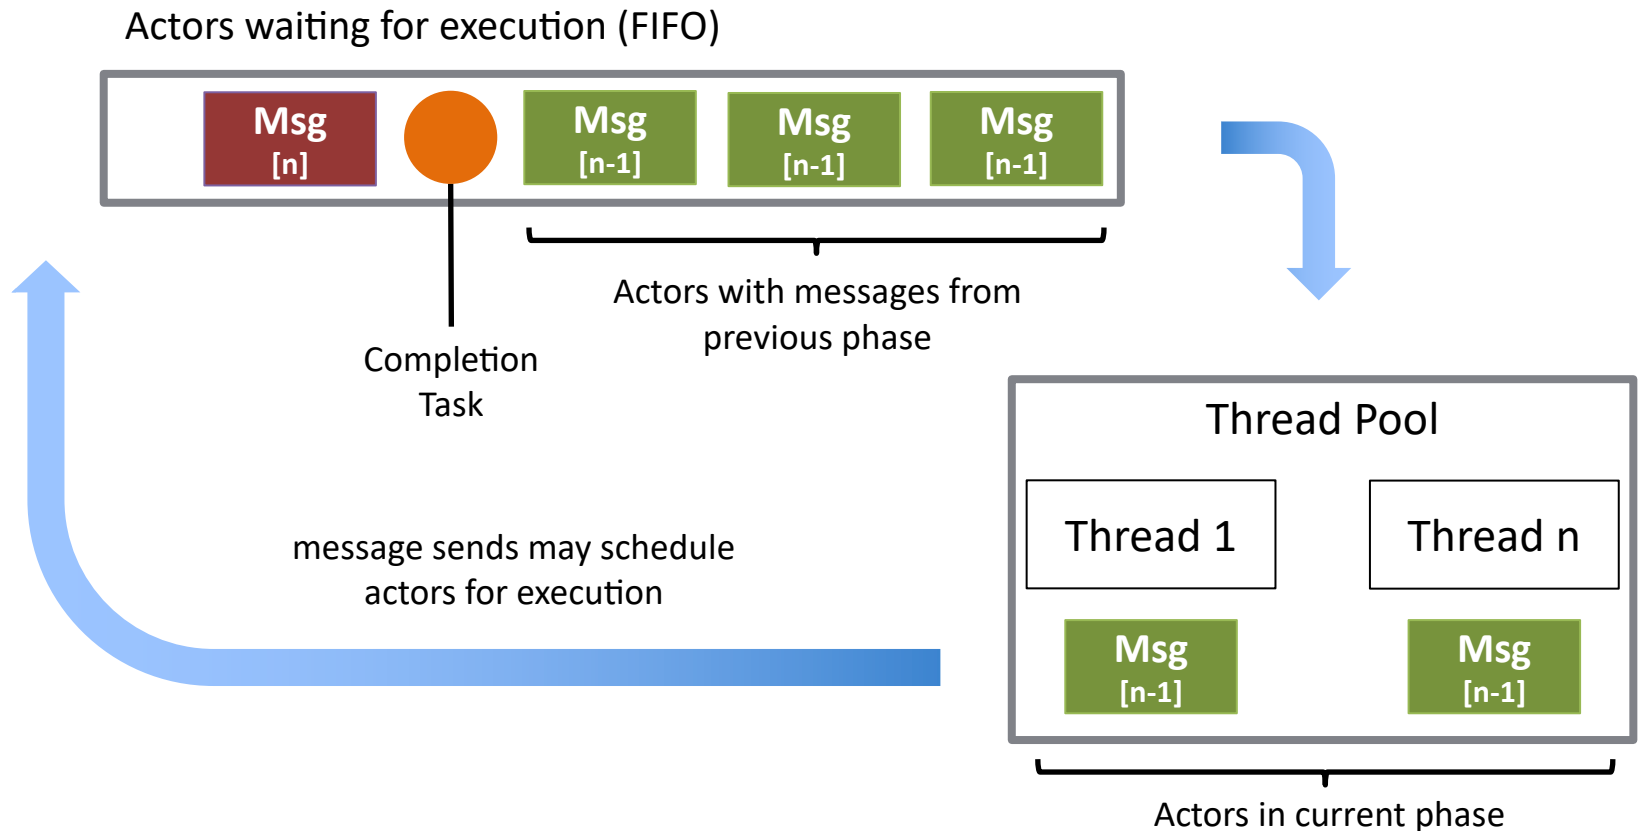


# Detecting Message Crossovers

- Attach send phase number to messages
- Messages sent in Phase n (previous) are captured

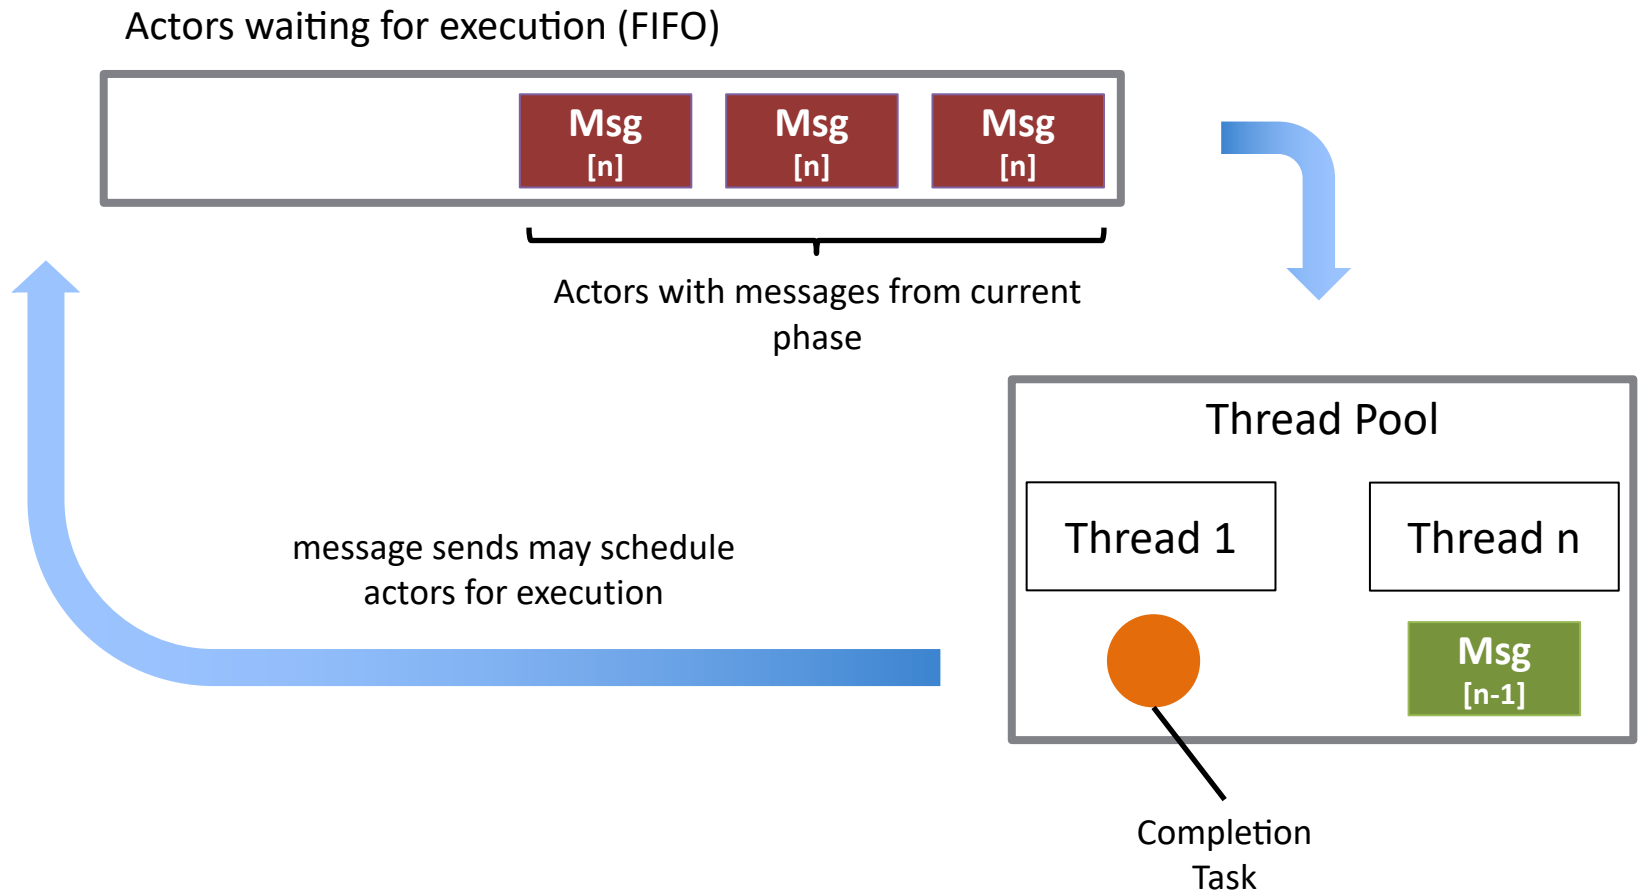


# Detecting Snapshot Completion (2)



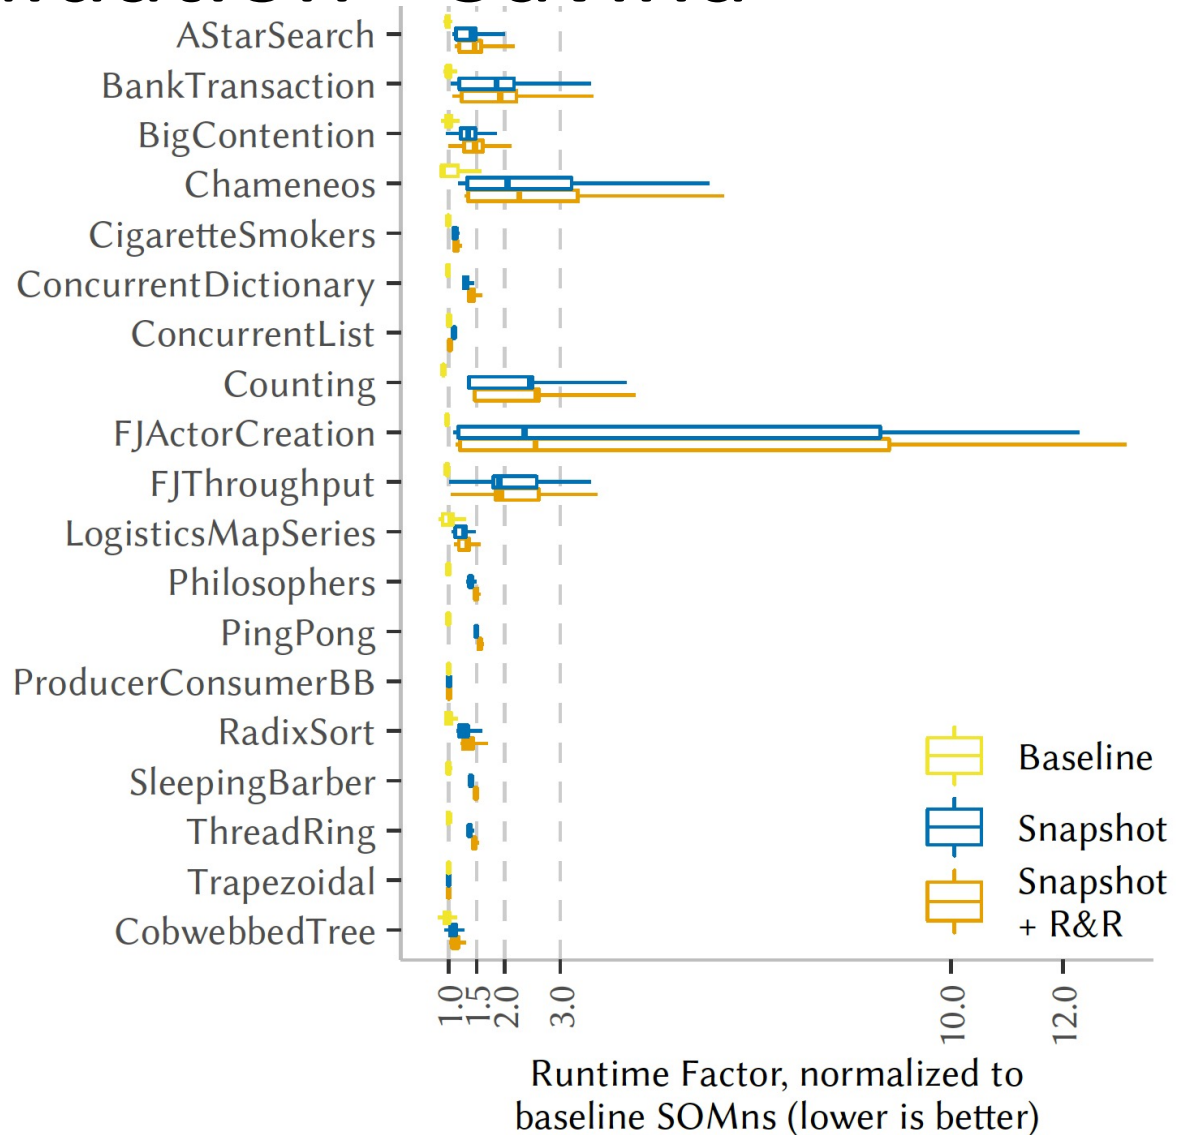


# Detecting Snapshot Completion (3)



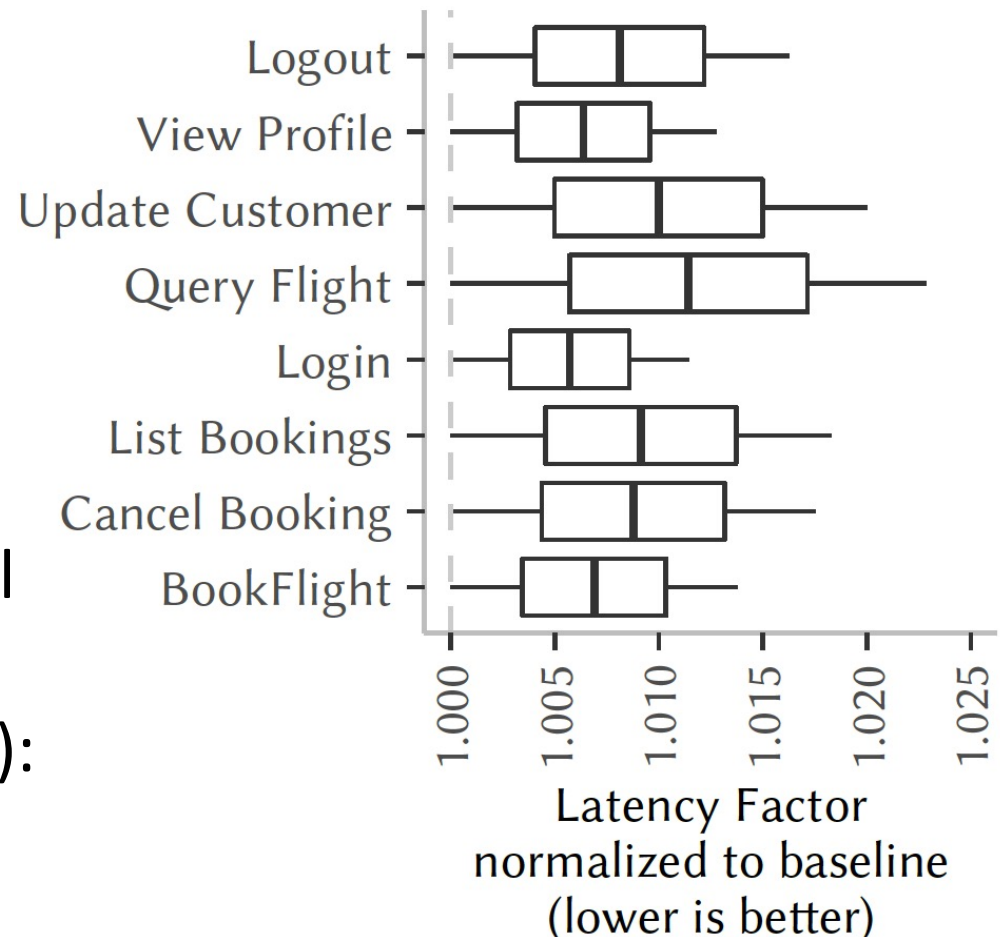
# Evaluation - Savina

- Snapshot every second iteration
- Worst-case scenario

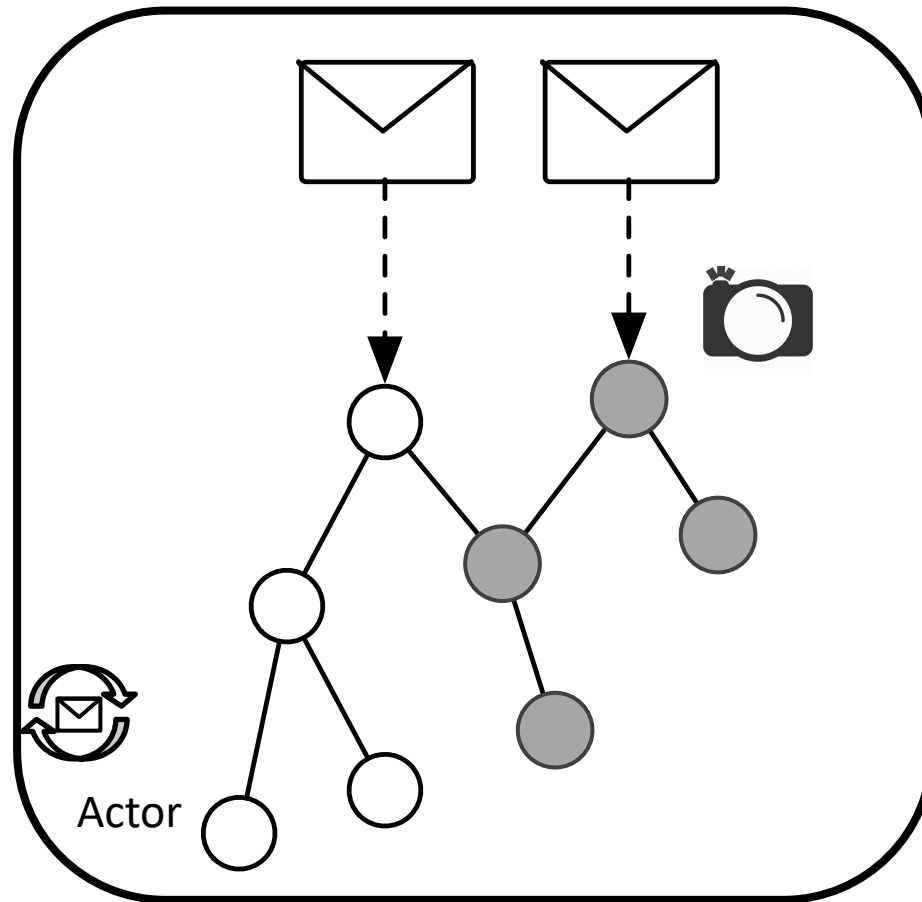


# Evaluation – AcmeAir Web Application

- Snapshot every 1000 requests
- Latency increases minimally (1,66% geo mean)
- 20 Million requests total
- Slow requests (> 100ms): 5.43% increase (0.007% of total requests)



# Snapshots can be Low-Overhead, Without Stop-the-World Pause



If it fails only 1 in 10 times, can we avert failure?

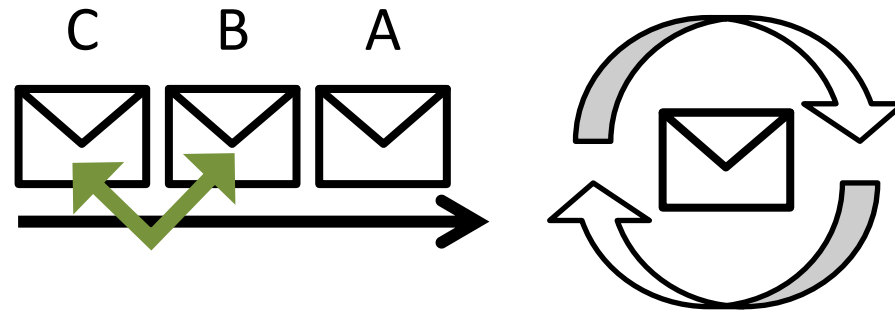
## BUG MITIGATION





# Bug Mitigation: Basic Idea

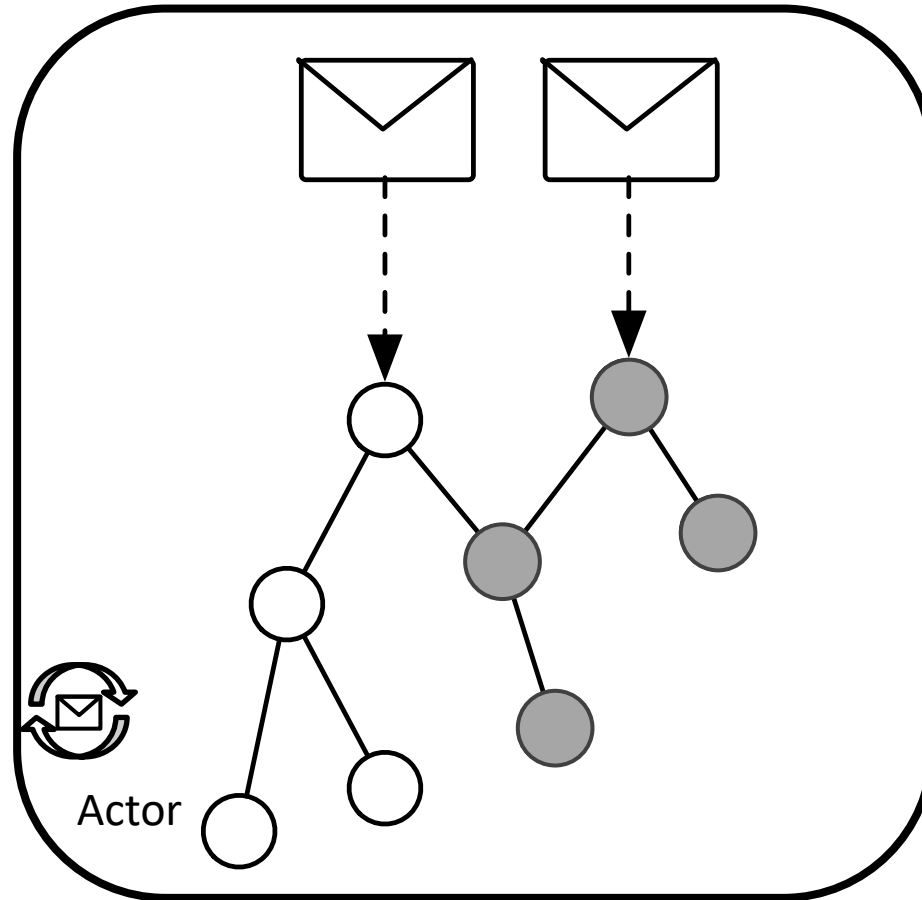
## Detect Event Races At Run Time



Order A -> B -> C problematic?

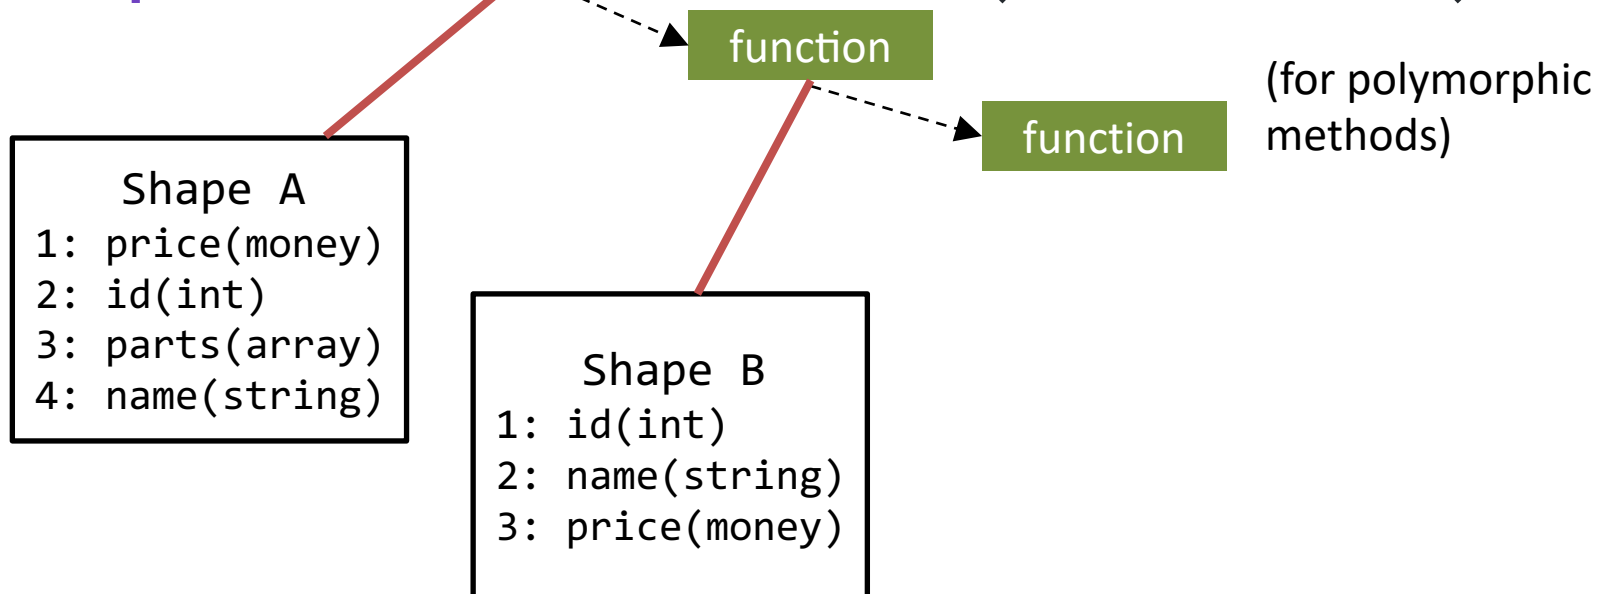
Let's swap them!

# Messages Usually Access Predictable Parts of the Heap



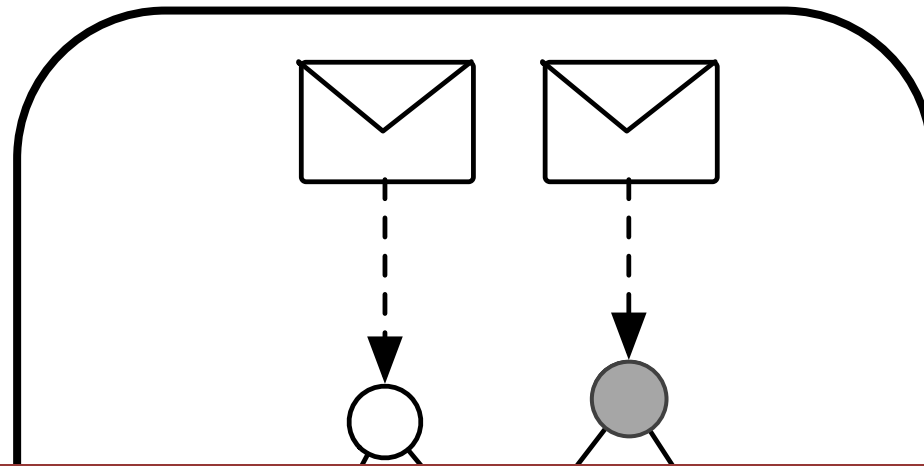
# Use Existing VM Techniques to Minimize Race Detection Overhead

`product.setPrice(newPrice)`





# Restrict Monitoring to Parts that can Race



**Very Early, but:  
Heap Access Patterns promising for  
light-weight, low-precision  
race-possibility detection**

# **WRAP-UP/CONCLUSION**

We're Looking for a Postdoc!

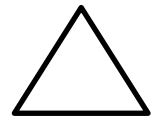
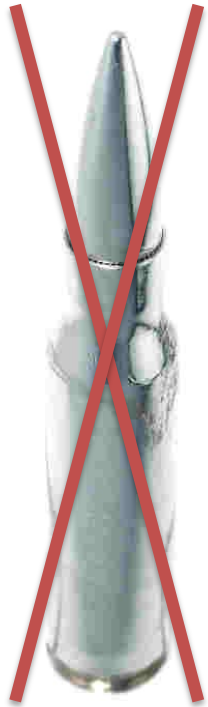
**Job Ad**

# Project CaMELot: Catch and Mitigate Event-Loop Concurrency Issues

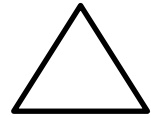


**Please get  
in touch!**

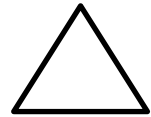
# Maybe there are no Silver Bullets?



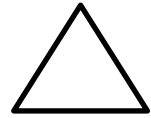
**Actors**



**CSP**



**Locks, Monitors, ...**



**Fork/Join**



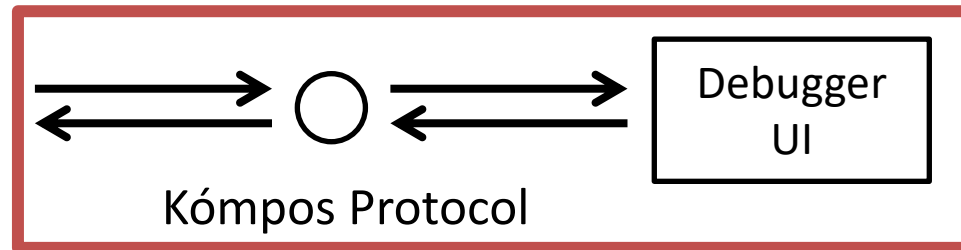
**Transactional Memory**



**Data Flow**

...

# Debuggers can be Great for High-level Concurrency Models!

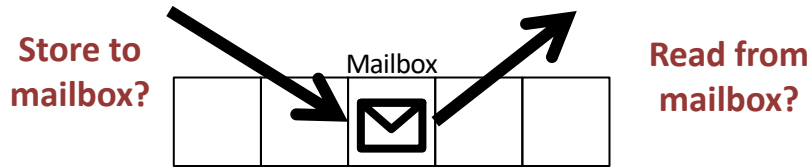


**Make tools agnostic**

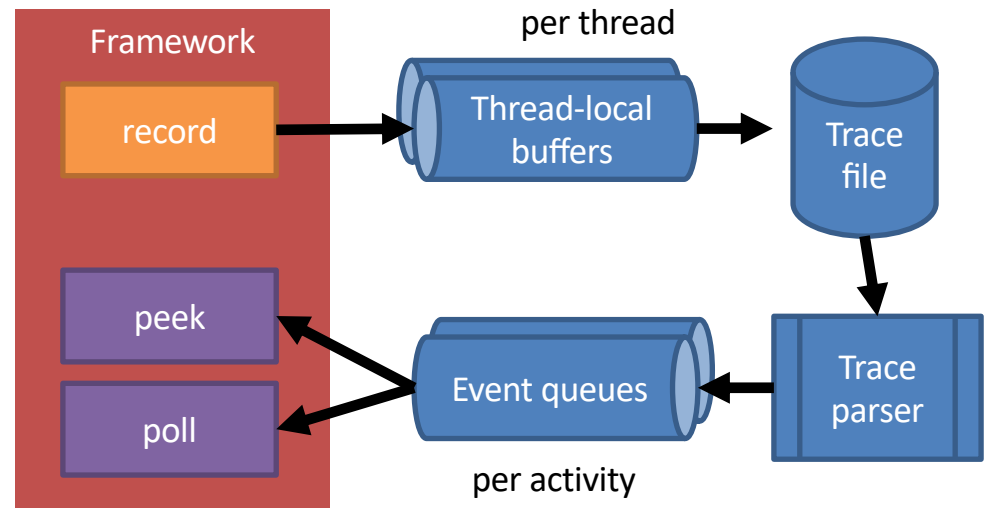
```
promise resolver  
promise resolution  
prom whenResolved: [:r |  
r println ].
```

**Offer the Key Features  
as Breakpoints/Steps**

# Agnostic Record&Replay is Practical!

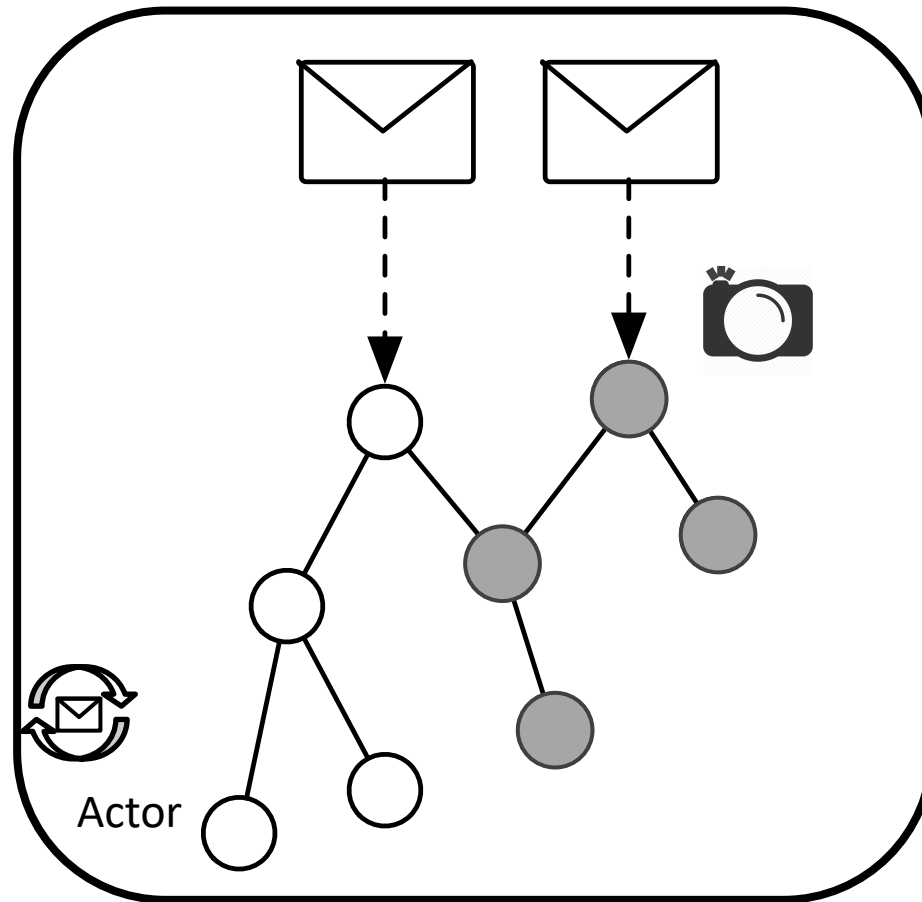


## Capture Non-determinism Per Concurrency Model

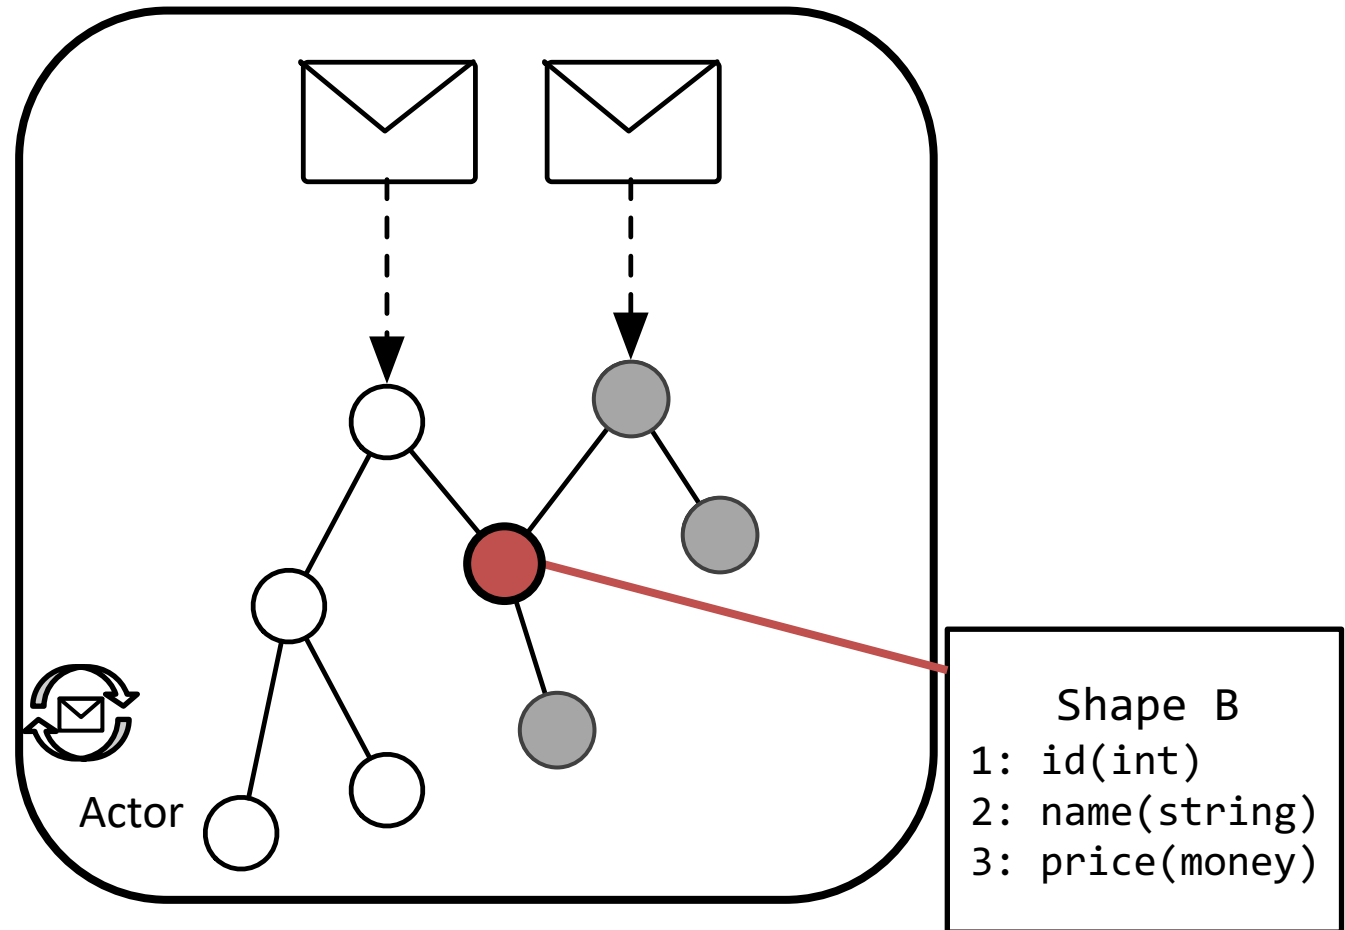


**Keep Framework  
Agnostic**

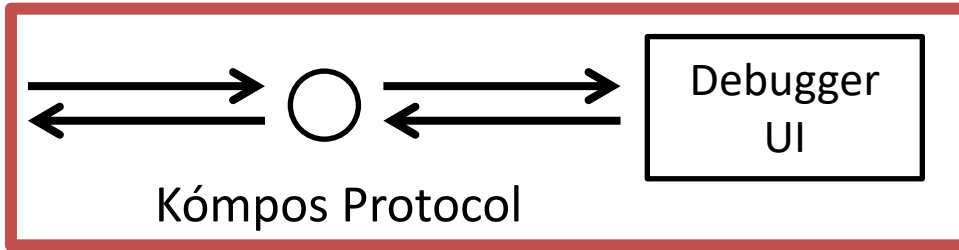
# Snapshots can be Low-Overhead, Without Stop-the-World Pause



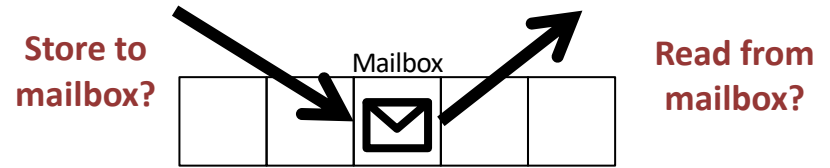
# And maybe, we can use it to do race-mitigation!



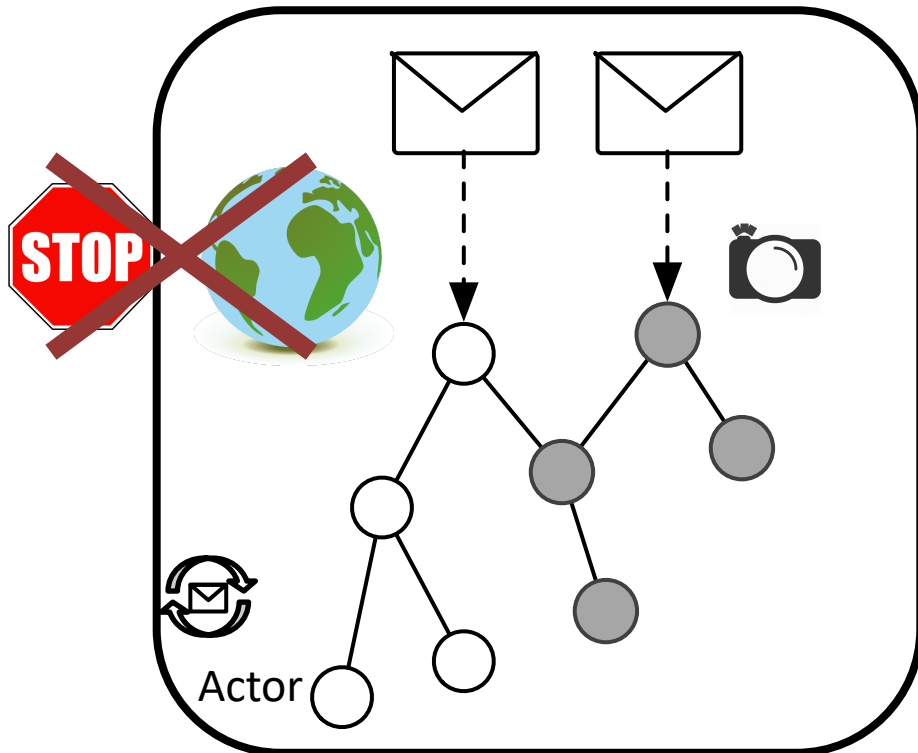




## Make tools agnostic



## Capture Non-determinism Per Concurrency Model



And don't stop the world  
for snapshotting!



# References

- **Capturing High-level Nondeterminism in Concurrent Programs for Practical Concurrency Model Agnostic Record & Replay** ([pdf](#))  
D. Aumayr, S. Marr, S. Kaleba, E. Gonzalez Boix, H. Mössenböck, <Programming>, p. 39, AOSA Inc., 2021. doi: [10.22152/programming-journal.org/2021/5/14](https://doi.org/10.22152/programming-journal.org/2021/5/14)
- **Asynchronous Snapshots of Actor Systems for Latency-Sensitive Applications** ([pdf](#))  
D. Aumayr, S. Marr, E. Gonzalez Boix, H. Mössenböck, **MPLR'19**, p. 157–171, ACM, 2019. doi: [10.1145/3357390.3361019](https://doi.org/10.1145/3357390.3361019)
- **Efficient and Deterministic Record & Replay for Actor Languages** ([pdf](#))  
D. Aumayr, S. Marr, C. Béra, E. Gonzalez Boix, H. Mössenböck, **ManLang'18**, ACM, 2018. doi: [10.1145/3237009.3237015](https://doi.org/10.1145/3237009.3237015)
- **A Concurrency-Agnostic Protocol for Multi-Paradigm Concurrent Debugging Tools** ([pdf](#))  
S. Marr, C. Torres Lopez, D. Aumayr, E. Gonzalez Boix, H. Mössenböck, **DLS'17**, p. 3–14, ACM, 2017. doi: [10.1145/3133841.3133842](https://doi.org/10.1145/3133841.3133842)
- **Kómpos: A Platform for Debugging Complex Concurrent Applications** ([pdf](#))  
S. Marr, C. Torres Lopez, D. Aumayr, E. Gonzalez Boix, H. Mössenböck, <Programming Demo'17>, p. 2:1–2:2, ACM, 2017. Demo. doi: [10.1145/3079368.3079378](https://doi.org/10.1145/3079368.3079378)
- **A Study of Concurrency Bugs and Advanced Development Support for Actor-based Programs** ([pdf](#))  
C. Torres Lopez, S. Marr, H. Mössenböck, E. Gonzalez Boix, **AGERE!'16 (LNCS)**, p. 155–185, Springer, 2018. doi: [10.1007/978-3-030-00302-9\\_6](https://doi.org/10.1007/978-3-030-00302-9_6)
- **Towards Advanced Debugging Support for Actor Languages: Studying Concurrency Bugs in Actor-based Programs** ([pdf](#))  
C. Torres Lopez, S. Marr, H. Mössenböck, E. Gonzalez Boix, **AGERE!'16**, 2016.