

Trace Partitioning and Local Monitoring for Asynchronous Components

Duncan Attard¹ and Adrian Francalanza¹ · 8th September 2017

¹CS, ICT, University of Malta, Malta

Outline

1. Monitoring Asynchronous Component Systems
2. Trace Partitioning and Local Monitoring in Practice

Introduction

Static pre-deployment techniques have limited applicability.

RV allows us to verify systems post-deployment

- it uses dedicated programs called **monitors**
- to analyse the **current execution** of system
- in a dynamic and **incremental** manner...

...**but** depends on the execution path taken by the system.

The traditional RV setup

Monitor generation

The traditional RV setup

Monitor generation

Property

The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow

The traditional RV setup


Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

The traditional RV setup

Monitor generation


Property \rightarrow *compilation* \rightarrow Monitor



The traditional RV setup

Monitor generation


Property \rightarrow compilation \rightarrow Monitor



The traditional RV setup

Monitor generation

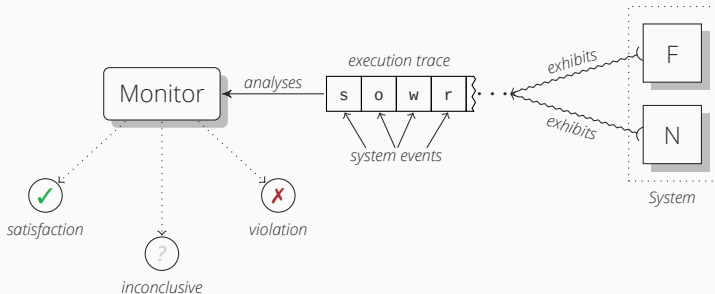
Property \rightarrow compilation \rightarrow Monitor



The traditional RV setup

Monitor generation

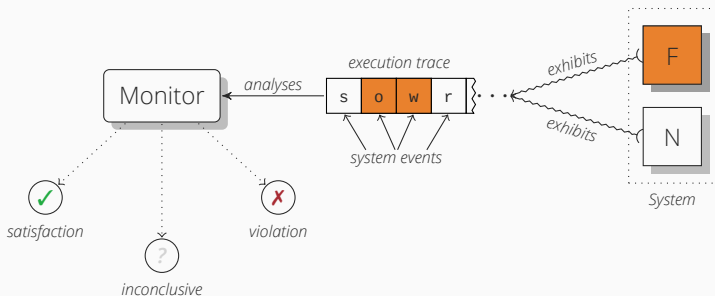
Property \rightarrow compilation \rightarrow Monitor



The traditional RV setup

Monitor generation

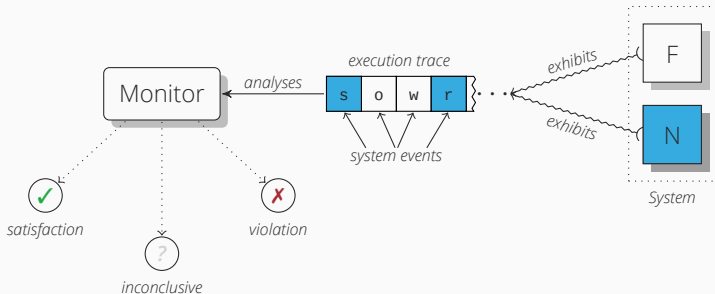
Property \rightarrow compilation \rightarrow Monitor



The traditional RV setup

Monitor generation


Property \rightarrow compilation \rightarrow Monitor



The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

“The system cannot send messages at startup.”




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

“The system *cannot* send messages at startup.”




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

"The system *cannot* send messages at startup."




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

"The system *cannot* send messages at startup."




The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

"The system cannot send messages at startup."




The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

"The system cannot send messages at startup."




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

“The system *cannot* send messages at startup.”




The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

"The system *cannot send messages at startup.*"




The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

“The system *cannot* send messages at startup.”




The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

"The system cannot send messages at startup."




The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

“The system cannot send messages at startup.”




The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

"The system cannot send messages at startup."




The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

"The system cannot send messages at startup."




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

“The system *cannot* send messages at startup.”




The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

"The system *cannot* send messages at startup."




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

“The system *cannot* send messages at startup.”




The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

“At startup, the system cannot open files **and** send messages .”




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

“At startup, the system *cannot* **o**pen files **and** **s**end messages .”




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

“At startup, the system *cannot* **o**pen files **and** **s**end messages .”




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

“At startup, the system *cannot* open files **and** send messages .”




The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

“At startup, the system cannot open files **and** send messages .”




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

“At startup, the system cannot open files **and** send messages .”




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

“At startup, the system cannot open files **and** send messages .”




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

“At startup, the system cannot open files **and** send messages .”




The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

“At startup, the system can neither open files **nor** send messages.”




The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

“At startup, the system can *neither* open files *nor* send messages.”




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

“At startup, the system can *neither open files nor send messages.*”




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

“At startup, the system can *neither* open files *nor* send messages.”




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

“At startup, the system can *neither* open files *nor* send messages.”




The traditional RV setup

Monitor generation

Property \rightarrow compilation \rightarrow Monitor

“At startup, the system can *neither* open files *nor* send messages.”




The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

"Message sending cannot issue any errors."




The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

“Message sending cannot issue any errors.”




The traditional RV setup

Monitor generation

Property \rightarrow *compilation* \rightarrow Monitor

"Message sending cannot issue any **e**rrors." ...?




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

Stop treating the system as one monolithic block...




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

Stop treating the system as one monolithic block...




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

Stop treating the system as one monolithic block...




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

Stop treating the system as one monolithic block...




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

“The system cannot send messages at startup.”




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

“The system cannot send messages at startup.”




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

"The system cannot send messages at startup."




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

“The system cannot send messages at startup.”




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

*“At startup, the system cannot open files **and** send messages.”*




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

“At startup, the system *cannot open files and send messages.*”




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

“At startup, the system *cannot open files and send messages.*”




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

“At startup, the system *cannot open files and send messages.*”




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

*“At startup, the system cannot open files **and** send messages.”*




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

*“At startup, the system cannot open files **and** send messages.”*




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

*“At startup, the system cannot open files **and** send messages.”*




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

*“At startup, the system cannot open files **and** send messages.”*




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

*“At startup, the system can neither open files **nor** send messages.”*




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

"At startup, the system can *neither open files nor send messages.*"




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

*“At startup, the system can **neither open files nor send messages.**”*




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

*“At startup, the system can **neither open files nor send messages.**”*




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

*“At startup, the system can **neither open files nor send messages.**”*




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

*“At startup, the system can **neither open files nor send messages.**”*




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

“At startup, the system can *neither open files nor send messages.*”




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

“Message sending cannot issue any errors.”




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

“Message sending cannot issue any errors.”




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

“Message sending cannot issue any errors.”




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

“Message sending cannot issue any errors.”




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

“Message sending cannot issue any errors.”




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

“Message sending cannot issue any errors.”




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

*“At startup, the system cannot open files **and** send messages.”*




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

“At startup, the system *cannot open files and send messages.*”




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

“At startup, the system *cannot open files and send messages.*”




Partitioned traces

Approach valid but expensive

- We need an unbounded buffer to store inferred traces.
- Inferring new traces requires computational resources...

“At startup, the system *cannot open files and send messages.*”



Tracing vs. monitoring

Tracing differs from monitoring

- Trace partitioning operates at the **instrumentation** level...
...and deals with **what** events are extracted from **where**.
- Monitors deal with **how** these events are analysed.

Tracing vs. monitoring

Tracing differs from monitoring

- Trace partitioning operates at the **instrumentation** level...
...and deals with **what** events are extracted from **where**.
- Monitors deal with **how** these events are analysed.

Component monitors are but one method of trace analysis.

Tracing vs. monitoring

Tracing differs from monitoring

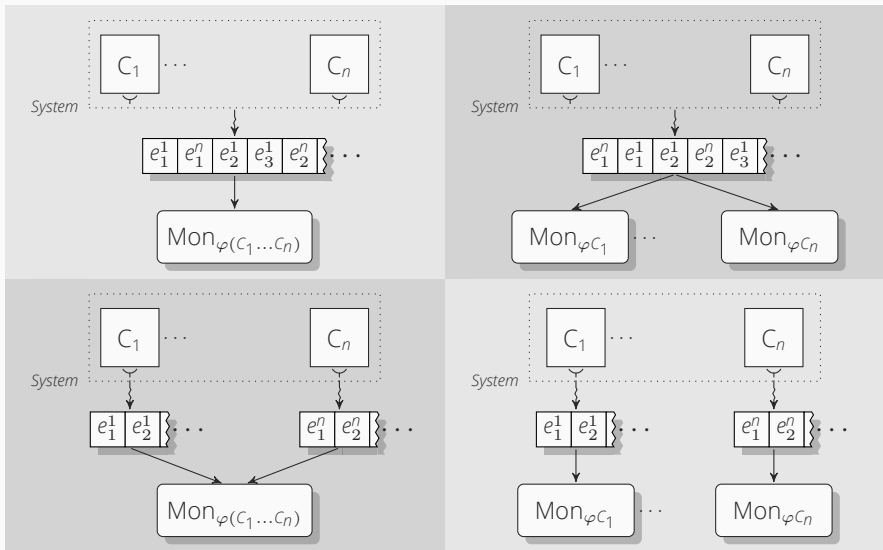
- Trace partitioning operates at the **instrumentation** level...
...and deals with **what** events are extracted from **where**.
- Monitors deal with **how** these events are analysed.

Component monitors are but one method of trace analysis.

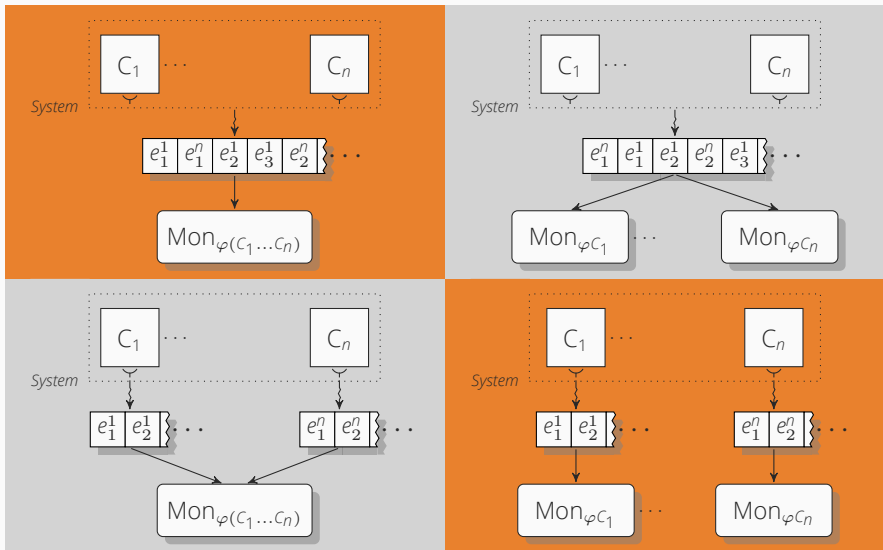
This separation of concerns is nice:

- it allows us to reason separately on both;
- it makes this reasoning manageable.

What and where, how...



What and where, how...



The advantages of trace partitioning

1. **compactly encodes** interleaving behaviour,
2. inherently **filters** extraneous events from a partition,
3. can be **selectively applied** to components of interest,
4. is easily applicable to non-replicated components,
5. **preserves the origin** of indistinguishable events,
6. may be naturally extended to distributed settings.

The advantages of trace partitioning

1. **compactly encodes** interleaving behaviour,
2. inherently **filters** extraneous events from a partition,
3. can be **selectively applied** to components of interest,
4. is easily applicable to non-replicated components,
5. **preserves the origin** of indistinguishable events,
6. may be naturally extended to distributed settings.

Outline

1. Monitoring Asynchronous Component Systems
2. Trace Partitioning and Local Monitoring in Practice

An empirical evaluation of local monitoring

The study was conducted on a third-party app called Ranch.

An empirical evaluation of local monitoring

The study was conducted on a third-party app called Ranch.

Aim: to test local monitoring in a **real-world** setting and

1. explore its **applicability** in industry-level apps,
2. investigate its **feasibility** in terms of performance.

An empirical evaluation of local monitoring

The study was conducted on a third-party app called Ranch.

Aim: to test local monitoring in a **real-world** setting and

1. explore its **applicability** in industry-level apps,
2. investigate its **feasibility** in terms of performance.

Focus: global properties that can be written as local ones.

An empirical evaluation of local monitoring

The study was conducted on a third-party app called Ranch.

Aim: to test local monitoring in a **real-world** setting and

1. explore its **applicability** in industry-level apps,
2. investigate its **feasibility** in terms of performance.

Focus: global properties that can be written as local ones.

Method: identified a set of components within Ranch and:


1. formulated a number of global properties,
2. monitored the components using global monitors,
3. reformulated the global properties as local ones,
4. monitored the same components using local monitors.

Experiments and results

100 components (recommended number for production use).

Global monitoring **consistently** led to resource exhaustion.

Local monitoring scaled well in relation to the **baseline**.




Experiments and results

< 5 components (used to be able to test global monitoring).

Global monitoring scaled poorly.

Local monitoring scaled well in relation to the baseline.



Conclusion

Trace partitioning \Rightarrow maximising information

- Extra **information** is extracted from the system execution.
- This helps **improve** the **monitor analysis**.

Conclusion

Trace partitioning \Rightarrow maximising information

- Extra **information** is extracted from the system execution.
- This helps **improve** the **monitor analysis**.

Local monitoring

- The instrumentation technique **fits local monitors** naturally.
- Local monitoring + partitioned traces = **lower overheads**.

Thank you