

Debugging Spark Applications

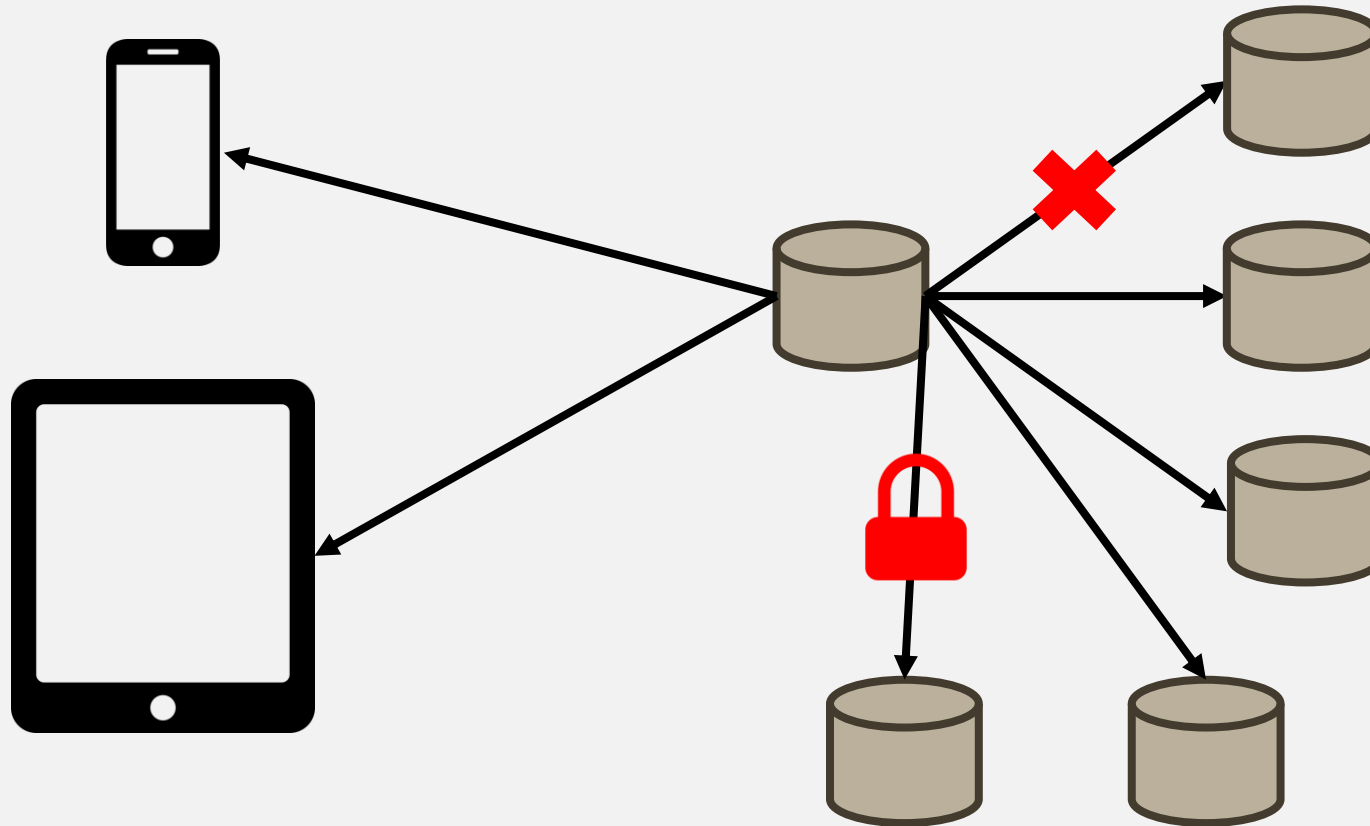
Melike GEÇER

Supervised by

Haidar Osman

25.06.2019

Why is debugging hard in distributed systems?



Why is debugging hard?

Lack of tooling

```
Java HotSpot(TM) 64-Bit Server VM warning: ignoring option MaxPermSize=256m; support was removed in 8.0
SLF4J: Class path contains multiple SLF4J bindings.
SLF4J: Found binding in [jar:file:/tmp/hadoop-melike/nm-local-dir/usercache/melike/filecache/10/__spark_libs__7564521103123317951.zip/slf4j-log
SLF4J: Found binding in [jar:file:/opt/hadoop-3.1.2/share/hadoop/common/lib/slf4j-log4j12-1.7.25.jar!/org/slf4j/impl/StaticLoggerBinder.class]
SLF4J: See http://www.slf4j.org/codes.html#multiple_bindings for an explanation.
SLF4J: Actual binding is of type [org.slf4j.impl.Log4jLoggerFactory]
2019-06-24 00:40:02,129 INFO util.SignalUtils: Registered signal handler for TERM
2019-06-24 00:40:02,131 INFO util.SignalUtils: Registered signal handler for HUP
2019-06-24 00:40:02,131 INFO util.SignalUtils: Registered signal handler for INT
2019-06-24 00:40:02,720 WARN util.NativeCodeLoader: Unable to load native-hadoop library for your platform... using builtin-java classes where
2019-06-24 00:40:02,973 INFO yarn.ApplicationMaster: Preparing Local resources
2019-06-24 00:40:03,935 INFO yarn.ApplicationMaster: Prepared Local resources Map(__app__.jar -> resource { scheme: "hdfs" host: "leela.unibe.c
2019-06-24 00:40:03,970 INFO yarn.ApplicationMaster: ApplicationAttemptId: appattempt_1561329314867_0002_000002
2019-06-24 00:40:03,980 INFO spark.SecurityManager: Changing view acls to: melike
2019-06-24 00:40:03,981 INFO spark.SecurityManager: Changing modify acls to: melike
2019-06-24 00:40:03,981 INFO spark.SecurityManager: Changing view acls groups to:
2019-06-24 00:40:03,982 INFO spark.SecurityManager: Changing modify acls groups to:
2019-06-24 00:40:03,983 INFO spark.SecurityManager: SecurityManager: authentication disabled; ui acls disabled; users with view permissions: S
2019-06-24 00:40:04,005 INFO yarn.ApplicationMaster: Starting the user application in a separate Thread
2019-06-24 00:40:04,022 INFO yarn.ApplicationMaster: Waiting for spark context initialization
2019-06-24 00:40:04,023 INFO yarn.ApplicationMaster: Waiting for spark context initialization ...
2019-06-24 00:40:04,048 INFO spark.SparkContext: Running Spark version 2.0.2
2019-06-24 00:40:04,099 INFO spark.SecurityManager: Changing view acls to: melike
2019-06-24 00:40:04,099 INFO spark.SecurityManager: Changing modify acls to: melike
2019-06-24 00:40:04,099 INFO spark.SecurityManager: Changing view acls groups to:
2019-06-24 00:40:04,099 INFO spark.SecurityManager: Changing modify acls groups to:
2019-06-24 00:40:04,099 INFO spark.SecurityManager: SecurityManager: authentication disabled; ui acls disabled; users with view permissions: S
2019-06-24 00:40:04,427 INFO util.Utils: Successfully started service 'sparkDriver' on port 42097.
2019-06-24 00:40:04,457 INFO spark.SparkEnv: Registering MapOutputTracker
2019-06-24 00:40:04,484 INFO spark.SparkEnv: Registering BlockManagerMaster
2019-06-24 00:40:04,502 INFO storage.DiskBlockManager: Created local directory at /tmp/hadoop-melike/nm-local-dir/usercache/melike/appcache/app
2019-06-24 00:40:04,523 INFO memory.MemoryStore: MemoryStore started with capacity 93.3 MB
2019-06-24 00:40:04,641 INFO spark.SparkEnv: Registering OutputCommitCoordinator
2019-06-24 00:40:04,758 INFO util.log: Logging initialized @3466ms
2019-06-24 00:40:04,813 INFO ui.JettyUtils: Adding filter: org.apache.hadoop.yarn.server.webproxy.amfilter.AmIpFilter
2019-06-24 00:40:04,880 INFO server.Server: jetty-9.2.z-SNAPSHOT
2019-06-24 00:40:04,908 INFO handler.ContextHandler: Started o.s.j.s.ServletContextHandler@557d5a49{/jobs,null,AVAILABLE}
2019-06-24 00:40:04,908 INFO handler.ContextHandler: Started o.s.j.s.ServletContextHandler@205eac11{/jobs/job,null,AVAILABLE}
```

Why is debugging hard?

Data is always dirty

Student Name	Report File Size	Submission Date
Skyla Wong	10 GB	13.10.2018
Karla Parkes	2 MB	12 Oct 2018
Archer Nunez	1.5 MB	Second of April
Danny Oakley	1.9 MB	A very very long sentence that is beyond the expected size of the field

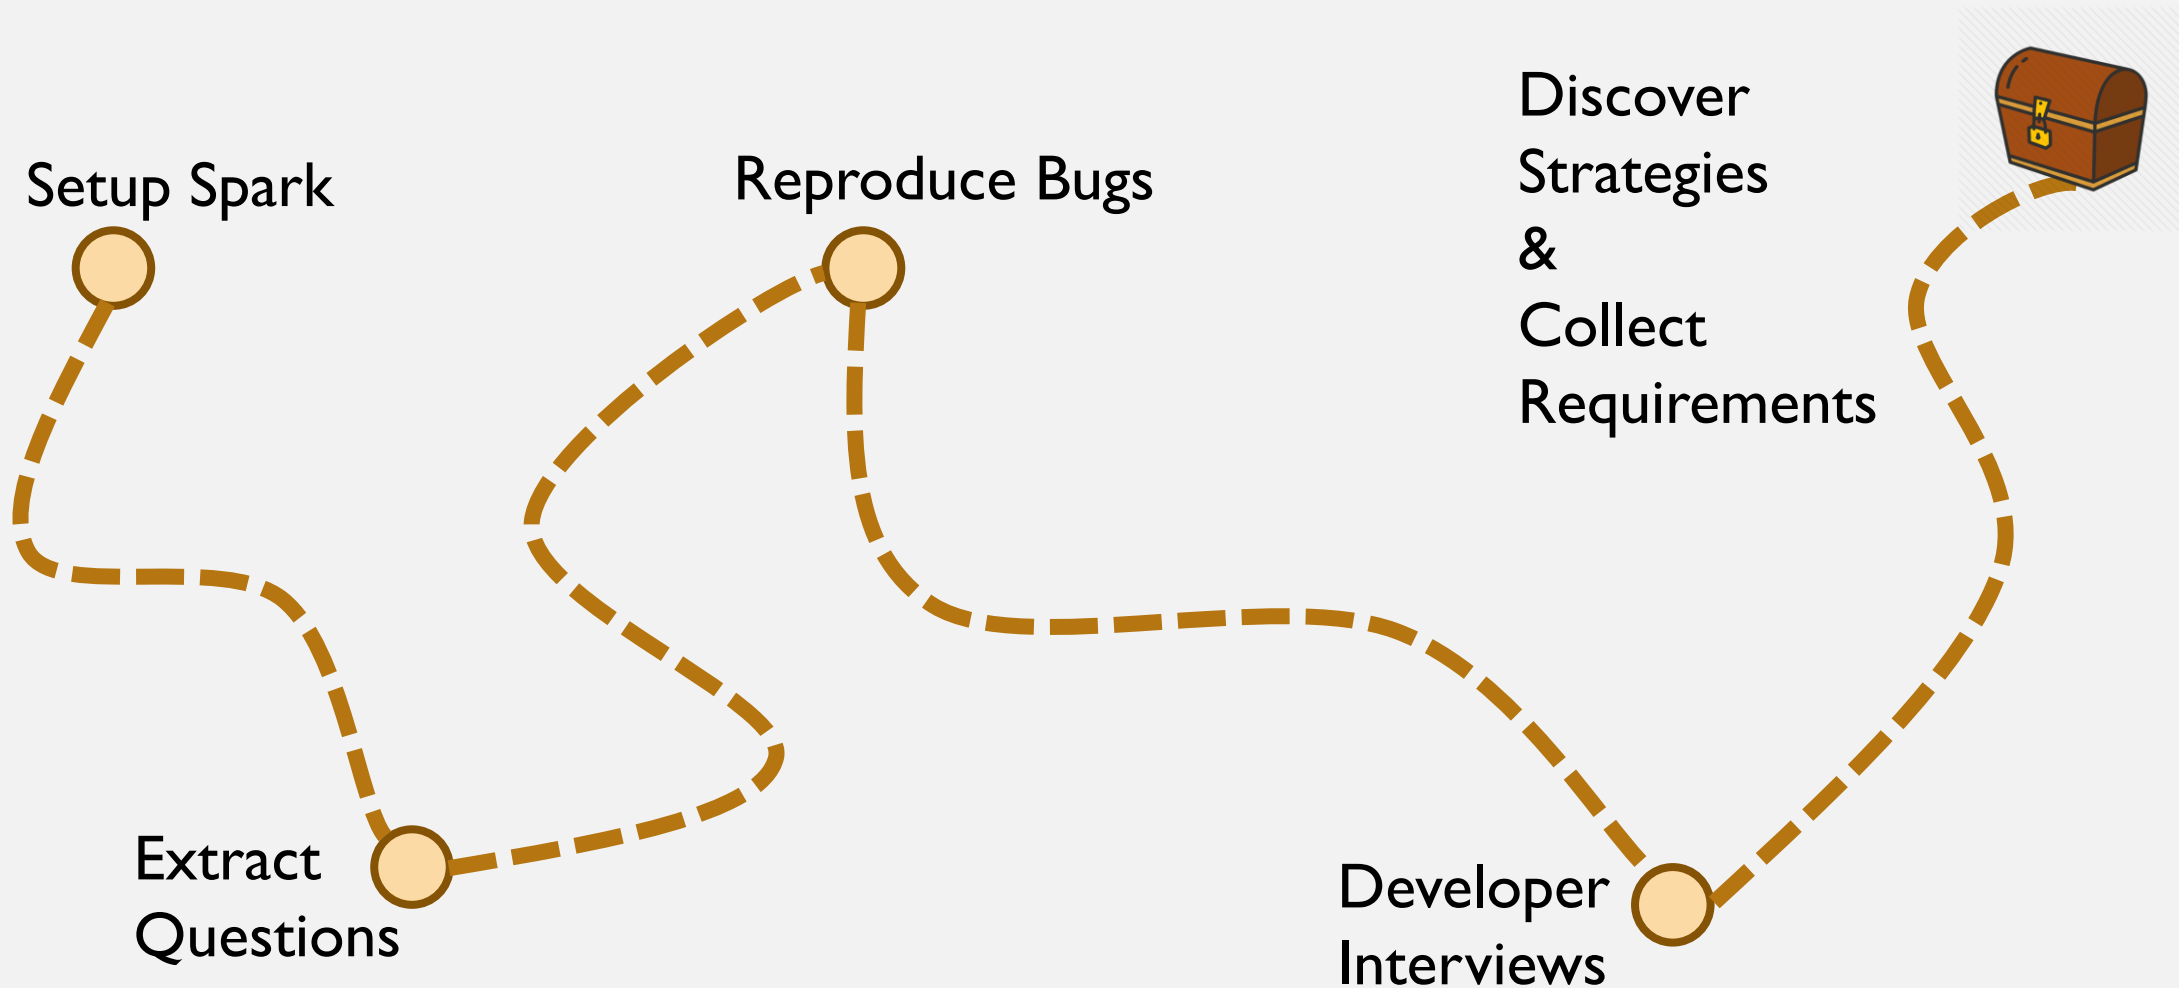
Research Questions

RQ1: How do developers debug distributed applications like Spark?

RQ2: What are the challenges that developers face when they debug such applications?

RQ3: How can we overcome these challenges?

Investigation Roadmap



Setup Spark



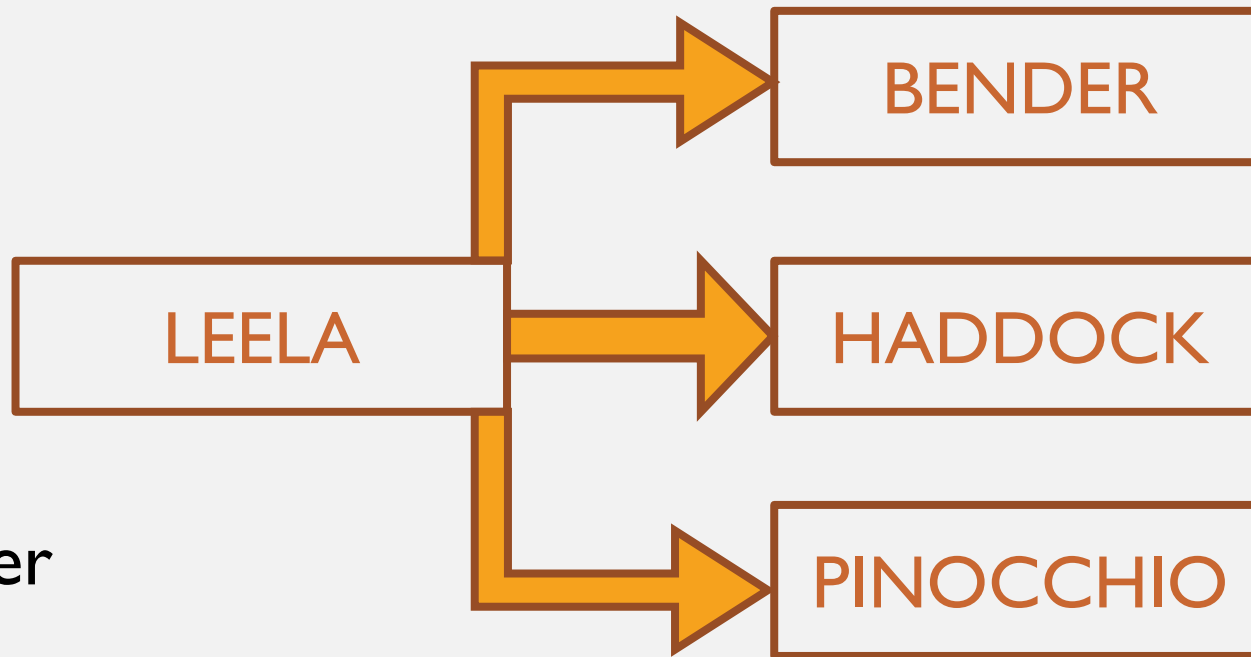
ID	User	Name	Application Type	Queue
application_1561329314867_0002	melike	MainClass	SPARK	default

```
Java HotSpot(TM) 64-Bit Server VM warning: ignoring option MaxPermSize=256m; support was removed in 8.0
SLF4J: Class path contains multiple SLF4J bindings.
SLF4J: Found binding in [jar:file:/tmp/hadoop-melike/nm-local-dir/usercache/melike/filecache/10/_spark_libs_7564521103123317951.zip/slf4j-log4j12-1.7.16.jar!/org/slf4j/
SLF4J: Found binding in [jar:file:/opt/hadoop-3.1.2/share/hadoop/common/lib/slf4j-log4j12-1.7.25.jar!/org/slf4j/impl/StaticLoggerBinder.class]
SLF4J: See http://www.slf4j.org/codes.html#multiple_bindings for an explanation.
SLF4J: Actual binding is of type [org.slf4j.impl.Log4jLoggerFactory]
2019-06-24 00:40:02,129 INFO util.SignalUtils: Registered signal handler for TERM
2019-06-24 00:40:02,131 INFO util.SignalUtils: Registered signal handler for HUP
2019-06-24 00:40:02,131 INFO util.SignalUtils: Registered signal handler for INT
2019-06-24 00:40:02,720 WARN util.NativeCodeLoader: Unable to load native-hadoop library for your platform... using builtin-java classes where applicable
2019-06-24 00:40:02,973 INFO yarn.ApplicationMaster: Preparing Local resources
2019-06-24 00:40:03,935 INFO yarn.ApplicationMaster: Prepared Local resources Map(__app__.jar -> resource { scheme: "hdfs" host: "leela.unibe.ch" port: 9000 file: "/user/
2019-06-24 00:40:03,970 INFO yarn.ApplicationMaster: ApplicationAttemptId: appattempt_1561329314867_0002_000002
2019-06-24 00:40:03,980 INFO spark.SecurityManager: Changing view acls to: melike
2019-06-24 00:40:03,981 INFO spark.SecurityManager: Changing modify acls to: melike
2019-06-24 00:40:03,981 INFO spark.SecurityManager: Changing view acls groups to:
2019-06-24 00:40:03,982 INFO spark.SecurityManager: Changing modify acls groups to:
```

Setup Spark

Master Node:


- HistoryServer
- DataNode
- NameNode
- NodeManager
- ResourceManager



Worker Nodes:

- NodeManager
- DataNode

Extract Questions

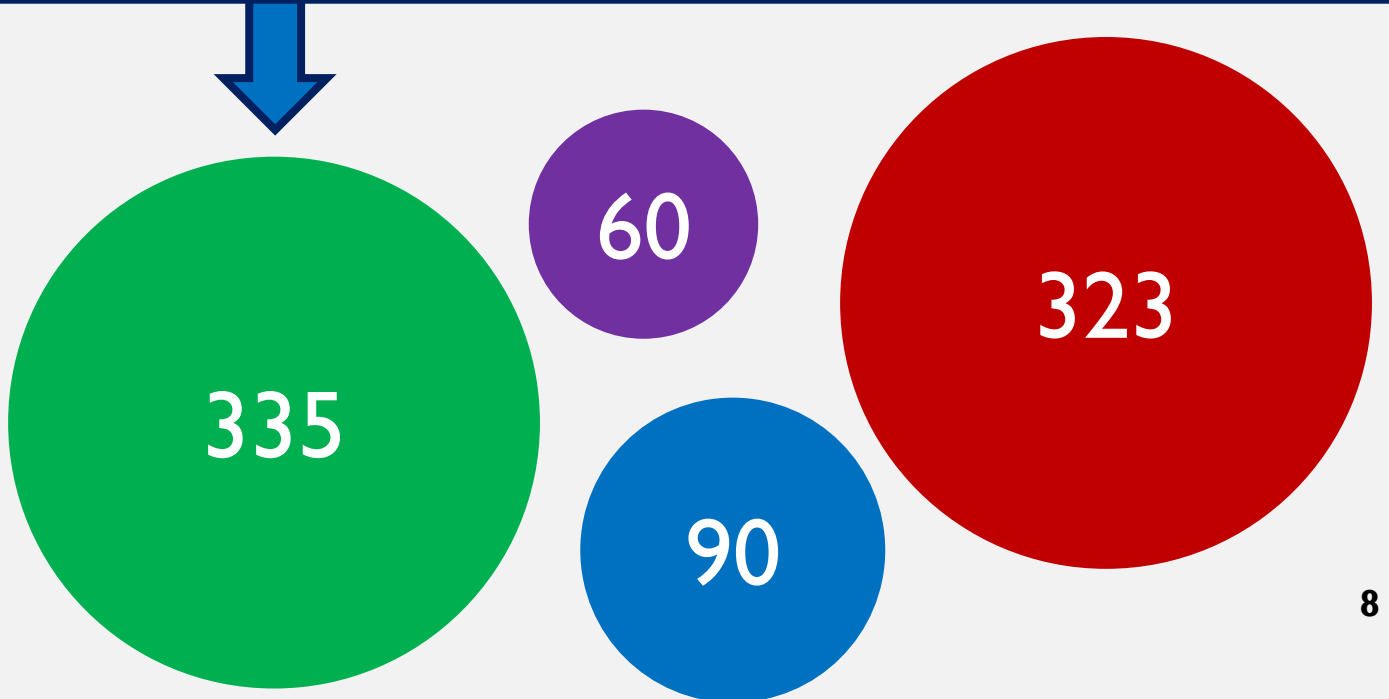


Not Serializable Exception

Out of Memory Error

Class Not Found Exception

Connect Exception



Extract Questions



Stack Overflow Exception is recommended by an expert.

Reproduction of Error Messages: a Sample Case

Not Serializable Exception

...

```
def main (args:Array[String]){  
    val listOfNumbers = Spark.sc.parallelize( List ( 192, 35, 78, 84 ) )  
    new MathOperations().divisibleBy3( listOfNumbers )  
}
```

...

```
class MathOperations {  
    def divisibleBy3 ( list: RDD [ Int ] ) = { ... }  
    def checkDivisibility ( dividend: Int , divider: Int ): Boolean = ...  
}
```

...

Reproduction of Error Messages: a Sample Case

JAR file is generated using Sbt Assembly and the task is submitted via following script:

```
spark-submit \
```

```
--num-executors 3 \
```

```
--master yarn \
```

```
--deploy-mode cluster \
```

```
--class MainClass \
```

```
CodeSampleOne-assembly-0.1.jar
```

- Number of executors to launch

- Run on a YARN cluster

- Deploy on the worker nodes

- The entry point for your application

- Application JAR

```
org.apache.spark.SparkException: Task not serializable
  at org.apache.spark.util.ClosureCleaner$.ensureSerializable(ClosureCleaner.scala:298)
  at org.apache.spark.util.ClosureCleaner$.org$apache$spark$util$ClosureCleaner$$clean(ClosureCleaner.scala:288)
  at org.apache.spark.util.ClosureCleaner$.clean(ClosureCleaner.scala:108)
  at org.apache.spark.SparkContext.clean(SparkContext.scala:2039)
  at org.apache.spark.rdd.RDD$$anonfun$map$1.apply(RDD.scala:366)
  at org.apache.spark.rdd.RDD$$anonfun$map$1.apply(RDD.scala:365)
  at org.apache.spark.rdd.RDDOperationScope$.withScope(RDDOperationScope.scala:151)
  at org.apache.spark.rdd.RDDOperationScope$.withScope(RDDOperationScope.scala:112)
  at org.apache.spark.rdd.RDD.withScope(RDD.scala:358)
  at org.apache.spark.rdd.RDD.map(RDD.scala:365)
  at MathOperations.divisibleBy3(MainClass.scala:27)
  at MainClass$.main(MainClass.scala:13)
  at MainClass.main(MainClass.scala)
  at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
  at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)
  at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
  at java.lang.reflect.Method.invoke(Method.java:498)
  at org.apache.spark.deploy.yarn.ApplicationMaster$$anon$2.run(ApplicationMaster.scala:627)
```

Caused by: `java.io.NotSerializableException: MathOperations`

Serialization stack:

- object not serializable (class: MathOperations, value: MathOperations@54b55111)
 - field (class: MathOperations\$\$anonfun\$1, name: \$outer, type: class MathOperations)
 - object (class MathOperations\$\$anonfun\$1, <function1>)
- ```
at org.apache.spark.serializer.SerializationDebugger$.improveException(SerializationDebugger.scala:40)
at org.apache.spark.serializer.JavaSerializationStream.writeObject(JavaSerializer.scala:46)
at org.apache.spark.serializer.JavaSerializerInstance.serialize(JavaSerializer.scala:100)
at org.apache.spark.util.ClosureCleaner$.ensureSerializable(ClosureCleaner.scala:295)
... 17 more
```

# Investigation Roadmap

