



**Better a Microbenchmark on a Cluster than a  
User at the Office:  
Flink Cluster Benchmarking**

David Georg Reichelt <sup>1</sup>    Lars-Peter Meyer <sup>1</sup>  
Stefan Kühne <sup>1</sup>

<sup>1</sup>Universität Leipzig, University Computing Centre,  
Research and Development



- 1 Apache Flink
- 2 Microbenchmark Suite
- 3 Case Study
- 4 Related Work
- 5 Summary





# Apache Flink





...



...

```
1  DataStream<WordWithCount> windowCounts = text
2    .flatMap(new FlatMapFunction<String,WordWithCount>()
3    {
4    @Override
5    public void flatMap(String value,
6                        Collector<WordWithCount> out) {
7        for (String word : value.split("\\s")) {
8            out.collect(new WordWithCount(word, 1L));
9        }
10   } } )
11   .keyBy("word")
12   .timeWindow(Time.seconds(5))
13   .reduce(new ReduceFunction<WordWithCount>() {
14   @Override
15   public WordWithCount reduce(WordWithCount a,
16                               WordWithCount b) {
17       return new WordWithCount(a.word,
18                               a.count + b.count);
19   } } );
```

```
1  DataStream<WordWithCount> windowCounts = text  
2      .flatMap(new FlatMapFunction<String, WordWithCount>()  
3      {  
4      @Override  
5      public void flatMap(String value,  
6                          Collector<WordWithCount> out) {  
7          for (String word : value.split("\\s")) {  
8              out.collect(new WordWithCount(word, 1L));  
9          }  
10     } } )  
11     .keyBy("word")  
12     .timeWindow(Time.seconds(5))  
13     .reduce(new ReduceFunction<WordWithCount>() {  
14     @Override  
15     public WordWithCount reduce(WordWithCount a,  
16                                 WordWithCount b) {  
17         return new WordWithCount(a.word,  
18                                 a.count + b.count);  
19     } } );
```

```
1  DataStream<WordWithCount> windowCounts = text
2  .flatMap(new FlatMapFunction<String, WordWithCount>()
3  {
4  @Override
5  public void flatMap(String value,
6  Collector<WordWithCount> out) {
7  for (String word : value.split("\\s")) {
8  out.collect(new WordWithCount(word, 1L));
9  }
10 } } )
11 .keyBy("word")
12 .timeWindow(Time.seconds(5))
13 .reduce(new ReduceFunction<WordWithCount>() {
14 @Override
15 public WordWithCount reduce(WordWithCount a,
16 WordWithCount b) {
17 return new WordWithCount(a.word,
18 a.count + b.count);
19 } } );
```

```
1  DataStream<WordWithCount> windowCounts = text
2    .flatMap(new FlatMapFunction<String, WordWithCount>()
3    {
4    @Override
5    public void flatMap(String value,
6                        Collector<WordWithCount> out) {
7        for (String word : value.split("\\s")) {
8            out.collect(new WordWithCount(word, 1L));
9        }
10   } } )
11   .keyBy("word")
12   .timeWindow(Time.seconds(5))
13   .reduce(new ReduceFunction<WordWithCount>() {
14   @Override
15   public WordWithCount reduce(WordWithCount a,
16                               WordWithCount b) {
17       return new WordWithCount(a.word,
18                               a.count + b.count);
19   } } );
```



```
1  DataStream<WordWithCount> windowCounts = text
2    .flatMap(new FlatMapFunction<String, WordWithCount>()
3    {
4    @Override
5    public void flatMap(String value,
6                        Collector<WordWithCount> out) {
7        for (String word : value.split("\\s")) {
8            out.collect(new WordWithCount(word, 1L));
9        }
10   } } )
11   .keyBy("word")
12   .timeWindow(Time.seconds(5))
13   .reduce(new ReduceFunction<WordWithCount>() {
14   @Override
15   public WordWithCount reduce(WordWithCount a,
16                               WordWithCount b) {
17       return new WordWithCount(a.word,
18                               a.count + b.count);
19   } } );
```

```
1  DataStream<WordWithCount> windowCounts = text
2    .flatMap(new FlatMapFunction<String, WordWithCount>()
3    {
4    @Override
5    public void flatMap(String value,
6                        Collector<WordWithCount> out) {
7        for (String word : value.split("\\s")) {
8            out.collect(new WordWithCount(word, 1L));
9        }
10   } } )
11   .keyBy("word")
12   .timeWindow(Time.seconds(5))
13   .reduce(new ReduceFunction<WordWithCount>() {
14   @Override
15   public WordWithCount reduce(WordWithCount a,
16                               WordWithCount b) {
17       return new WordWithCount(a.word,
18                               a.count + b.count);
19   } } );
```

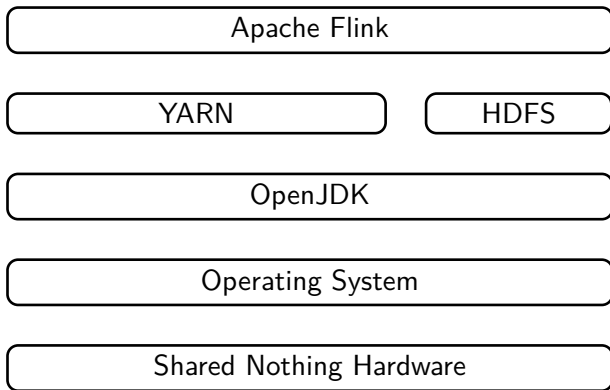


Figure: Tools and our Cluster Setup of Flink





# Microbenchmark Suite





- Goal
  - Identify performance problems in Flink clusters
  - Usage from operators perspective
  - Short-running benchmarks





- Goal
  - Identify performance problems in Flink clusters
  - Usage from operators perspective
  - Short-running benchmarks
- Makrobenchmark: Measure the performance of a whole application
- Micro-Benchmark: Measure the performance of a small code unit



- Each workload: Job executed on cluster
- Parameters: *size* and *parallelism*
- Computing performance: Execute *size \* factor* additions
- RAM performance: Reserve *size \* factor* bytes
- HDFS write performance
  - Write *size \* factor* bytes
  - return throughput
- HDFS read performance
  - Read *size \* factor* bytes from previously created files
  - return throughput

- Steady state needs to be assured
- Assumption: Normal distribution  
⇒ Steady state reached if t-test does not reject that values are equal
- Setup: 100 Jobs with 10.000, 20.000, 40.000 size
- Measurements 15 to 25 not unequal 5 to 15  
⇒ 25 Measurements, use last 10



- Steady state needs to be assured




Figure: Average Execution Times (Size 10.000)



# Case Study



- 2 Clusters á 16 Workers
- Both Flink 1.5 with YARN, OpenJDK 1.8

- 2 Clusters á 16 Workers
- Both Flink 1.5 with YARN, OpenJDK 1.8

	Worker@Cluster1	Worker@Cluster2
CPU	2 Xeon <sup>®</sup> , 2.4GHz	1 Xeon <sup>®</sup> , 2.5GHz
Model	E5-2620v3	E5-2430v2
RAM	128 GByte	48 GByte
HDDs	5x4TB SATA	2x4TB SATA
Ethernet	10 GBit/s	1 GBit/s
OS	CentOS 7.6	openSUSE 13.2

**Table:** Hardware of the Clusters Used




Figure: Add-Benchmark Execution Time




Figure: Histogram of Parallel Execution Durations




Figure: Write-Benchmark Execution Time




Figure: Histogram of Execution Durations





## Related Work





- Benchmarking from users perspective
  - HiBench suite
  - BigBench
  - Production scenarios
- Benchmarking HDFS
  - TestDFSIO
  - Enhanced TestDFSIO





# Summary





- Built benchmark from operators perspective
- Goal: Reproduce problems with microbenchmarks
- Case study: Problems reproducible
  
- Next step: Facilitate root cause analysis
  - Run on standalone cluster
  - Search logs
  - Add workloads





# Thanks for your attention!

David Georg Reichelt  
Universitätsrechenzentrum  
Universität Leipzig  
reichelt@informatik.uni-leipzig.de

Funded by:



Hanns  
Seidel  
Stiftung



FKZ 01IS14014B

SPONSORED BY THE



Federal Ministry  
of Education  
and Research