



Profiling and Optimizing Java Streams

Eduardo Rosales, Matteo Basso, Andrea Rosà, Walter Binder

Università della Svizzera italiana

<Programming> 2023
March 16, 2023



INTRODUCTION




- ***Stream API*** (package `java.util.stream`)
 - Data processing
 - MapReduce-style transformations
 - Two key abstractions:
 - ***Stream***
 - ***Stream pipeline***



```
transactionList.stream()  
  .parallel()  
  .filter(t -> t.getStatus() == Transaction.VALID)  
  .map(Transaction::getID)  
  .collect(Collectors.toSet());
```

```
transactionList.stream()
    .parallel()
    .filter(t -> t.getStatus() == Transaction.VALID)
    .map(Transaction::getID)
    .collect(Collectors.toSet());
```




- A ***stream*** represents a sequence of data elements coming from a ***data source***

```
transactionList.stream()
```


```
.parallel()
.filter(t -> t.getStatus() == Transaction.VALID)
.map(Transaction::getID)
.collect(Collectors.toSet());
```

Data source




```

transactionList.stream() → Stream
    .parallel()              creation
    .filter(t -> t.getStatus() == Transaction.VALID)
    .map(Transaction::getID)
    .collect(Collectors.toSet());
  
```



- A *pipeline* is associated with the stream

```
transactionList.stream()
    .parallel()
    .filter(t -> t.getStatus() == Transaction.VALID)
    .map(Transaction::getID)
    .collect(Collectors.toSet());
```



- The pipeline can contain **operations**


```
transactionList.stream()
```


```
.parallel()
```

```
.filter(t -> t.getStatus() == Transaction.VALID)
```

```
.map(Transaction::getID)
```


```
.collect(Collectors.toSet());
```

Execution Mode	
sequential	parallel




```
transactionList.stream()
    .parallel()
    .filter(t -> t.getStatus() == Transaction.VALID)
    .map(Transaction::getID)
    .collect(Collectors.toSet());
```

Execution Mode	
sequential	parallel



- Parallel streams may spawn **tasks**
- Tasks are executed in a **fork/join pool** by threads called **workers**





- Stream code suffers from important performance penalties when compared to imperative code [1, 2, 3]
- These penalties are mainly caused by:
 - **Abstraction overheads**
 - Due to extra object allocations and reclamations
 - **Abundant virtual method calls**
 - Can prevent *Just-In-Time* (JIT) compiler optimizations

[1] Biboudis et al. *Clash of the Lambdas*. IC00OLPS'14.

[2] Kiselyov et al. *Stream Fusion, to Completeness*. ACM SIGPLAN Notices, 2017.

[3] Møller et al. *Eliminating Abstraction Overhead of Java Stream Pipelines Using Ahead-of-Time Program Optimization*. OOPSLA'20.

To mitigate stream-related overheads and optimize streams
developers need means to study the **runtime behavior** of streams



Current approaches to stream optimization:

🗨️ Mainly rely on static analysis techniques

🗨️ Overlook **runtime information** key to spot stream-related performance issues

🗨️ Suffer from important limitations to detect all streams used by a Java application

- **Research gap:**

There is a lack of dedicated tools able to dynamically analyze stream processing on the JVM to help developers locate streams that impair good performance





- Propose a technique enabling cycle-accurate stream profiling
 - Accurately measure the computations performed by a stream in terms of **reference cycles** (*cycles* for short)
- Analyze stream processing in *Renaissance* [4]
- Optimize stream-related performance issues in *Renaissance*
- Conduct an evaluation on accuracy and overhead

[4] Prokopec et al. *Renaissance: Benchmarking Suite for Parallel Applications on the JVM*. PLDI'19.



PROFILING AND OPTIMIZING JAVA STREAMS




- Our technique to profile stream executions is implemented in ***StreamProf***, the first dedicated stream profiler for the JVM
- Features:
 - Detects every form of stream execution
 - Shows the impact of stream processing on application performance
 - Generates accurate profiles that help developers detect problematic streams


StreamProf targets **stream executions**:

- **Sequential stream executions:**
 - Carried out by the current thread

- **Parallel stream executions:**
 - Typically involve the execution of **tasks**
 - Tasks can be executed by multiple **workers**




- ***Nested stream***: A stream whose execution is triggered by (and occurs during the execution of) another ***outer stream***



- Multiple nesting levels are possible


- **Location:** Fully qualified name of the caller of the method executing the stream

<i>ExampleClass.exampleMethod</i> 	
1290	transactionList.stream()
1291	.parallel()
1292	.filter(t -> t.getStatus() == Transaction.VALID)
1293	.map(Transaction::getID)
1294	.collect(Collectors.toSet());



We model stream execution around the concept of *span*


Span: The interval in which a stream is executed by a thread





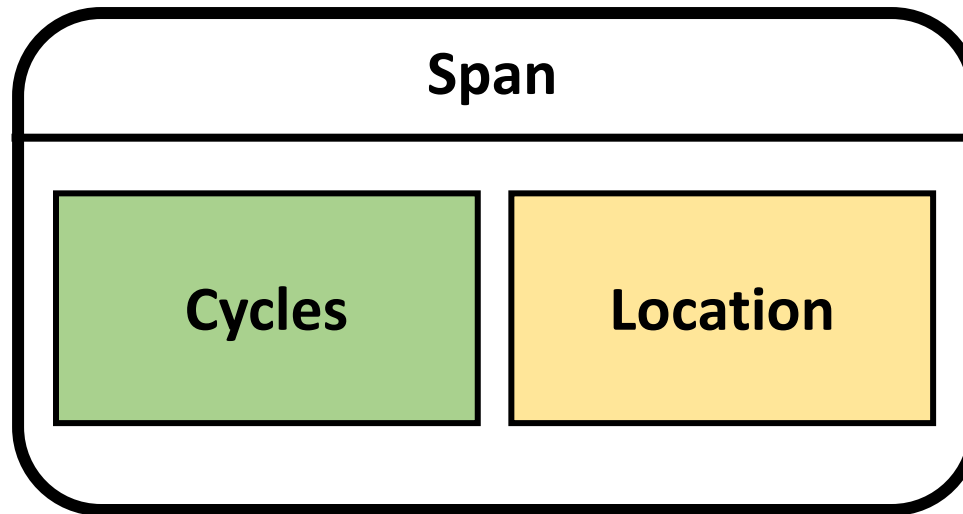


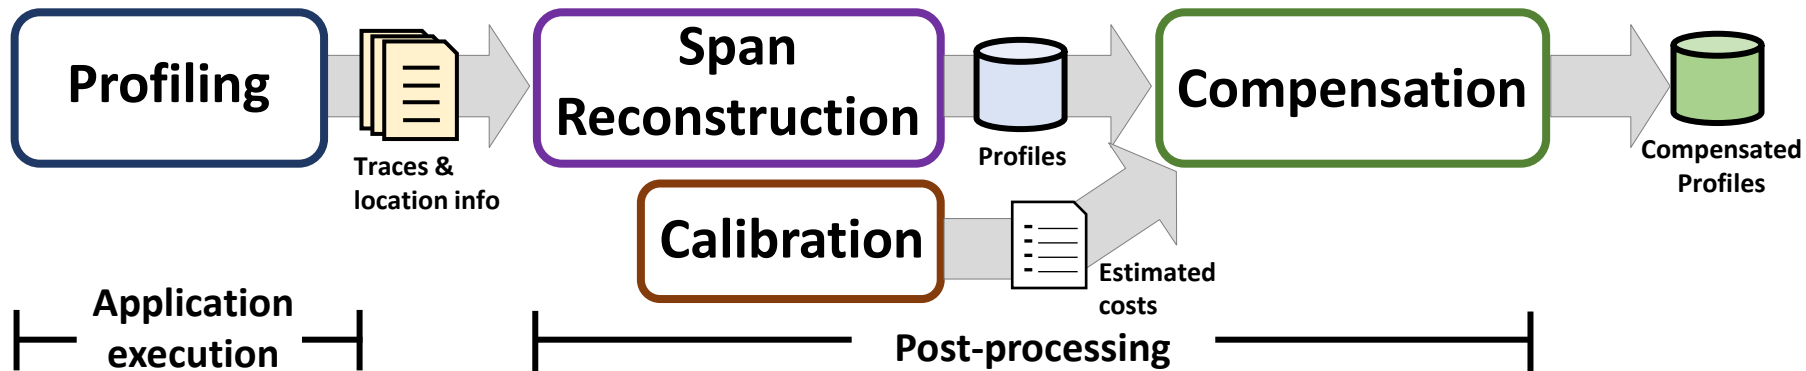


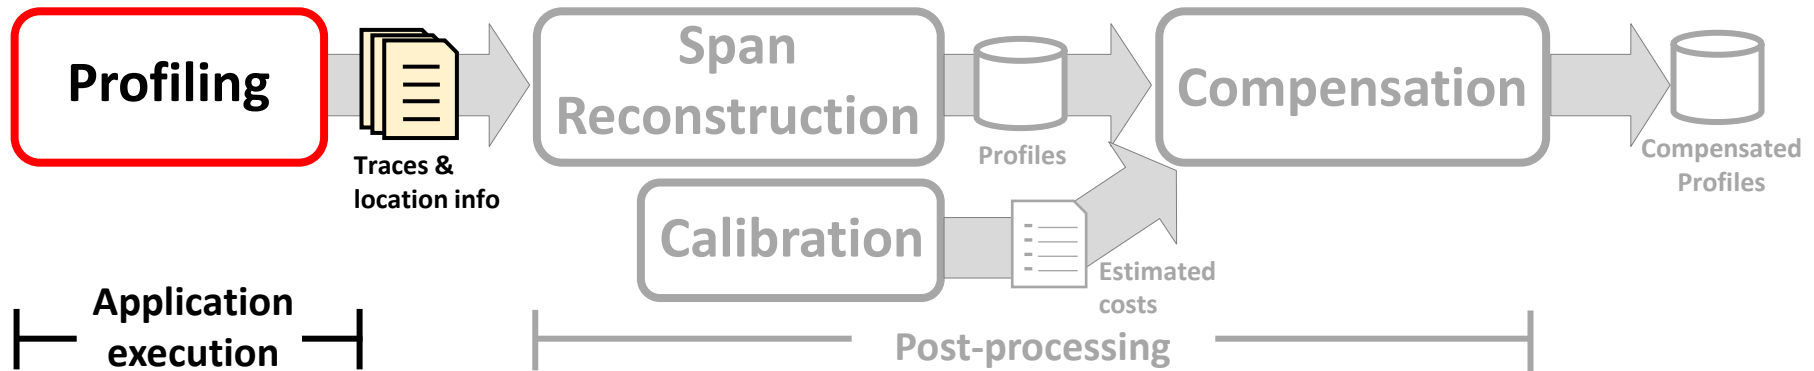


$$measured\ cycles = cycles_{end} - cycles_{begin}$$

- For each span we compute the corresponding cycles and location








- **Profiling:** Profile every span and associate it with the respective measured cycles and location

Difference when profiling sequential and parallel streams:

- **Sequential stream execution:**

- Carried out by a single thread

- Execution represented by a single span



- The measured cycles in the span can be directly attributed to the sequential stream execution

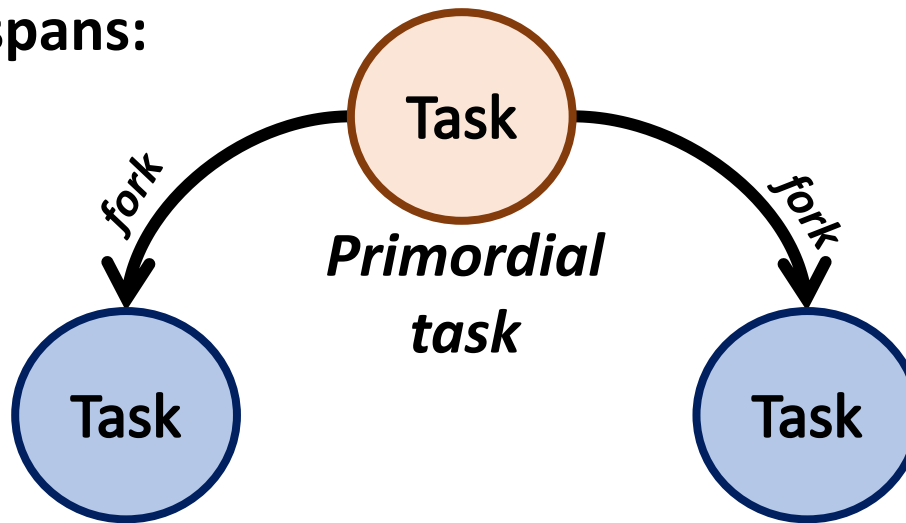
- **No stream ID required** to identify a sequential stream

- ***Anonymous span:***

- Represents a sequential stream execution

- **Parallel stream execution:**
 - Typically carried out by multiple tasks → multiple spans
 - To aggregate the measured cycles of all spans
 - A **stream ID required** to identify the parallel stream execution
 - **Named spans:** All spans associated with the same parallel stream execution


- Two named spans:



- Primordial span:** Represents the primordial task
 - Associated with a new unique stream ID
 - Support span:** Any named span other than the primordial span
 - Retrieves the already generated stream ID


The instrumentation relies on a *tracer*:

- Reads cycle counters and stores the values in the form of 2 events:



- span-begin* → *cycles_{begin}*


- span-end* → *cycles_{end}*




Trace

- The tracer maintains thread-local buffers
 - Stores span information associated to a thread in memory
 - No synchronization (only when allocating new buffers)

- The tracer dumps data upon application completion:
 - Dumps one trace per thread participating in stream processing




Trace - Thread 1



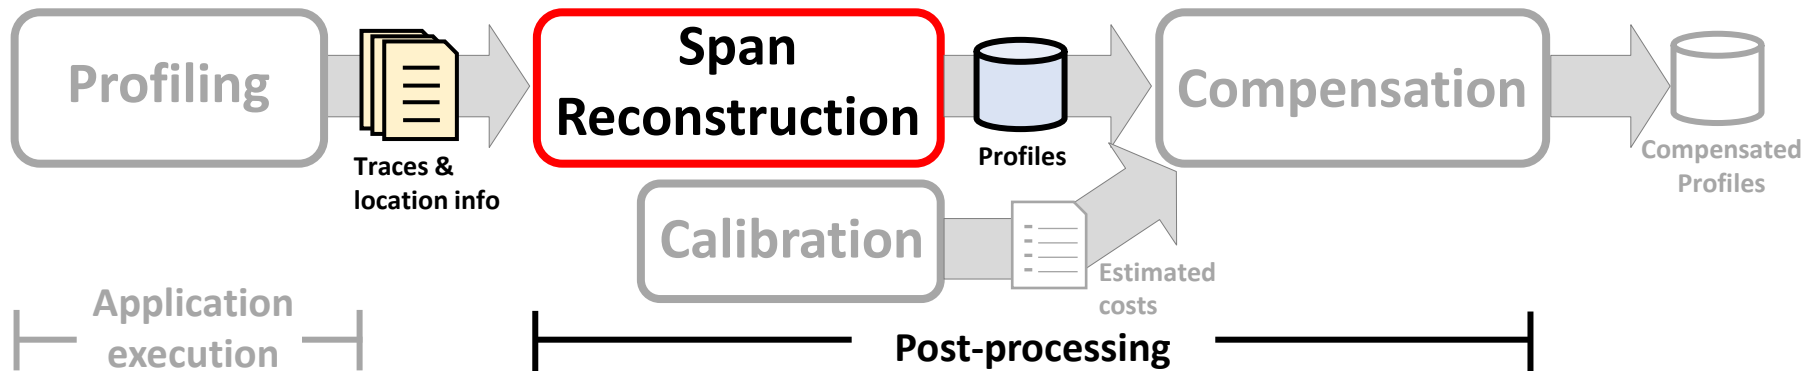
Trace - Thread 2

...




Trace - Thread n

- Traces are **post-processed offline**



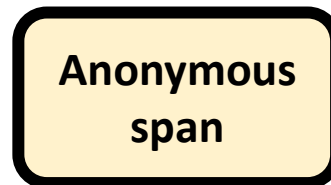
- **Span reconstruction:** each stream profile is reconstructed from the *span-begin* and *span-end* events stored in the dumped traces



Measured cycles computation:


- **Sequential streams:**

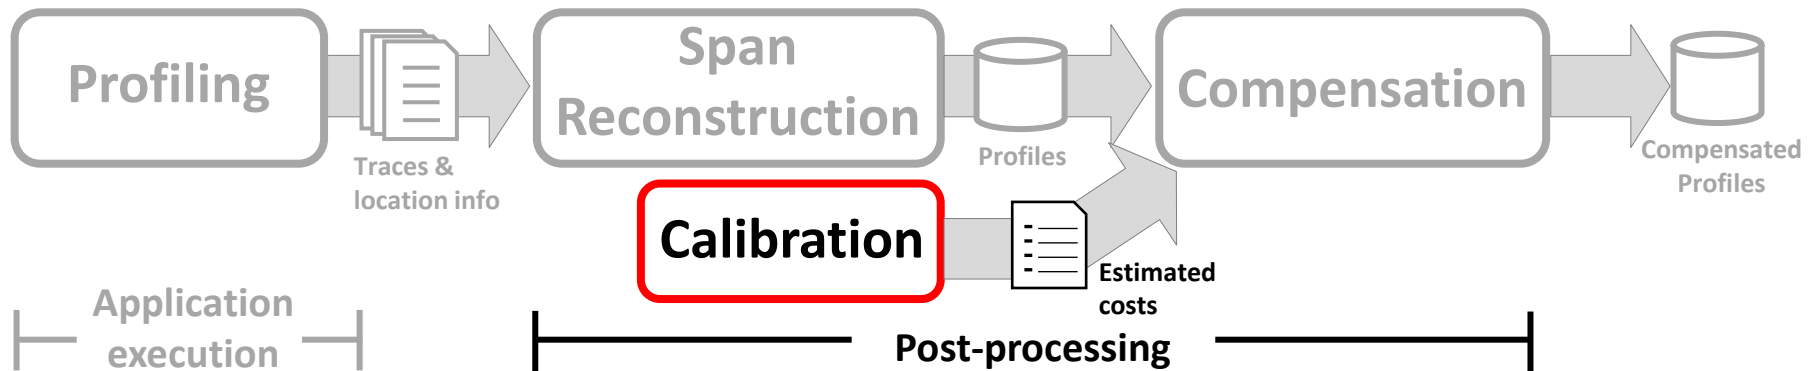
- Computed from the single anonymous span



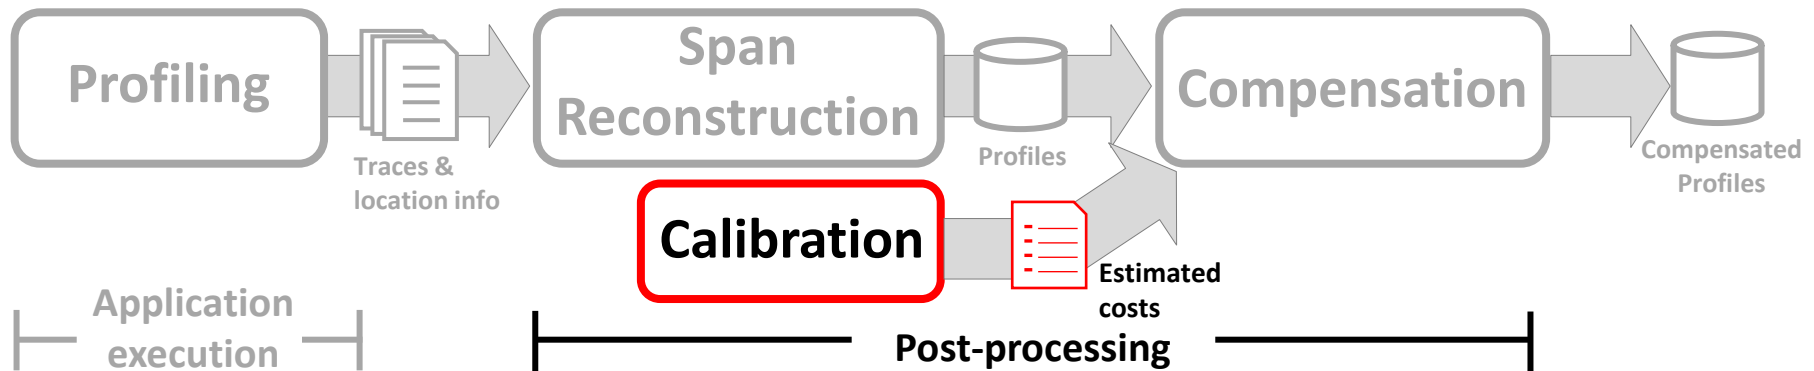
- **Parallel streams:**

- Computed from the respective multiple named spans

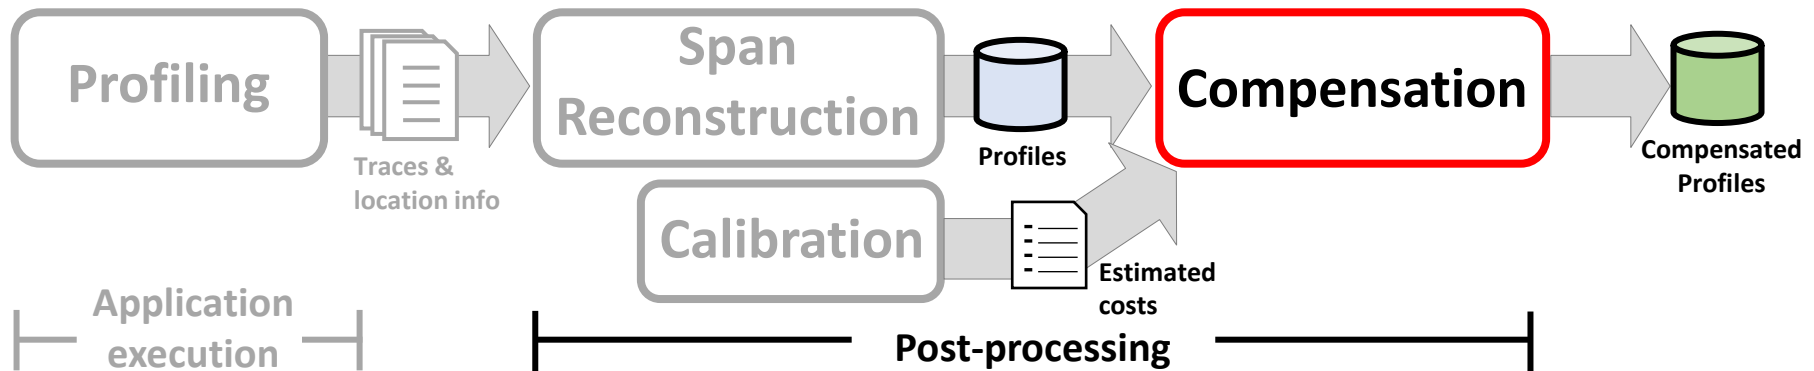




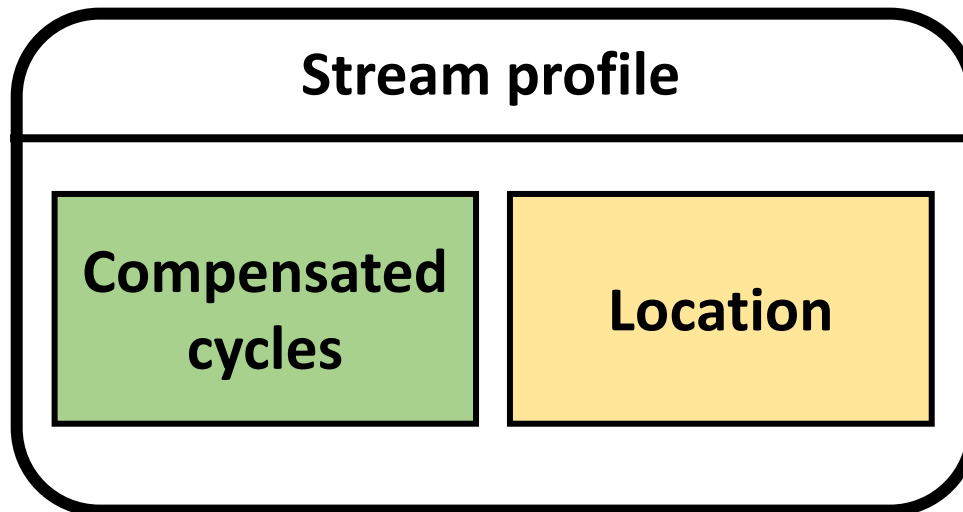
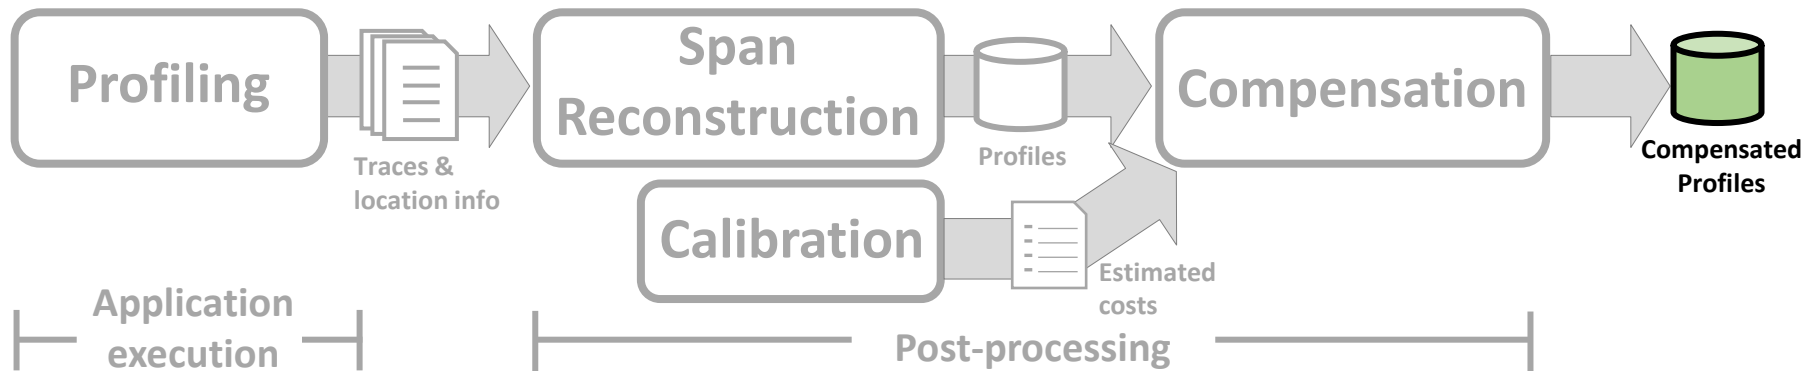
- The inserted instrumentation code introduces extra cycles that are included in the measured cycles of each span
 - Decreasing profile accuracy
- To remove these extra cycles they need to be estimated, which is the aim of the **calibration** phase



- **Calibration:** Produces *estimated costs*, i.e., estimations of the extra cycles required to profile spans
 - The estimated costs are approximated by constants



- **Compensation:** Produce *compensated cycles*, i.e., the measured cycles for a span after the removal of both:
 - Estimated costs (the output of the calibration)
 - *Nested cycles*, i.e., cycles elapsed due to nested spans






Using **StreamProf** to optimize stream processing:

- Target applications:
 - Stream-based workloads from *Renaissance* [8]
 - *mnemonics*
 - *par-mnemonics*
- Analysis targets **steady state**
- Testbed:
 - M₁: **8 core**, 128 GB RAM
 - M₂: **18 cores**, 256 GB RAM




OPTIMIZATION 1: REMOVING UNNEEDED STREAMS




- Customized heatmap of stream executions in *mnemonics*

Profiling and Optimizing Streams – Optimization I



x-axis: streams are grouped by their *compensated cycles*

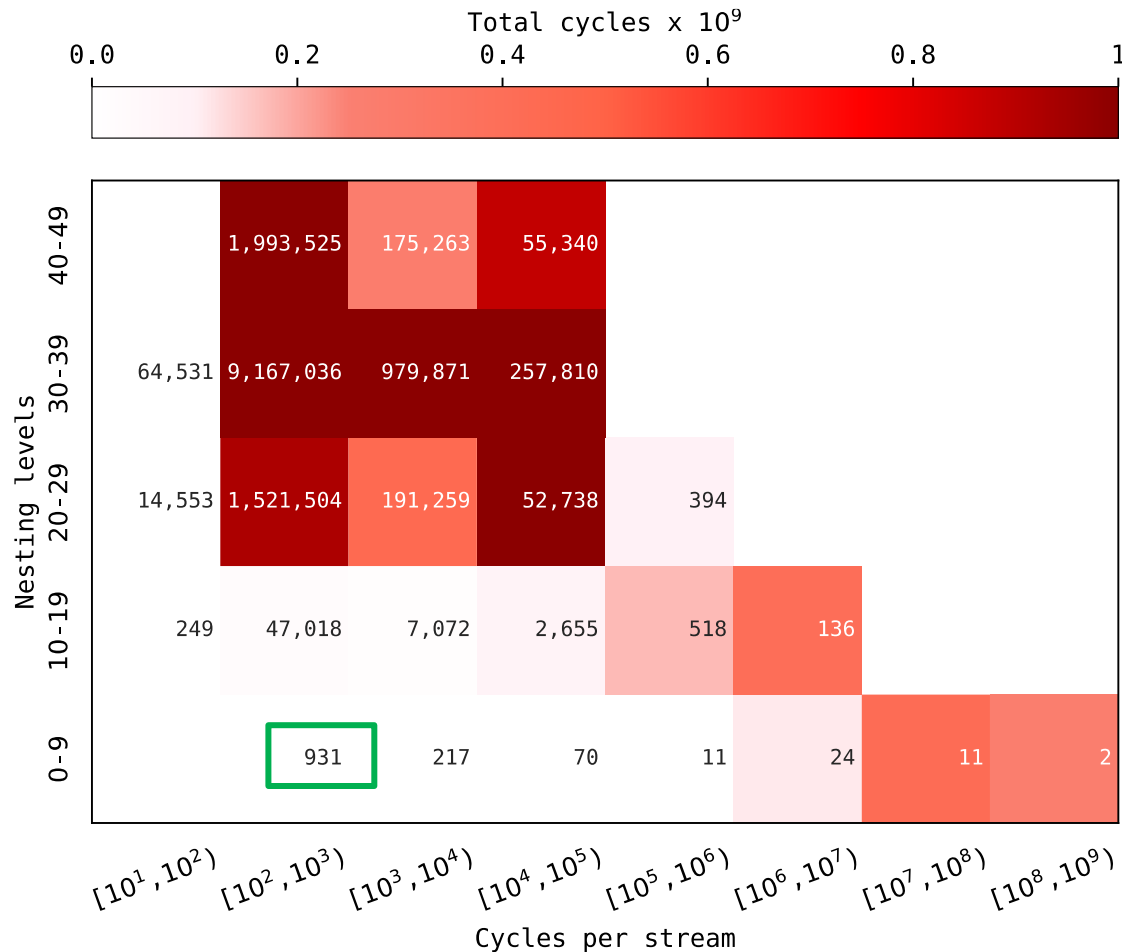
Profiling and Optimizing Streams – Optimization I



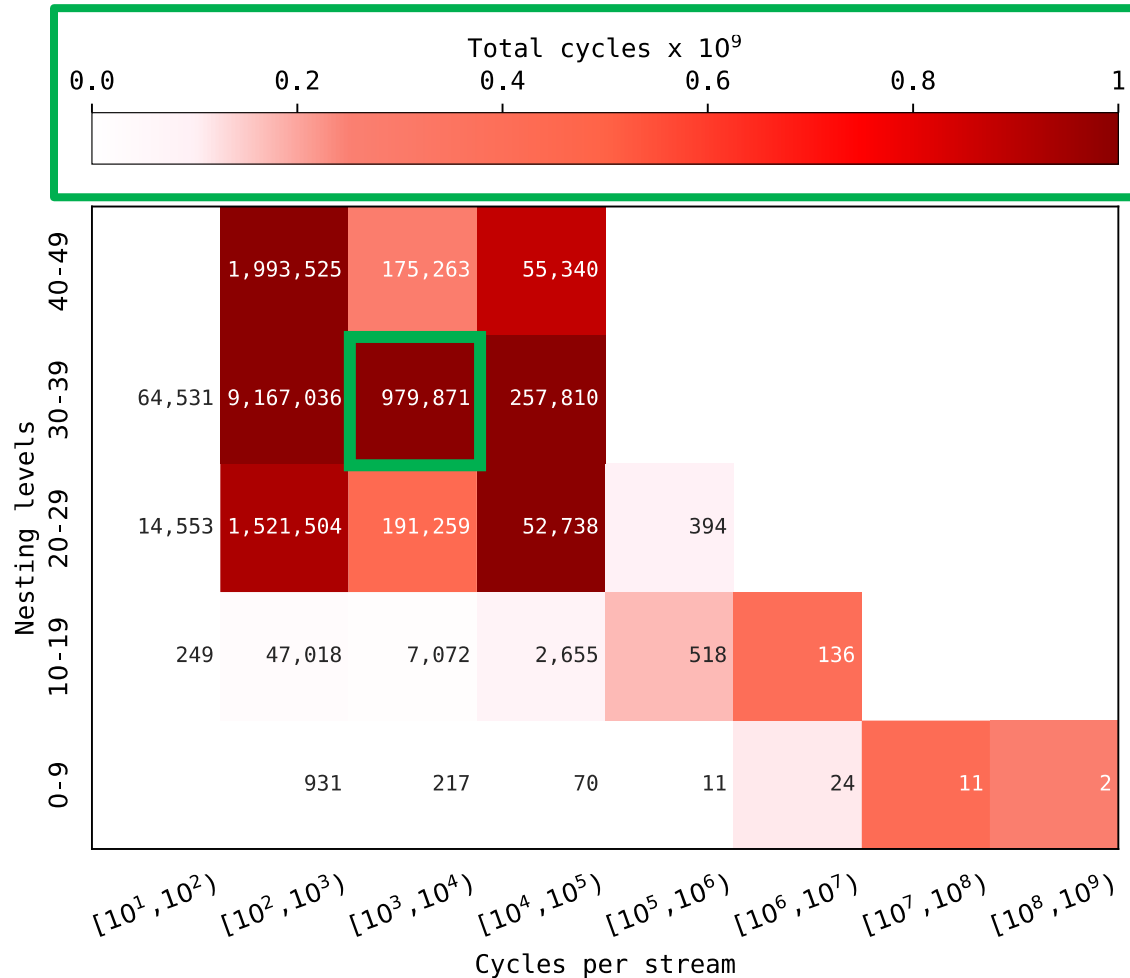
y-axis: streams are grouped by their *nesting level*



Profiling and Optimizing Streams – Optimization I




Cell: reports the number of stream executions for a given range of *nesting levels* and *cycles*

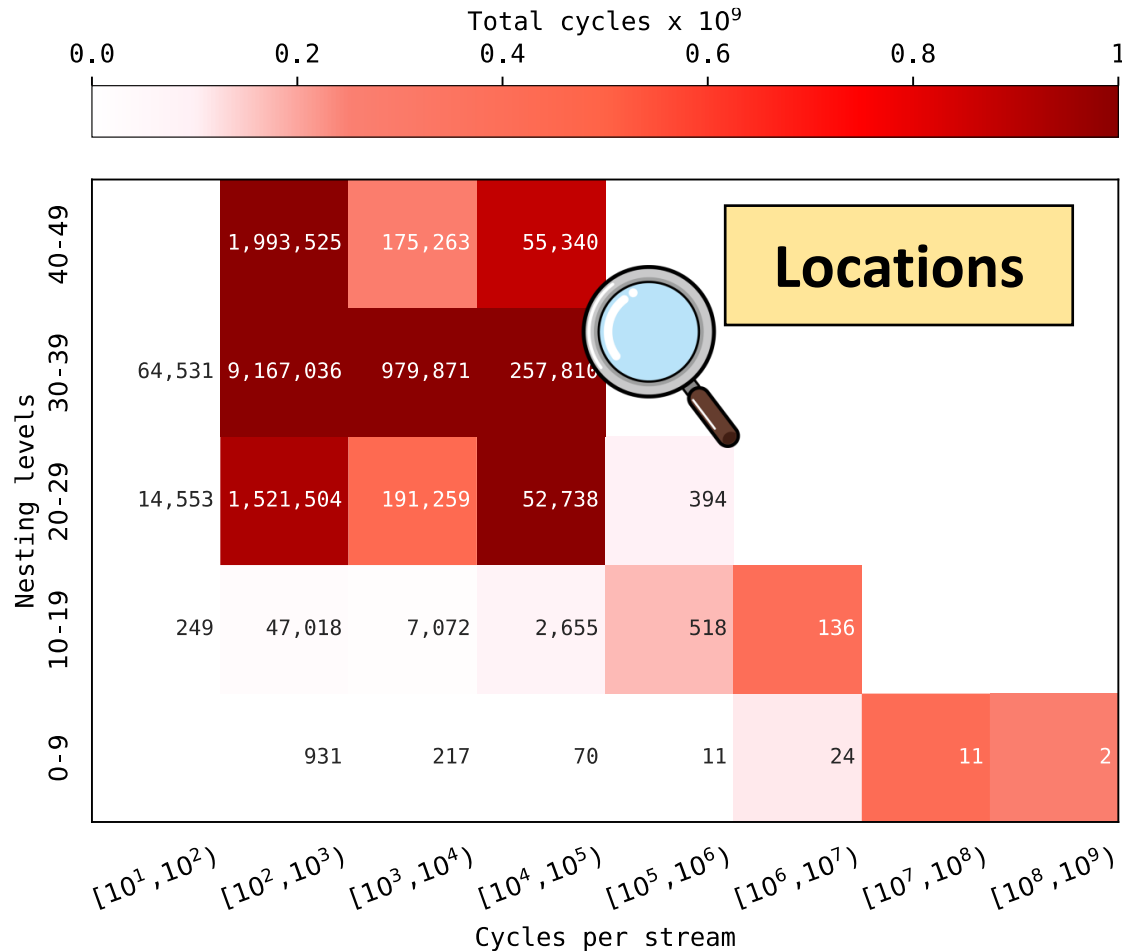


Heat: the color of a cell indicates the *total cycles of all the stream executions in the group*

Profiling and Optimizing Streams – Optimization I



Darker regions: good targets for stream code optimizations



Hot locations: locations in application code responsible for most of the stream processing



- **Optimization I:**

- **Target:** *mnemonics* and *par-mnemonics*

- *MnemonicsCoderWithStream*.**wordForNum**:

- Recursive method
- Many stream executions

-  **Finding:** These streams always produce the same result


- The result does not depend on any input of the current recursion level

- **Optimization:** Move the streams out the recursive code



-  **Benefit:** Reduce overheads due to unneeded stream executions



- *MnemonicsCoderWithStream.lambdas\$encode\$9*:
 - Many stream executions
 - ⚠ **Finding:** These streams are always created from an empty set
 - The empty set never changes
 - The streams do not contribute in any form to the workload output
- **Optimization:** Remove the unneeded streams
 - ✓ **Benefit:** Reduce overheads due to unneeded stream executions



- Heatmaps of the original (*left*) and optimized (*right*) mnemonics



- Heatmaps of the original (*left*) and optimized (*right*) mnemonics



OPTIMIZATION 2: IMPROVING LOAD IMBALANCE

- **Optimization II:**

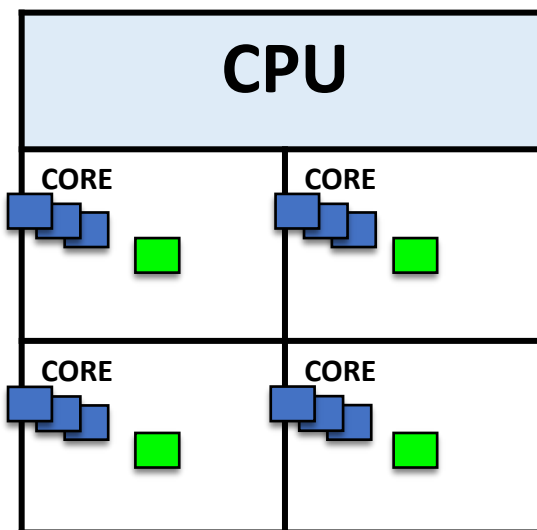

- **Target:** *par-mnemonics*

- *StreamProf* reports the distribution of cycles per worker

- ⚠ **Finding:** Only 2 workers execute more than **99.99%** of the total cycles (on both machines)

- **Optimization:** Tune *task granularity* to improve load balance

- ✓ **Benefit:** Improve parallel stream execution performance



Benchmark	Version	Machine	Time [ms]	Speedup	
				Factor	95% CI
<i>mnemonics</i>	orig	M ₁	4,944.72		
		M ₂	3,303.64		
	opt 1	M ₁	1,200.70	4.09	(3.99, 4.19)
		M ₂	981.97	3.36	(3.32, 3.40)
<i>par-mnemonics</i>	orig	M ₁	4,419.53		
		M ₂	2,977.83		
	opt 1	M ₁	1,106.22	3.98	(3.89, 4.04)
		M ₂	905.19	3.27	(3.20, 3.33)
	opt 2	M ₁	880.09	5.00	(4.91, 5.10)
		M ₂	764.63	3.88	(3.80, 3.96)

Benchmark	Version	Machine	Time [ms]	Speedup	
				Factor	95% CI
<i>mnemonics</i>	orig	M ₁	4,944.72		
		M ₂	3,303.64		
	opt 1	M ₁	1,200.70	4.09	(3.99, 4.19)
		M ₂	981.97	3.36	(3.32, 3.40)
<i>par-mnemonics</i>	orig	M ₁	4,419.53		
		M ₂	2,977.83		
	opt 1	M ₁	1,106.22	3.98	(3.89, 4.04)
		M ₂	905.19	3.27	(3.20, 3.33)
	opt 2	M ₁	880.09	5.00	(4.91, 5.10)
		M ₂	764.63	3.88	(3.80, 3.96)

Benchmark	Version	Machine	Time [ms]	Speedup	
				Factor	95% CI
<i>mnemonics</i>	orig	M ₁	4,944.72		
		M ₂	3,303.64		
	opt 1	M ₁	1,200.70	4.09	(3.99, 4.19)
		M ₂	981.97	3.36	(3.32, 3.40)
<i>par-mnemonics</i>	orig	M ₁	4,419.53		
		M ₂	2,977.83		
	opt 1	M ₁	1,106.22	3.98	(3.89, 4.04)
		M ₂	905.19	3.27	(3.20, 3.33)
	opt 2	M ₁	880.09	5.00	(4.91, 5.10)
		M ₂	764.63	3.88	(3.80, 3.96)

Benchmark	Version	Machine	Time [ms]	Speedup		
				Factor	95% CI	
<i>mnemonics</i>	orig	M ₁	4,944.72			
		M ₂	3,303.64			
	opt 1	M ₁	1,200.70	4.09	(3.99, 4.19)	
		M ₂	981.97	3.36	(3.32, 3.40)	
	<i>par-mnemonics</i>	orig	M ₁	4,419.53		
			M ₂	2,977.83		
opt 1		M ₁	1,106.22	3.98	(3.89, 4.04)	
		M ₂	905.19	3.27	(3.20, 3.33)	
opt 2	M ₁	880.09	5.00	(4.91, 5.10)		
	M ₂	764.63	3.88	(3.80, 3.96)		

Benchmark	Version	Machine	Time [ms]	Speedup	
				Factor	95% CI
<i>mnemonics</i>	orig	M ₁	4,944.72		
		M ₂	3,303.64		
	opt 1	M ₁	1,200.70	4.09	(3.99, 4.19)
		M ₂	981.97	3.36	(3.32, 3.40)
<i>par-mnemonics</i>	orig	M ₁	4,419.53		
		M ₂	2,977.83		
	opt 1	M ₁	1,106.22	3.98	(3.89, 4.04)
		M ₂	905.19	3.27	(3.20, 3.33)
	opt 2	M ₁	880.09	5.00	(4.91, 5.10)
		M ₂	764.63	3.88	(3.80, 3.96)

Benchmark	Version	Machine	Time [ms]	Speedup	
				Factor	95% CI
<i>mnemonics</i>	orig	M ₁	4,944.72		
		M ₂	3,303.64		
	opt 1	M ₁	1,200.70	4.09	(3.99, 4.19)
		M ₂	981.97	3.36	(3.32, 3.40)
<i>par-mnemonics</i>	orig	M ₁	4,419.53		
		M ₂	2,977.83		
	opt 1	M ₁	1,106.22	3.98	(3.89, 4.04)
		M ₂	905.19	3.27	(3.20, 3.33)
	opt 2	M ₁	880.09	5.00	(4.91, 5.10)
		M ₂	764.63	3.88	(3.80, 3.96)

- $$speedup = \frac{exec\ time\ original\ workload}{exec\ time\ optimized\ workload}$$

Benchmark	Version	Machine	Time [ms]	Speedup	
				Factor	95% CI
<i>mnemonics</i>	orig	M ₁	4,944.72		
		M ₂	3,303.64		
	opt 1	M ₁	1,200.70	4.09	(3.99, 4.19)
		M ₂	981.97	3.36	(3.32, 3.40)
<i>par-mnemonics</i>	orig	M ₁	4,419.53		
		M ₂	2,977.83		
	opt 1	M ₁	1,106.22	3.98	(3.89, 4.04)
		M ₂	905.19	3.27	(3.20, 3.33)
	opt 2	M ₁	880.09	5.00	(4.91, 5.10)
		M ₂	764.63	3.88	(3.80, 3.96)

- 95% confidence intervals (*CI*)

Benchmark	Version	Machine	Time [ms]	Speedup	
				Factor	95% CI
<i>mnemonics</i>	orig	M ₁	4,944.72		
		M ₂	3,303.64		
	opt 1	M ₁	1,200.70	4.09	(3.99, 4.19)
		M ₂	981.97	3.36	(3.32, 3.40)
<i>par-mnemonics</i>	orig	M ₁	4,419.53		
		M ₂	2,977.83		
	opt 1	M ₁	1,106.22	3.98	(3.89, 4.04)
		M ₂	905.19	3.27	(3.20, 3.33)
	opt 2	M ₁	880.09	5.00	(4.91, 5.10)
		M ₂	764.63	3.88	(3.80, 3.96)

- ✓ Removing unneeded stream processing (**opt1**) improves the performance of both *mnemonics* and *par-mnemonics*
- ✓ Only 6 lines of code changed

Benchmark	Version	Machine	Time [ms]	Speedup	
				Factor	95% CI
<i>mnemonics</i>	orig	M ₁	4,944.72		
		M ₂	3,303.64		
	opt 1	M ₁	1,200.70	4.09	(3.99, 4.19)
		M ₂	981.97	3.36	(3.32, 3.40)
<i>par-mnemonics</i>	orig	M ₁	4,419.53		
		M ₂	2,977.83		
	opt 1	M ₁	1,106.22	3.98	(3.89, 4.04)
		M ₂	905.19	3.27	(3.20, 3.33)
	opt 2	M ₁	880.09	5.00	(4.91, 5.10)
		M ₂	764.63	3.88	(3.80, 3.96)

✓ Improving load balance (**opt2**) results in performance gains in *par-mnemonics*

✓ 2 lines of code changed

Benchmark	Version	Machine	Time [ms]	Speedup	
				Factor	95% CI
<i>mnemonics</i>	orig	M ₁	4,944.72		
		M ₂	3,303.64		
	opt 1	M ₁	1,200.70	4.09	(3.99, 4.19)
		M ₂	981.97	3.36	(3.32, 3.40)
<i>par-mnemonics</i>	orig	M ₁	4,419.53		
		M ₂	2,977.83		
	opt 1	M ₁	1,106.22	3.98	(3.89, 4.04)
		M ₂	905.19	3.27	(3.20, 3.33)
	opt 2	M ₁	880.09	5.00	(4.91, 5.10)
		M ₂	764.63	3.88	(3.80, 3.96)

✓ Average speedup is **3.94x** considering all workloads on both machines





ACCURACY AND OVERHEAD EVALUATION

- **Average *Cycles Per Span* (CPS)**

- **Expectation:**

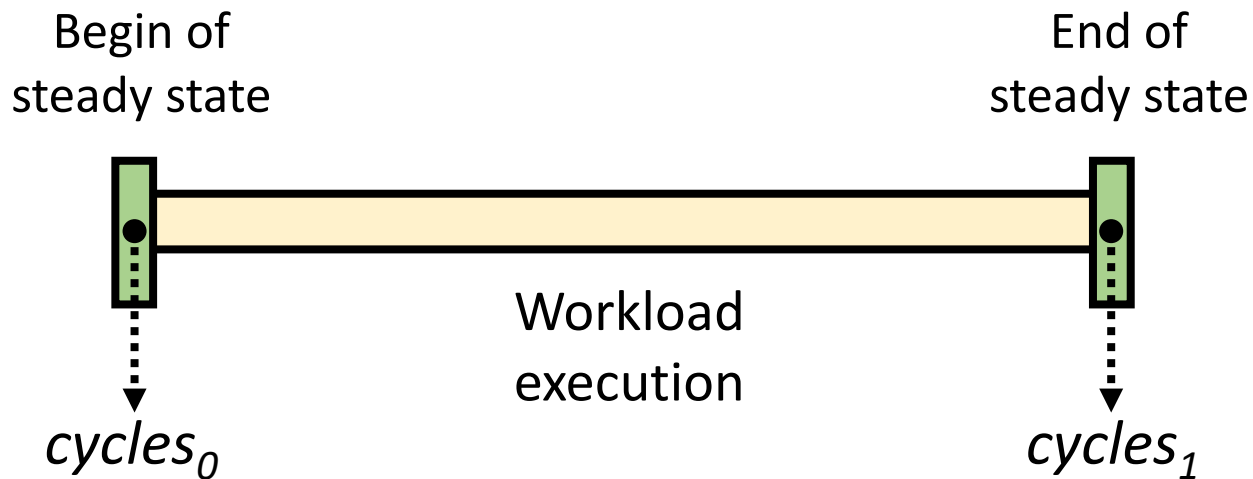
- For workloads with a low CPS

- the relative cost of the inserted instrumentation code
 - is higher than for workloads with a higher CPS



■ Baseline computation:

- To compute CPS and accuracy, we need a *baseline*
- All evaluated workloads take place only within stream executions
- Baseline: total cycles elapsed by all threads involved in stream processing



$$\mathbf{baseline = cycles_1 - cycles_0}$$

- **CPS computation:**

$$CPS = \textit{baseline} / \textit{total_spans}$$

- **Accuracy computation:**

$$RE = \left| \frac{\textit{baseline} - \textit{compensated_cycles}}{\textit{baseline}} \right|$$

$$\text{Accuracy} = 1 - RE$$



- Testbed:
 - All stream-based workloads from *Renaissance*
 - *mnemonics*
 - *par-mnemonics*
 - *scrabble*
 - OpenJDK [5]
 - Collection of stream-based workloads released by the developers of OpenJDK
 - JEDI [6]
 - Benchmark suite consisting of TPCH-queries recast to Java Streams

[5] <https://github.com/usi-dag/jdk20u/tree/master/test/micro/org/openjdk/bench/java/util/stream>

[6] <https://github.com/usi-dag/JEDI>



ACCURACY EVALUATION



Benchmark	Version	Machine	CPS	Accuracy [%]	
<i>mnemonics</i>	orig	M ₁	915	87.99	
		M ₂	679	95.18	
	opt 1	M ₁	2,972	98.99	
		M ₂	2,699	99.54	
	<i>par-mnemonics</i>	orig	M ₁	1,027	88.27
			M ₂	761	95.83
opt 1		M ₁	3,290	99.68	
		M ₂	2,989	99.66	
opt 2		M ₁	3,241	99.42	
		M ₂	3,436	94.98	
<i>scrabble</i>	orig	M ₁	1,481	89.97	
		M ₂	2,151	97.13	

Benchmark	Version	Machine	CPS	Accuracy [%]	
<i>mnemonics</i>	orig	M ₁	915	87.99	
		M ₂	679	95.18	
	opt 1	M ₁	2,972	98.99	
		M ₂	2,699	99.54	
	<i>par-mnemonics</i>	orig	M ₁	1,027	88.27
			M ₂	761	95.83
opt 1		M ₁	3,290	99.68	
		M ₂	2,989	99.66	
opt 2		M ₁	3,241	99.42	
		M ₂	3,436	94.98	
<i>scrabble</i>	orig	M ₁	1,481	89.97	
		M ₂	2,151	97.13	

Benchmark	Version	Machine	CPS	Accuracy [%]
<i>mnemonics</i>	orig	M ₁	915	87.99
		M ₂	679	95.18
	opt 1	M ₁	2,972	98.99
		M ₂	2,699	99.54
<i>par-mnemonics</i>	orig	M ₁	1,027	88.27
		M ₂	761	95.83
	opt 1	M ₁	3,290	99.68
		M ₂	2,989	99.66
	opt 2	M ₁	3,241	99.42
		M ₂	3,436	94.98
<i>scrabble</i>	orig	M ₁	1,481	89.97
		M ₂	2,151	97.13

- Accuracy shown as percentage

Benchmark	Version	Machine	CPS	Accuracy [%]
<i>mnemonics</i>	orig	M ₁	915	87.99
		M ₂	679	95.18
	opt 1	M ₁	2,972	98.99
		M ₂	2,699	99.54
<i>par-mnemonics</i>	orig	M ₁	1,027	88.27
		M ₂	761	95.83
	opt 1	M ₁	3,290	99.68
		M ₂	2,989	99.66
	opt 2	M ₁	3,241	99.42
		M ₂	3,436	94.98
<i>scrabble</i>	orig	M ₁	1,481	89.97
		M ₂	2,151	97.13

- Positive correlation between CPS and accuracy
 - PCC: **0.71**, considering all workloads on M₁ and M₂

Benchmark	Version	Machine	CPS	Accuracy [%]	
<i>mnemonics</i>	orig	M ₁	915	87.99	
		M ₂	679	95.18	
	opt 1	M ₁	2,972	98.99	
		M ₂	2,699	99.54	
	<i>par-mnemonics</i>	orig	M ₁	1,027	88.27
			M ₂	761	95.83
opt 1		M ₁	3,290	99.68	
		M ₂	2,989	99.66	
opt 2		M ₁	3,241	99.42	
		M ₂	3,436	94.98	
<i>scrabble</i>	orig	M ₁	1,481	89.97	
		M ₂	2,151	97.13	
JEDI (mean)		M ₁		91.55	
		M ₂		90.34	
OpenJDK (mean)		M ₁		94.82	
		M ₂		96.73	



OVERHEAD EVALUATION



Benchmark	Version	Machine	Time [ms]	CPS	Overhead	
					Factor	95% CI
<i>mnemonics</i>	orig	M ₁	4,944.72	915	1.55	(1.51, 1.60)
		M ₂	3,303.64	679	1.46	(1.44, 1.49)
	opt 1	M ₁	1,200.70	2,972	1.16	(1.15, 1.16)
		M ₂	981.97	2,699	1.12	(1.12, 1.13)
<i>par-mnemonics</i>	orig	M ₁	4,419.53	1,027	1.51	(1.47, 1.54)
		M ₂	2,977.83	761	1.41	(1.38, 1.45)
	opt 1	M ₁	1,106.22	3,290	1.13	(1.13, 1.14)
		M ₂	905.19	2,989	1.11	(1.11, 1.12)
	opt 2	M ₁	880.09	3,241	1.10	(1.10, 1.11)
		M ₂	764.63	3,436	1.09	(1.08, 1.09)
<i>scrabble</i>	orig	M ₁	310.84	1,481	1.43	(1.42, 1.44)
		M ₂	187.09	2,151	1.23	(1.22, 1.23)

Overhead:

- $$\text{slowdown factor} = \frac{\text{exec time workload with profiling}}{\text{exec time workload without profiling}}$$

Benchmark	Version	Machine	Time [ms]	CPS	Overhead	
					Factor	95% CI
<i>mnemonics</i>	orig	M ₁	4,944.72	915	1.55	(1.51, 1.60)
		M ₂	3,303.64	679	1.46	(1.44, 1.49)
	opt 1	M ₁	1,200.70	2,972	1.16	(1.15, 1.16)
		M ₂	981.97	2,699	1.12	(1.12, 1.13)
<i>par-mnemonics</i>	orig	M ₁	4,419.53	1,027	1.51	(1.47, 1.54)
		M ₂	2,977.83	761	1.41	(1.38, 1.45)
	opt 1	M ₁	1,106.22	3,290	1.13	(1.13, 1.14)
		M ₂	905.19	2,989	1.11	(1.11, 1.12)
	opt 2	M ₁	880.09	3,241	1.10	(1.10, 1.11)
		M ₂	764.63	3,436	1.09	(1.08, 1.09)
<i>scrabble</i>	orig	M ₁	310.84	1,481	1.43	(1.42, 1.44)
		M ₂	187.09	2,151	1.23	(1.22, 1.23)

- Negative correlation between CPS and overhead
 - PCC: **-0.96**, considering all workloads on M₁ and M₂



Benchmark	Version	Machine	Time [ms]	CPS	Overhead	
					Factor	95% CI
<i>mnemonics</i>	orig	M ₁	4,944.72	915	1.55	(1.51, 1.60)
		M ₂	3,303.64	679	1.46	(1.44, 1.49)
	opt 1	M ₁	1,200.70	2,972	1.16	(1.15, 1.16)
		M ₂	981.97	2,699	1.12	(1.12, 1.13)
<i>par-mnemonics</i>	orig	M ₁	4,419.53	1,027	1.51	(1.47, 1.54)
		M ₂	2,977.83	761	1.41	(1.38, 1.45)
	opt 1	M ₁	1,106.22	3,290	1.13	(1.13, 1.14)
		M ₂	905.19	2,989	1.11	(1.11, 1.12)
	opt 2	M ₁	880.09	3,241	1.10	(1.10, 1.11)
		M ₂	764.63	3,436	1.09	(1.08, 1.09)
<i>scrabble</i>	orig	M ₁	310.84	1,481	1.43	(1.42, 1.44)
		M ₂	187.09	2,151	1.23	(1.22, 1.23)
JEDI (mean)		M ₁			1.13	
		M ₂			1.15	
OpenJDK (mean)		M ₁			1.07	
		M ₂			1.06	

- **Limitations of the technique:**
 - Additional sources of perturbation (e.g., prevention of JIT compiler optimizations) are not compensated
 - The profiles produced using our technique are platform dependent and require the availability of per-thread virtualized reference-cycle counters



- New technique enabling cycle-accurate profiling of stream executions
 - Implemented in *StreamProf*, a novel stream profiler for the JVM
 - Uses a perturbation-compensation technique
- Analysis of *Renaissance*, revealing previously unknown stream-related performance issues
- Optimization of two stream-based workloads from *Renaissance*
 - Average speedup: 3.94x
- Evaluation results show that our profiling technique is efficient and yields accurate profiles



Thanks a lot for your attention!

