## **Java 8 Parallel Stream Internals**

(Part 5)

Douglas C. Schmidt <u>d.schmidt@vanderbilt.edu</u> www.dre.vanderbilt.edu/~schmidt



**Professor of Computer Science** 

Institute for Software Integrated Systems

Vanderbilt University Nashville, Tennessee, USA



#### Learning Objectives in this Part of the Lesson

- Understand parallel stream internals, e.g.
  - Know what can change & what can't
  - Partition a data source into "chunks"
  - Process chunks in parallel
  - Configure the Java 8 parallel stream common fork-join pool
  - Avoid pool starvation & improve performance w/ManagedBlocker
  - Perform a reduction that combines partial results into a single result



See <a href="https://www.ibm.com/developerworks/library/j-java-streams-3-brian-goetz">www.ibm.com/developerworks/library/j-java-streams-3-brian-goetz</a>



This discussion assumes a non-concurrent collector (more discussions follow)

• After the common fork-join pool finishes DataSource processing chunks their partial results are combined into a final result DataSource<sub>1</sub> DataSource<sub>2</sub> join() occurs in a single thread at each level DataSource<sub>1,2</sub> DataSource<sub>2,2</sub> DataSource<sub>1 1</sub> DataSource<sub>21</sub> • i.e., the "parent" Process Process Process Process sequentially sequentially sequentially sequentially "Children" join "Parent"

- After the common fork-join pool finishes processing chunks their partial results are combined into a final result
  - join() occurs in a single thread at each level
    - i.e., the "parent"



As a result, there's typically no need for synchronizers during the joining

• Different terminal operations combine partial results in different ways



Understanding these differences is particularly important for parallel streams

- Different terminal operations combine partial results in different ways, e.g.
  - reduce() creates a new immutable value



#### See <a href="https://docs.oracle.com/javase/tutorial/essential/concurrency/immutable.html">docs.oracle.com/javase/tutorial/essential/concurrency/immutable.html</a>



See github.com/douglascraigschmidt/LiveLessons/tree/master/Java8/ex16



reduce() combines two immutable values (e.g., long or Long) & produces a new one

- Different terminal operations combine partial results in different ways, e.g.
  - reduce() creates a new immutable value
  - collect() mutates an existing value



See greenteapress.com/thinkapjava/html/thinkjava011.html

- Different terminal operations combine All words in Shakespeare's works partial results in different ways, e.g. reduce() creates a new 1<sup>st</sup> half of words 2nd half of words immutable value collect() mutates an 2nd quarter of words 4<sup>th</sup> quarter of words 1<sup>st</sup> quarter of words 3rd quarter of words existing value Process Process Process Process sequentially sequentially sequentially sequentially List<CharSequence> uniqueWords = getInput(sSHAKESPEARE), "\\s+") .parallelStream()
  - • •

.collect(toCollection(TreeSet::new));

See github.com/douglascraigschmidt/LiveLessons/tree/master/Java8/ex14

• Different terminal operations combine All words in Shakespeare's works partial results in different ways, e.g. reduce() creates a new 1<sup>st</sup> half of words 2nd half of words immutable value collect() mutates an 2nd quarter of words 1<sup>st</sup> quarter of words 3rd quarter of words 4<sup>th</sup> quarter of words existing value Process Process Process Process sequentially sequentially sequentially sequentially List<CharSequence> uniqueWords = getInput(sSHAKESPEARE), "\\s+") collect() collect() .parallelStream() collect() .collect(toCollection(TreeSet::new)); collect() mutates a container to accumulate the result it's producing



• • •

.collect(ConcurrentHashSetCollector.toSet());

Concurrent collectors are different than non-concurrent collectors (covered later)

 More discussion about reduce() vs. collect() appears online



#### See <a href="https://www.youtube.com/watch?v=oWIWEKNM5Aw">www.youtube.com/watch?v=oWIWEKNM5Aw</a>

- More discussion about reduce() vs. collect() appears online, e.g.
  - Always test w/a parallel stream to detect mistakes wrt mutable vs. immutable reductions

```
void buggyStreamReduce
  (boolean parallel) {
```

```
Stream<String> wordStream =
   allWords.stream();
```

```
if (parallel)
  wordStream.parallel();
```

```
String words = wordStream
.reduce(new StringBuilder(),
    StringBuilder::append,
    StringBuilder::append)
.toString();
```

See github.com/douglascraigschmidt/LiveLessons/tree/master/Java8/ex17

- More discussion about reduce() vs. collect() appears online, e.g.
  - Always test w/a parallel stream to detect mistakes wrt mutable vs. immutable reductions

This code fails when parallel() is used since reduce() expects to do an "immutable" reduction

```
void buggyStreamReduce
  (boolean parallel) {
```

```
Stream<String> wordStream =
  allWords.stream();
if (parallel)
 wordStream.parallel();
String words = wordStream
  .reduce(new StringBuilder(),
          StringBuilder::append,
          StringBuilder::append)
  .toString();
```

- More discussion about reduce() vs. collect() appears online, e.g.
  - Always test w/a parallel stream to detect mistakes wrt mutable vs. immutable reductions

There are race conditions here since

there's just one shared StringBuilder,

which is not properly thread-safe ...

```
void buggyStreamReduce
  (boolean parallel) {
```

```
Stream<String> wordStream =
  allWords.stream();
if (parallel)
 wordStream.parallel();
String words = wordStream
  .reduce(new StringBuilder(),
          StringBuilder::append,
          StringBuilder::append)
```

```
.toString();
```

- More discussion about reduce() vs. collect() appears online, e.g.
  - Always test w/a parallel stream to detect mistakes wrt mutable vs. immutable reductions

```
A stream can be dynamically switched to "parallel" mode!
```

```
void buggyStreamReduce
   (boolean parallel) {
```

```
Stream<String> wordStream =
   allWords.stream();
```

```
if (parallel)
    wordStream.parallel();
```

```
String words = wordStream
  .reduce(new StringBuilder(),
        StringBuilder::append,
        StringBuilder::append)
  .toString();
```

- More discussion about reduce() vs. collect() appears online, e.g.
  - Always test w/a parallel stream to detect mistakes wrt mutable vs. immutable reductions
  - Beware of issues related to association & identity

```
void testDifferenceReduce(...) {
  long difference = LongStream
    .rangeClosed(1, 100)
    .parallel()
    .reduce(OL,
             (x, y) \rightarrow x - y);
}
void testSum(long identity, ...) {
  long sum = LongStream
    .rangeClosed(1, 100)
    .reduce(identity,
     // Could use (x, y) \rightarrow x + y
             Math::addExact);
```

See github.com/douglascraigschmidt/LiveLessons/tree/master/Java8/ex17

- More discussion about reduce() vs. collect() appears online, e.g.
  - Always test w/a parallel stream to detect mistakes wrt mutable vs. immutable reductions
  - Beware of issues related to association & identity

*This code fails for a parallel stream since subtraction is not associative*  

```
void testSum(long identity, ...) {
  long sum = LongStream
   .rangeClosed(1, 100)
   .reduce(identity,
   // Could use (x, y) -> x + y
        Math::addExact);
```

See <a href="https://www.ibm.com/developerworks/library/j-java-streams-2-brian-goetz">www.ibm.com/developerworks/library/j-java-streams-2-brian-goetz</a>

- More discussion about reduce() vs. collect() appears online, e.g.
  - Always test w/a parallel stream to detect mistakes wrt mutable vs. immutable reductions
  - Beware of issues related to association & identity

This code fails if

identity is not OL

void testDifferenceReduce(...) { long difference = LongStream .rangeClosed(1, 100) .parallel() .reduce(0L,  $(x, y) \rightarrow x - y);$ void testSum(long identity, ...) { long sum = LongStream .rangeClosed(1, 100) .reduce(identity, // Could use  $(x, y) \rightarrow x + y$ Math::addExact);

The "identity" of an OP is defined as "identity OP value == value"

• Collector defines an interface whose implementations can accumulate input elements in a mutable result container

#### Interface Collector<T,A,R>

#### **Type Parameters:**

- ${\sf T}$  the type of input elements to the reduction operation
- A the mutable accumulation type of the reduction operation (often hidden as an implementation detail)
- ${\sf R}$  the result type of the reduction operation

#### public interface Collector<T,A,R>

A mutable reduction operation that accumulates input elements into a mutable result container, optionally transforming the accumulated result into a final representation after all input elements have been processed. Reduction operations can be performed either sequentially or in parallel.

Examples of mutable reduction operations include: accumulating elements into a Collection; concatenating strings using a StringBuilder; computing summary information about elements such as sum, min, max, or average; computing "pivot table" summaries such as "maximum valued transaction by seller", etc. The class Collectors provides implementations of many common mutable reductions.

A **Collector** is specified by four functions that work together to accumulate entries into a mutable result container, and optionally perform a final transform on the result. They are:

#### See docs.oracle.com/javase/8/docs/api/java/util/stream/Collector.html

 Collector implementations can either be non-concurrent or concurrent based on their characteristics

#### **Enum Collector.Characteristics**

java.lang.Object

java.lang.Enum<Collector.Characteristics> java.util.stream.Collector.Characteristics

All Implemented Interfaces: Serializable, Comparable<Collector.Characteristics>

Enclosing interface:

Collector<T,A,R>

public static enum Collector.Characteristics
extends Enum<Collector.Characteristics>

Characteristics indicating properties of a  ${\tt Collector},$  which can be used to optimize reduction implementations.

#### **Enum Constant Summary**

#### **Enum Constants**

Enum Constant and Description

#### CONCURRENT

Indicates that this collector is *concurrent*, meaning that the result container can support the accumulator function being called concurrently with the same result container from multiple threads.

#### IDENTITY\_FINISH

Indicates that the finisher function is the identity function and can be elided.

#### UNORDERED

Indicates that the collection operation does not commit to preserving the encounter order of input elements.

See <a href="https://docs.oracle.com/javase/8/docs/api/java/util/stream/Collector.Characteristics.html">docs.oracle.com/javase/8/docs/api/java/util/stream/Collector.Characteristics.html</a>

- Collector implementations can either be non-concurrent or concurrent based on their characteristics
  - This distinction is only relevant for *parallel* streams



See "Overview of Java 8 Streams (Part 4)" for non-concurrent collector implementation

- Collector implementations can either be non-concurrent or concurrent based on their characteristics
  - This distinction is only relevant for *parallel* streams
  - A non-concurrent collector can be used for either a sequential stream or a parallel stream!



#### We'll just focus on parallel streams in the subsequent discussion

• A non-concurrent collector operates by merging sub-results



See <a href="stackoverflow.com/questions/22350288/parallel-streams-collectors-and-thread-safety">stackoverflow.com/questions/22350288/parallel-streams-collectors-and-thread-safety</a>

- A non-concurrent collector operates by merging sub-results
  - The input source is partitioned into chunks



- A non-concurrent collector operates by merging sub-results
  - The input source is partitioned into chunks
  - Each chunk is collected into a result container
    - e.g., a list or a map



- A non-concurrent collector operates by merging sub-results
  - The input source is partitioned into chunks
  - Each chunk is collected into a result container
    - e.g., a list or a map



Different threads operate on different instances of intermediate result containers

- A non-concurrent collector operates by merging sub-results
  - The input source is partitioned into chunks
  - Each chunk is collected into a result container
  - These sub-results are then merged into a final mutable result container
    - Only one thread in the fork-join pool is used to merge any pair of intermediate results



- A non-concurrent collector operates by merging sub-results
  - The input source is partitioned into chunks
  - Each chunk is collected into a result container
  - These sub-results are then merged into a final mutable result container
    - Only one thread in the fork-join pool is used to merge any pair of intermediate results



Thus there's no need for any synchronizers in a non-concurrent collector

- A non-concurrent collector operates by merging sub-results
  - The input source is partitioned into chunks
  - Each chunk is collected into a result container
  - These sub-results are then merged into a final mutable result container

*This process is safe & orderpreserving, but merging is costly for containers like maps & sets* 



 A concurrent collector creates one concurrent result container & inserts elements into it from multiple threads in a parallel stream



See <a href="stackoverflow.com/questions/22350288/parallel-streams-collectors-and-thread-safety">stackoverflow.com/questions/22350288/parallel-streams-collectors-and-thread-safety</a>

- A concurrent collector creates one concurrent result container & inserts elements into it from multiple threads in a parallel stream
  - As usual, the input source is partitioned into chunks



- A concurrent collector creates one concurrent result container & inserts elements into it from multiple threads in a parallel stream
  - As usual, the input source is partitioned into chunks
  - Each chunk is collected into one concurrent result container
    - e.g., a concurrent map or set



- A concurrent collector creates one concurrent result container & inserts elements into it from multiple threads in a parallel stream
  - As usual, the input source is partitioned into chunks
  - Each chunk is collected into one concurrent result container
    - e.g., a concurrent map or set



Different threads in a parallel stream share one concurrent result container

- A concurrent collector creates one concurrent result container & inserts elements into it from multiple threads in a parallel stream
  - As usual, the input source is partitioned into chunks
  - Each chunk is collected into one concurrent result container

Thus there's no need to merge any intermediate sub-results!



• A concurrent collector may perform better than a non-concurrent collector if merging costs are high





See github.com/douglascraigschmidt/LiveLessons/tree/master/Java8/ex14

- A concurrent collector may perform better than a non-concurrent collector if merging costs are high
  - e.g., for a highly optimized result container like ConcurrentHashMap vs. merging HashMaps



See <a href="https://www.quora.com/What-is-the-difference-between-synchronize-and-concurrent-collection-in-Java">www.quora.com/What-is-the-difference-between-synchronize-and-concurrent-collection-in-Java</a>

• The Collector interface defines three generic types



<<Java Interface>>

Collector<T,A,R>

- supplier():Supplier<A>
- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>
- characteristics():Set<Characteristics>

#### See <a href="http://www.baeldung.com/java-8-collectors">www.baeldung.com/java-8-collectors</a>

- The Collector interface defines three generic types
  - T The type of objects available in the stream
    - e.g., Integer, String, etc.

<<Java Interface>>
Collector<T,A,R>

- supplier():Supplier<A>
- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>
- characteristics():Set<Characteristics>

- The Collector interface defines three generic types
  - A The type of a mutable accumulator object for collection
    - e.g., ConcurrentHashSet or ArrayList of T

<<Java Interface>>
Collector<TAR>

- supplier():Supplier<A>
- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>
- characteristics():Set<Characteristics>

- The Collector interface defines three generic types
  - T
  - A
  - **R** The type of a final result
    - e.g., ConcurrentHashSet or ArrayList of T

<<Java Interface>>

Collector<T,AR>

- supplier():Supplier<A>
- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>
- characteristics():Set<Characteristics>

#### See <a href="https://www.baeldung.com/java-8-collectors">www.baeldung.com/java-8-collectors</a>

• Five methods are defined in the Collector interface



<<Java Interface>>

Collector<T,A,R>

- supplier():Supplier<A>
- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>
- characteristics():Set<Characteristics>

- Five methods are defined in the Collector interface
  - characteristics() provides a stream with additional information used for internal optimizations, e.g.
    - UNORDERED
      - The collector need not preserve the encounter order

<<Java Interface>>

Collector<T,A,R>

- supplier():Supplier<A>
- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>

characteristics():Set<Characteristics>

A concurrent collector *should* be UNORDERED, but a non-concurrent collector *can* be ORDERED

- Five methods are defined in the Collector interface
  - characteristics() provides a stream with additional information used for internal optimizations, e.g.

UNORDERED

- IDENTIFY\_FINISH
  - The finisher() is the identity function so it can be a no-op
    - e.g. finisher() just returns null

<<Java Interface>>

Collector<T,A,R>

- supplier():Supplier<A>
- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>

characteristics():Set<Characteristics>

A concurrent collector *should* be IDENTITY\_FINISH, whereas a non-concurrent collector *could* be

- Five methods are defined in the Collector interface
  - characteristics() provides a stream with additional information used for internal optimizations, e.g.
    - UNORDERED
    - IDENTIFY\_FINISH
    - CONCURRENT

<<Java Interface>>

Collector<T,A,R>

- supplier():Supplier<A>
- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>

characteristics():Set<Characteristics>

- The accumulator() method is called concurrently on the result container
  - Naturally, the mutable result container must be synchronized!!

A concurrent collector *should* be CONCURRENT, but a non-concurrent collector should *not* be!

- Five methods are defined in the Collector interface
  - characteristics() provides a stream with additional information used for internal optimizations, e.g.
    - UNORDERED
    - IDENTIFY\_FINISH
    - CONCURRENT

<<Java Interface>>

Collector<T,A,R>

- supplier():Supplier<A>
- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>

characteristics():Set<Characteristics>

- The accumulator() method is called concurrently on the result container
- The combiner() method is a no-op since it's not called at all

- Five methods are defined in the Collector interface
  - characteristics() provides a stream with additional information used for internal optimizations, e.g.
    - UNORDERED
    - IDENTIFY\_FINISH
    - CONCURRENT

<<Java Interface>>

Collector<T,A,R>

- supplier():Supplier<A>
- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>

characteristics():Set<Characteristics>



A non-concurrent collector can be used with either sequential or parallel streams

- Five methods are defined in the Collector interface
  - characteristics() provides a stream with additional information used for internal optimizations, e.g.
    - UNORDERED
    - IDENTIFY\_FINISH
    - CONCURRENT

return Collections.unmodifiableSet
 (EnumSet.of(Collector.Characteristics.CONCURRENT,

Collector.Characteristics.UNORDERED,

Collector.Characteristics.IDENTITY\_FINISH));

See github.com/douglascraigschmidt/LiveLessons/tree/master/Java8/ex14

<<Java Class>>

GConcurrentHashSetCollector<T>
ConcurrentHashSetCollector()
supplier
Supplier

- accumulator():BiConsumer<ConcurrentHashSet<T>,T>
- ocombiner():BinaryOperator<ConcurrentHashSet<T>>
- ofinisher():Function<ConcurrentHashSet<T>,ConcurrentHashSet<T>>

characteristics():Set

toSet():Collector<E,?,ConcurrentHashSet<E>>

- Five methods are defined in the Collector interface
  - characteristics()
  - supplier() returns a supplier instance that generates an empty result container

<<Java Interface>>

Collector<T,A,R>

supplier():Supplier<A>

- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>

characteristics():Set<Characteristics>

- Five methods are defined in the Collector interface
  - characteristics()
  - **supplier()** returns a supplier instance that generates an empty result container, e.g.
    - return ArrayList::new

<<Java Interface>>

Collector<T,A,R>

#### supplier():Supplier<A>

- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>
- characteristics():Set<Characteristics>

A non-concurrent collector has a different result container for each thread in a parallel stream

- Five methods are defined in the Collector interface
  - characteristics()
  - **supplier()** returns a supplier instance that generates an empty result container, e.g.
    - return ArrayList::new
    - return ConcurrentHashSet::new

<<Java Class>>
ConcurrentHashSetCollector<T>
ConcurrentHashSetCollector<T>
ConcurrentHashSetCollector()
Supplier():Supplier<ConcurrentHashSet<T>>
Combiner():BiConsumer<ConcurrentHashSet<T>,T>
Combiner():BinaryOperator<ConcurrentHashSet<T>>
finisher():Function<ConcurrentHashSet<T>,ConcurrentHashSet<T>>
ConcurrentHashSet<T>,ConcurrentHashSet<T>>
ConcurrentHashSet<T>>
Concur

A concurrent collector has one result container shared by each thread in a parallel stream

- Five methods are defined in the Collector interface
  - characteristics()
  - supplier()
  - accumulator() returns a biconsumer that adds a new element to result container

<<Java Interface>>

Collector<T,A,R>

- supplier():Supplier<A>
- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>
- characteristics():Set<Characteristics>

- Five methods are defined in the Collector interface
  - characteristics()
  - supplier()
  - accumulator() returns a biconsumer that adds a new element to result container, e.g.
    - return ArrayList::add

<<Java Interface>>

Collector<T,A,R>

- supplier():Supplier<A>
- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>
- characteristics():Set<Characteristics>

A non-concurrent collector's methods should *not* be synchronized!

- Five methods are defined in the Collector interface
  - characteristics()
  - supplier()
  - accumulator() returns a biconsumer that adds a new element to result container, e.g.
    - return ArrayList::add
    - return ConcurrentHashSet::add

< <java class="">&gt;</java>	
GConcurrentHashSetCollector <t></t>	
ConcurrentHashSetCollector()	
supplier():Supplier <concurrenthashset<t>&gt;</concurrenthashset<t>	
accumulator():BiConsumer <concurrenthashset<t>,T&gt;</concurrenthashset<t>	
combiner():BinaryOperator <concurrenthashset<t>&gt;</concurrenthashset<t>	•
finisher():Function <concurrenthashset<t>,ConcurrentHashSet<t>&gt;</t></concurrenthashset<t>	
ocharacteristics():Set	

toSet():Collector<E.?.ConcurrentHashSet<E>>

A concurrent collector's methods *must* be synchronized!

- Five methods are defined in the Collector interface
  - characteristics()
  - supplier()
  - accumulator()
  - combiner() returns a function that merges two result containers together

<<Java Interface>>

Collector<T,A,R>

- supplier():Supplier<A>
- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>
- characteristics():Set<Characteristics>

- Five methods are defined in the Collector interface
  - characteristics()
  - supplier()
  - accumulator()
  - **combiner()** returns a function that merges two result containers together, e.g.

60

- Five methods are defined in the Collector interface
  - characteristics()
  - supplier()
  - accumulator()
  - **combiner()** returns a function that merges two result containers together, e.g.

<<Java Interface>>

Collector<T,A,R>

- supplier():Supplier<A>
- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>
- characteristics():Set<Characteristics>

• return null

The combiner() method is not called when CONCURRENT is set

- Five methods are defined in the Collector interface
  - characteristics()
  - supplier()
  - accumulator()
  - combiner()
  - finisher() returns a function that converts the result container to final result type

<<Java Interface>>

Collector<T,A,R>

- supplier():Supplier<A>
- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>

characteristics():Set<Characteristics>

- Five methods are defined in the Collector interface
  - characteristics()
  - supplier()
  - accumulator()
  - combiner()
  - **finisher()** returns a function that converts the result container to final result type, e.g.

• Function.identity() or something much more interesting!

<<Java Interface>>

Collector<T,A,R>

- supplier():Supplier<A>
- accumulator():BiConsumer<A,T>
- combiner():BinaryOperator<A>
- finisher():Function<A,R>
- characteristics():Set<Characteristics>

- Five methods are defined in the Collector interface
  - characteristics()
  - supplier()
  - accumulator()
  - combiner()
  - **finisher()** returns a function that converts the result container to final result type, e.g.
    - Function.identity()
    - return null

<<Java Class>>

G ConcurrentHashSetCollector<T>

G ConcurrentHashSetCollector()

supplier():Supplier<ConcurrentHashSet<T>>

accumulator():BiConsumer<ConcurrentHashSet<T>,T>

combiner():BinaryOperator<ConcurrentHashSet<T>>

finisher():Function<ConcurrentHashSet<T>,ConcurrentHashSet<T>>

characteristics():Set

StoSet():Collector<E,?,ConcurrentHashSet<E>>

#### The finisher() method is not called when IDENTITY\_FINISHER is set

# End of Java 8 Parallel Stream Internals (Part 5)