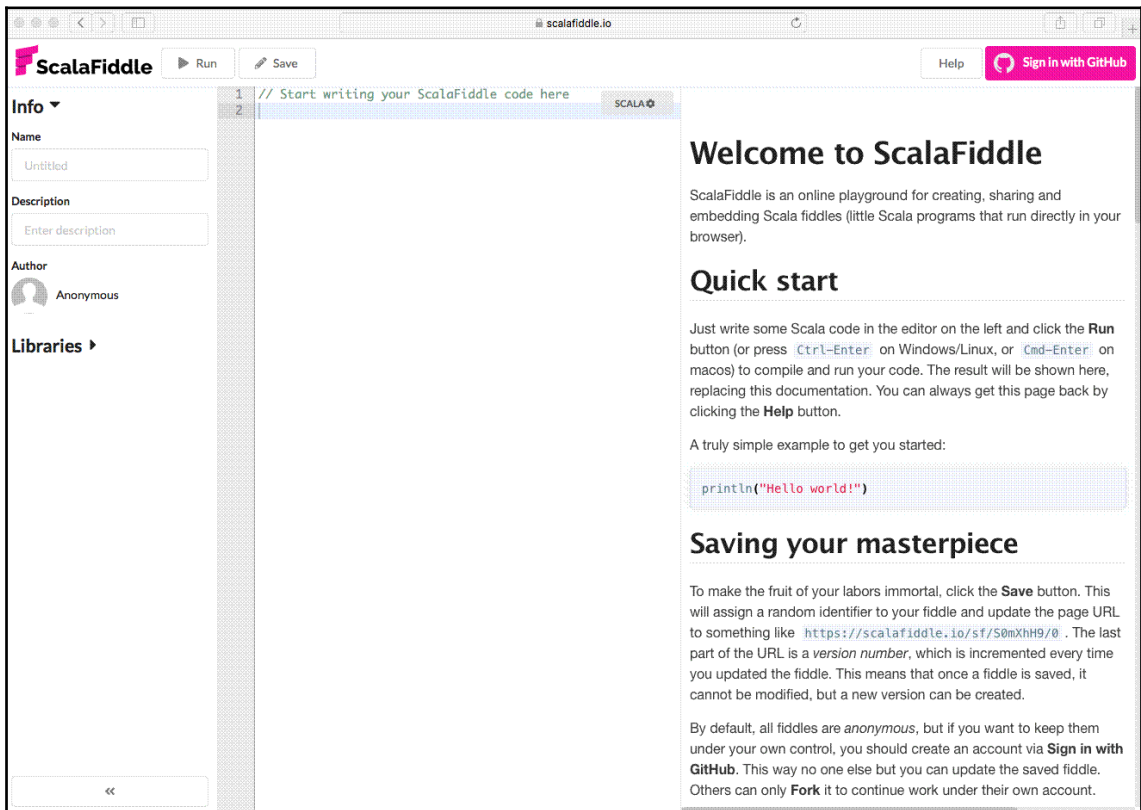


Chapter 1: Scala Overview



The screenshot shows the ScalaFiddle web application interface. At the top, there is a browser window with the address bar showing "scalafiddle.io". The application header includes the ScalaFiddle logo, a "Run" button, a "Save" button, a "Help" button, and a "Sign in with GitHub" button.

On the left side, there is an "Info" panel with a dropdown arrow. It contains fields for "Name" (with "Untitled" as the default value), "Description" (with "Enter description" as the default value), and "Author" (with a profile icon and the name "Anonymous"). Below the "Info" panel is a "Libraries" section with a right-pointing arrow.

The main area is a code editor with a light blue background. It shows two lines of code: line 1 is "// Start writing your ScalaFiddle code here" and line 2 is empty. The editor is titled "SCALA" with a small icon.

On the right side, there is a "Welcome to ScalaFiddle" section. It includes a paragraph: "ScalaFiddle is an online playground for creating, sharing and embedding Scala fiddles (little Scala programs that run directly in your browser)." Below this is a "Quick start" section with a paragraph: "Just write some Scala code in the editor on the left and click the **Run** button (or press **Ctrl-Enter** on Windows/Linux, or **Cmd-Enter** on macos) to compile and run your code. The result will be shown here, replacing this documentation. You can always get this page back by clicking the **Help** button." Below the paragraph is a code block containing the text:

```
println("Hello world!")
```


Below the "Quick start" section is a "Saving your masterpiece" section. It includes a paragraph: "To make the fruit of your labors immortal, click the **Save** button. This will assign a random identifier to your fiddle and update the page URL to something like <https://scalafiddle.io/sf/50mXhH9/0>. The last part of the URL is a *version number*, which is incremented every time you updated the fiddle. This means that once a fiddle is saved, it cannot be modified, but a new version can be created." Below this paragraph is another paragraph: "By default, all fiddles are *anonymous*, but if you want to keep them under your own control, you should create an account via **Sign in with GitHub**. This way no one else but you can update the saved fiddle. Others can only **Fork** it to continue work under their own account."


Plugins


Marketplace Installed Updates (0) ⚙️


🔍 Search plugins in marketplace


Featured [Show All](#)


**Scala**
Languages
Adds support for the Scala language. The following features are available for free with IntelliJ...
🕒 Apr 19, 2019 ↓ 11.9M ☆ 4.5
[Restart IDE](#) ←


**TeamCity**
Tools integration
Provides integration with JetBrains TeamCity.
🕒 Dec 12, 2018 ↓ 622.1K ☆ 3.8
[Install](#)


**IdeaVim**
Editor
Vim emulation plug-in for IDEs based on the IntelliJ platform. IdeaVim supports many Vim...
🕒 Mar 27, 2019 ↓ 5.1M ☆ 4.4
[Install](#)


**BashSupport**
Languages
Bash language support for the IntelliJ platform. Supports run configurations, syntax highlighting,..
🕒 Mar 29, 2019 ↓ 10M ☆ 4.6
[Install](#)

**Dark Purple Theme**
UI
A dark theme in purple tones. For version 2019.1 and above. To install: Go to Settings (Preferences).
🕒 Mar 22, 2019 ↓ 24.9K ☆ 4.9
[Install](#)

**Snyk Vulnerability Scanni...**
Security
This is an IDE extension that helps you detect and fix security issues in your project. By scanning...
🕒 Dec 16, 2018 ↓ 4.6K ☆ 4.7
[Install](#)

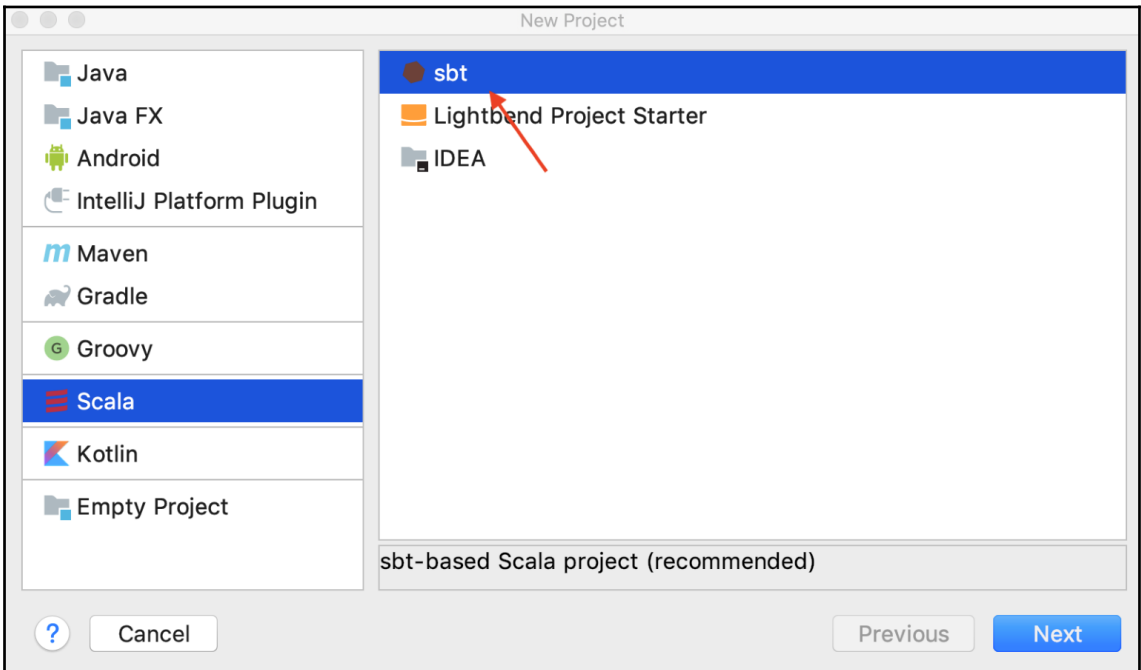
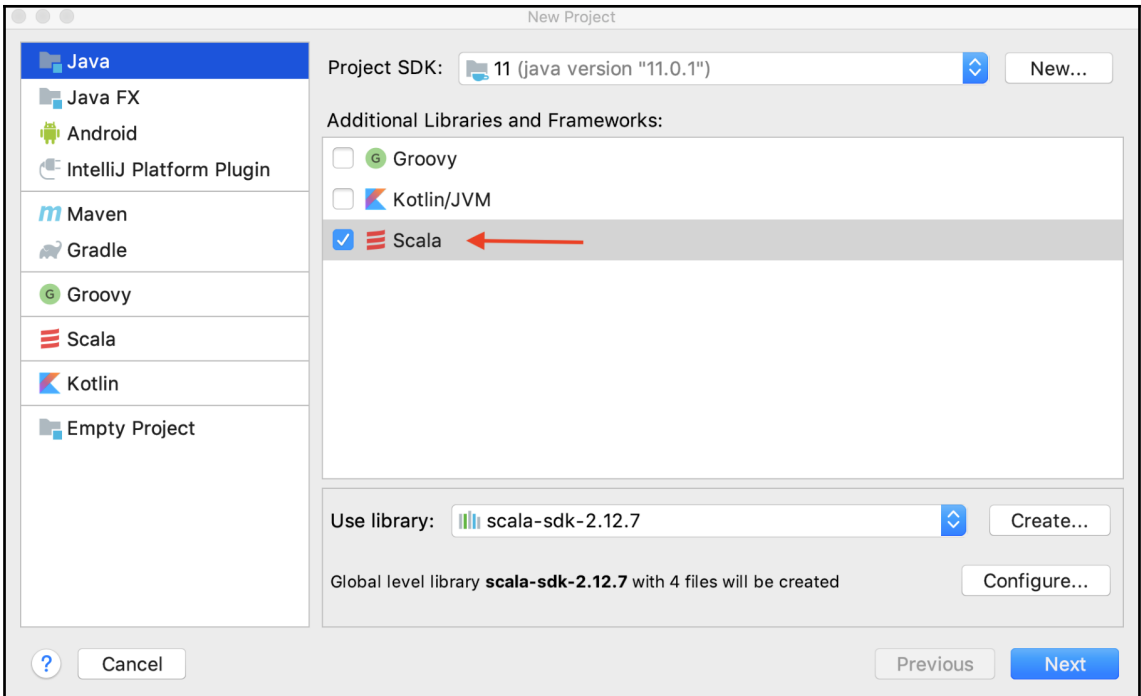
**EduTools**
Code tools
With the EduTools plugin, you can learn and teach programming languages such as Kotlin, Java,...
🕒 Mar 28, 2019 ↓ 249.6K ☆ 3.8
[Install](#)

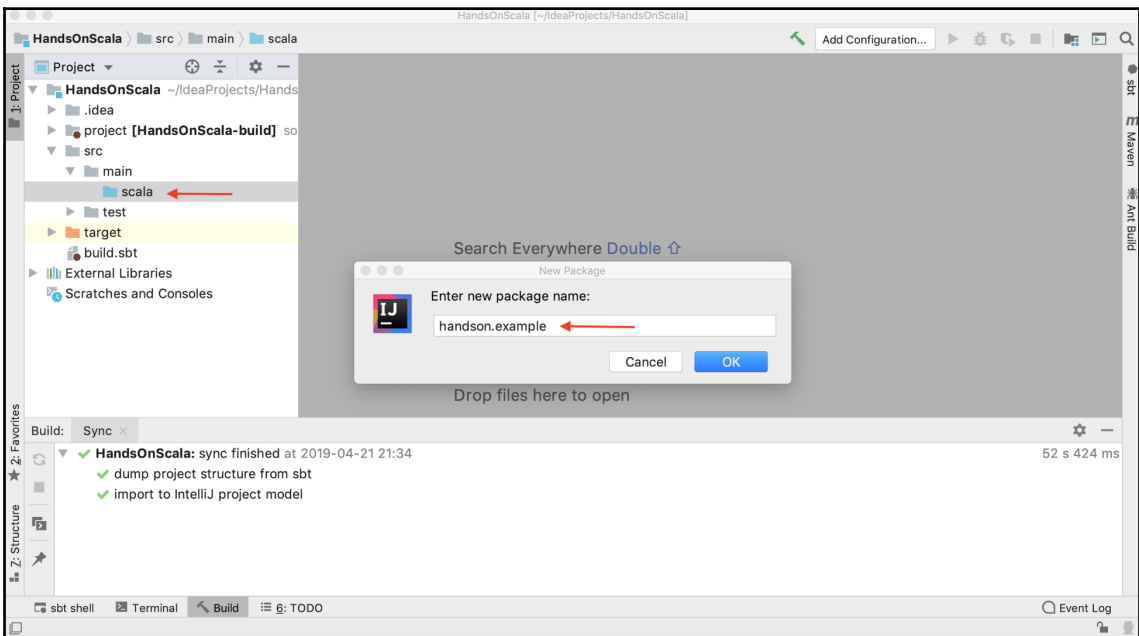
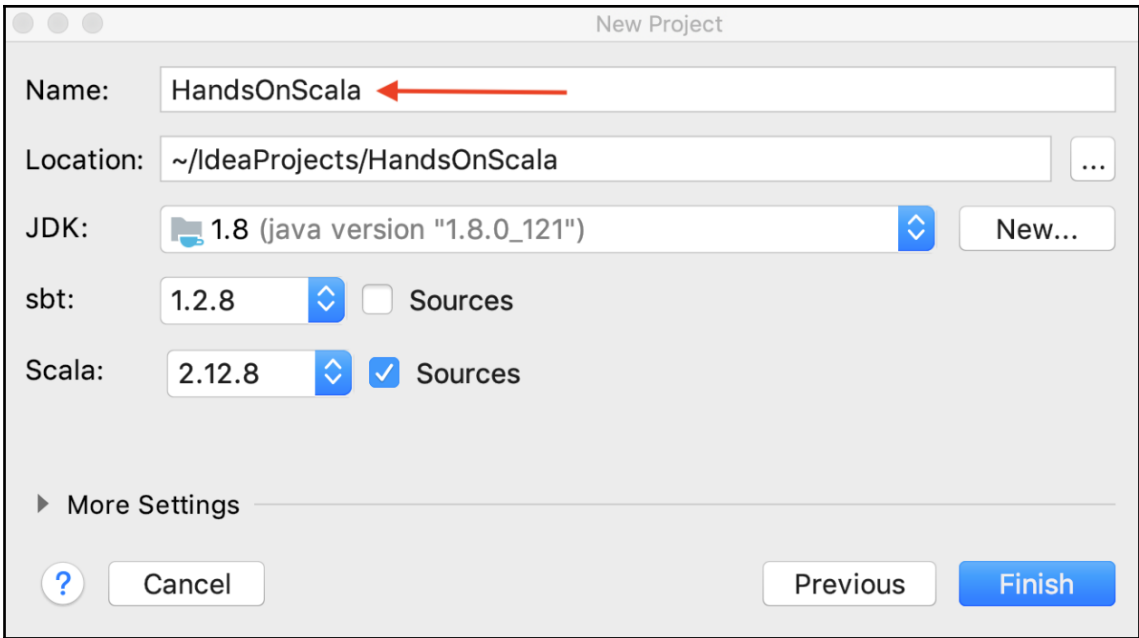
**Key Promoter X**
Apps, Notification and Interaction
A plugin to learn the IntelliJ IDEA shortcuts. The Key Promoter X helps you to learn essential...
🕒 Jan 30, 2019 ↓ 198.6K ☆ 5
[Install](#)

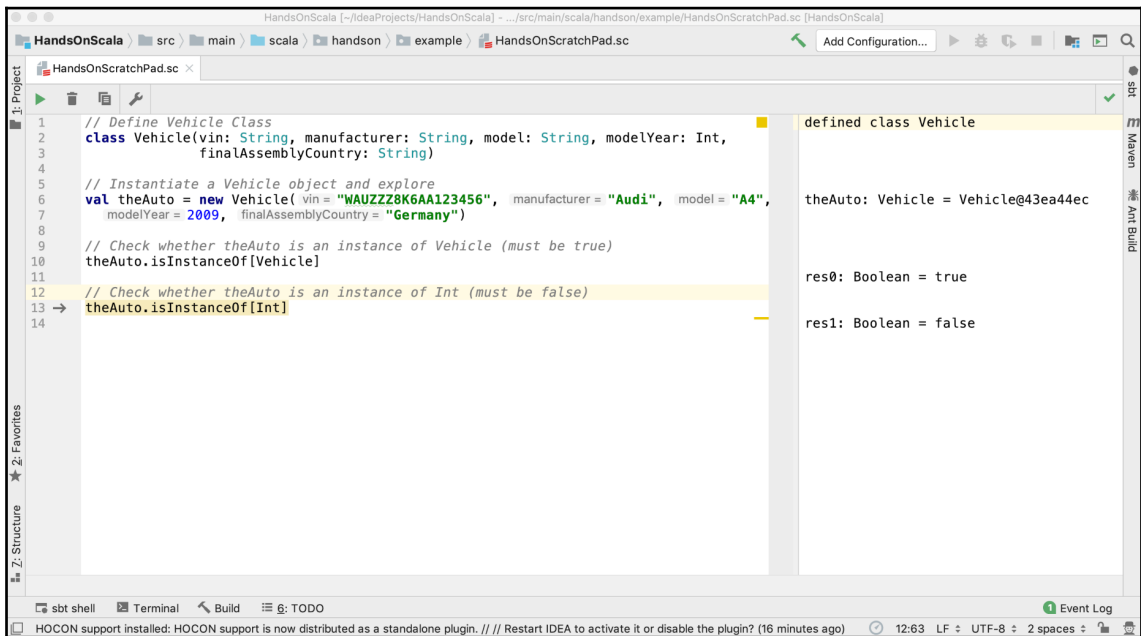
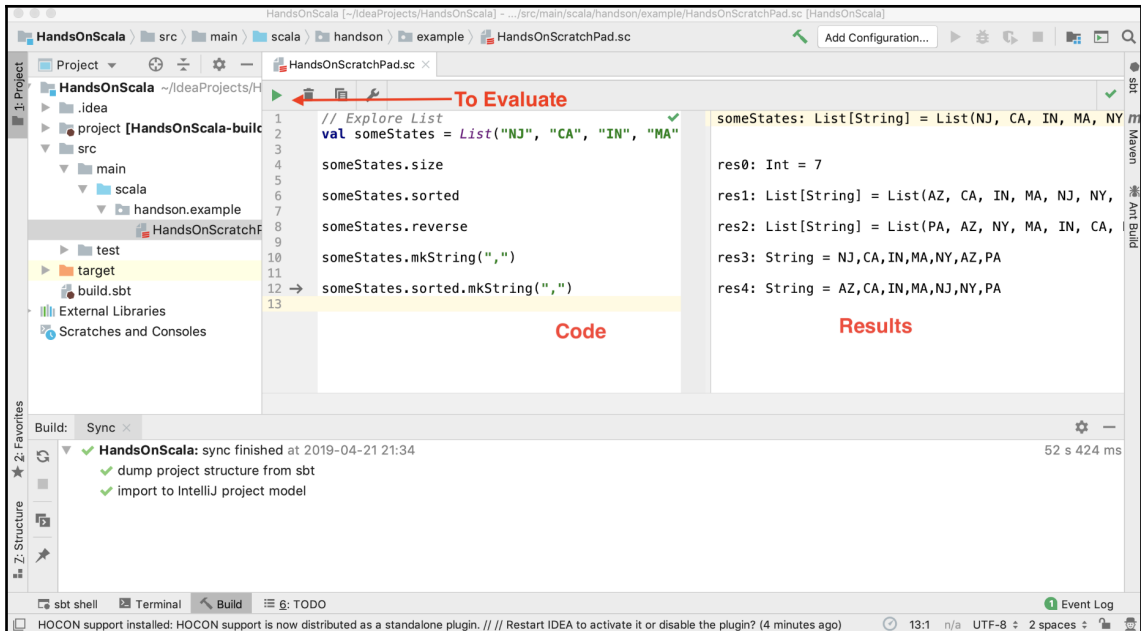
**IDE Features Trainer**
Code tools
Learn basic shortcuts and essential features interactively – right inside the IDE. No need to read long...
🕒 Mar 27, 2019 ↓ 1.1M ☆ 5
[Install](#)

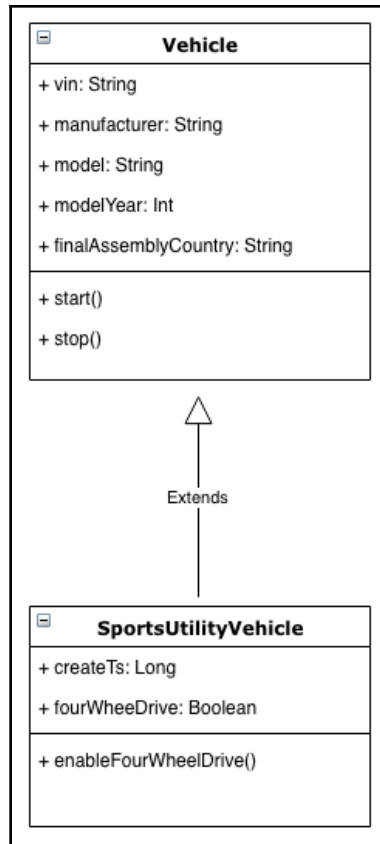
[Cancel](#) [OK](#)











Structure

```

package handson.example

class Vehicle(vin: String, manufacturer: String, model: String, modelYear: Int, finalAssemblyCountry: String) { // class
  private val createTs = System.currentTimeMillis() // example of encapsulation (hiding internals)
  def start(): Unit = { println("Starting...") } // behavior
  def stop(): Unit = { println("Stopping...") } // behavior
}

class SportsUtilityVehicle(vin: String, manufacturer: String, model: String, modelYear: Int, finalAssemblyCountry: String)
  extends Vehicle(vin, manufacturer, model, modelYear, finalAssemblyCountry) { // inheritance example
  def enableFourWheelDrive(): Unit = { if (fourWheelDrive) println("Enabling 4 wheel drive") }
  override def start(): Unit = {
    enableFourWheelDrive()
    println("Starting SUV...")
  }
}

object InheritanceExample {
  def main(args: Array[String]): Unit = {
    val theAuto = new Vehicle(vin = "MAUZZZ8K6AA123456", manufacturer = "Audi", model = "A4", modelYear = 2009, finalAssemblyCountry = "USA")
    theAuto.start()
    theAuto.stop()
    val anotherAuto: Vehicle = new SportsUtilityVehicle(vin = "MAUZZZ8K6A654321", manufacturer = "Audi", model = "Q7", modelYear = 2010, finalAssemblyCountry = "USA")
    anotherAuto.start() // polymorphism example
  }
}

```

Source Code

```

Run: InheritanceExample
/Library/Java/JavaVirtualMachines/jdk1.8.0_181.jdk/Contents/Home/bin/java ...
Starting...
Stopping...
Enabling 4 wheel drive
Starting SUV...
Process finished with exit code 0

```

Runtime

Structure

```

package handson.example

import scala.annotation.tailrec

object RecursionExample {
  def factorial(n: Int): Long = if (n <= 1) 1 else n * factorial(n-1)
  def optimizedFactorial(n: Int): Long = {
    @tailrec
    def factorialAcc(acc: Long, i: Int): Long = {
      if (i <= 1) acc else factorialAcc(acc * i, i - 1)
    }
    factorialAcc(acc = 1, n)
  }
  def main(args: Array[String]): Unit = {
    println(factorial(10))
    println(optimizedFactorial(10))
  }
}

```

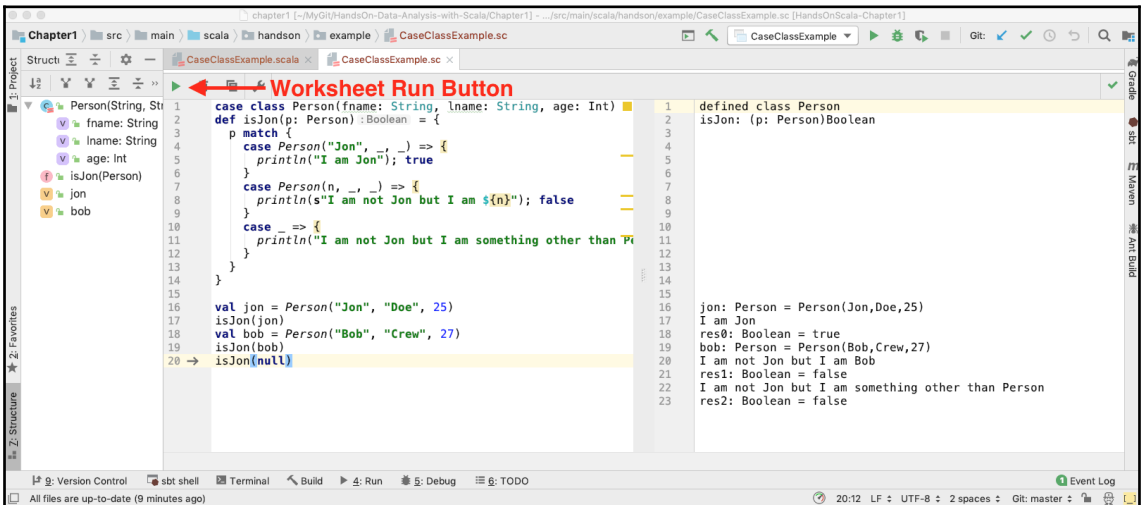
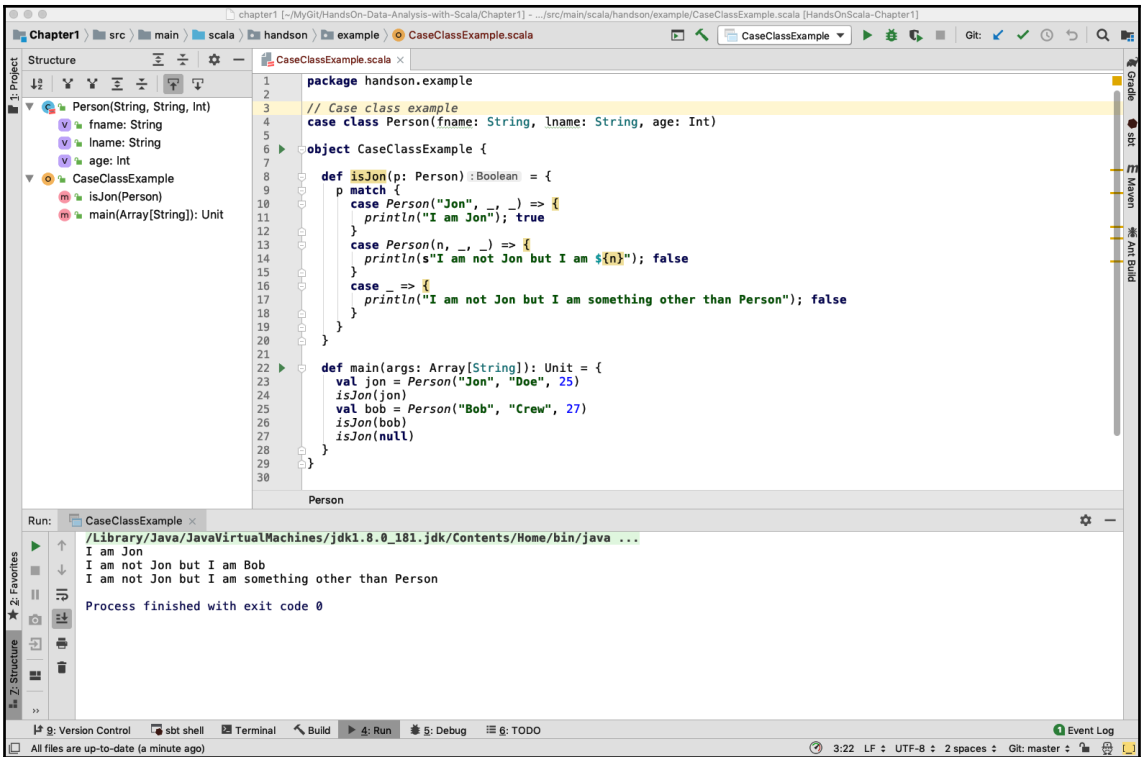
Recursive

Tail Recursive

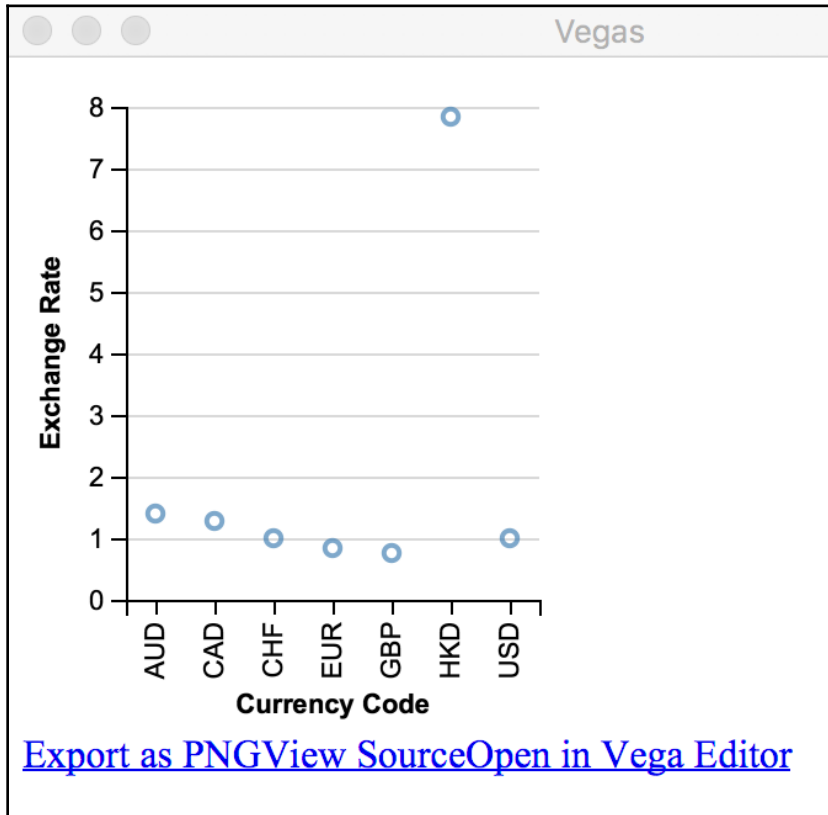
```

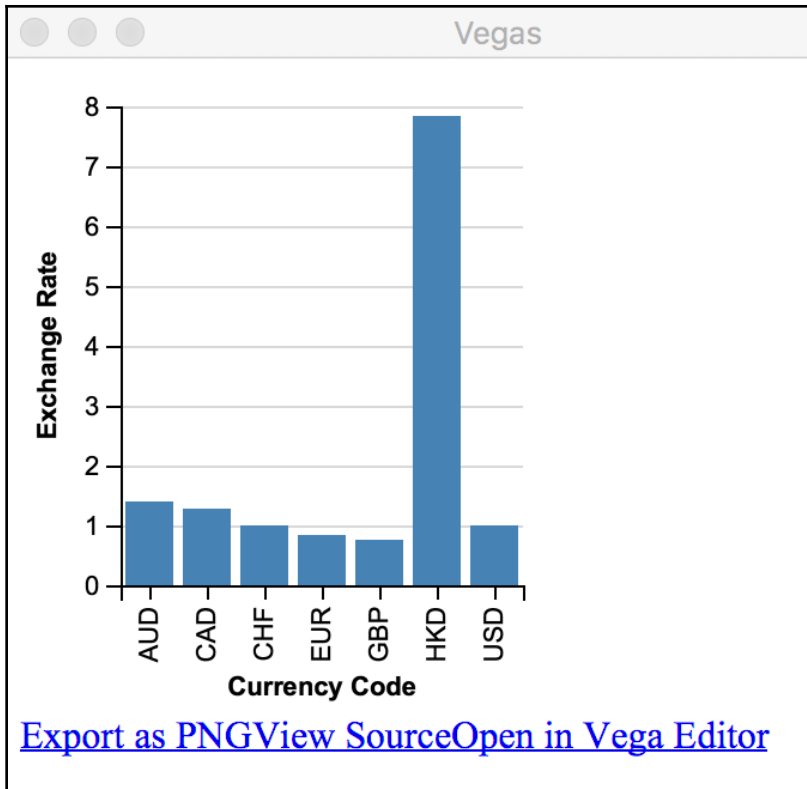
Run: RecursionExample
/Library/Java/JavaVirtualMachines/jdk1.8.0_181.jdk/Contents/Home/bin/java ...
3628800
3628800
Process finished with exit code 0

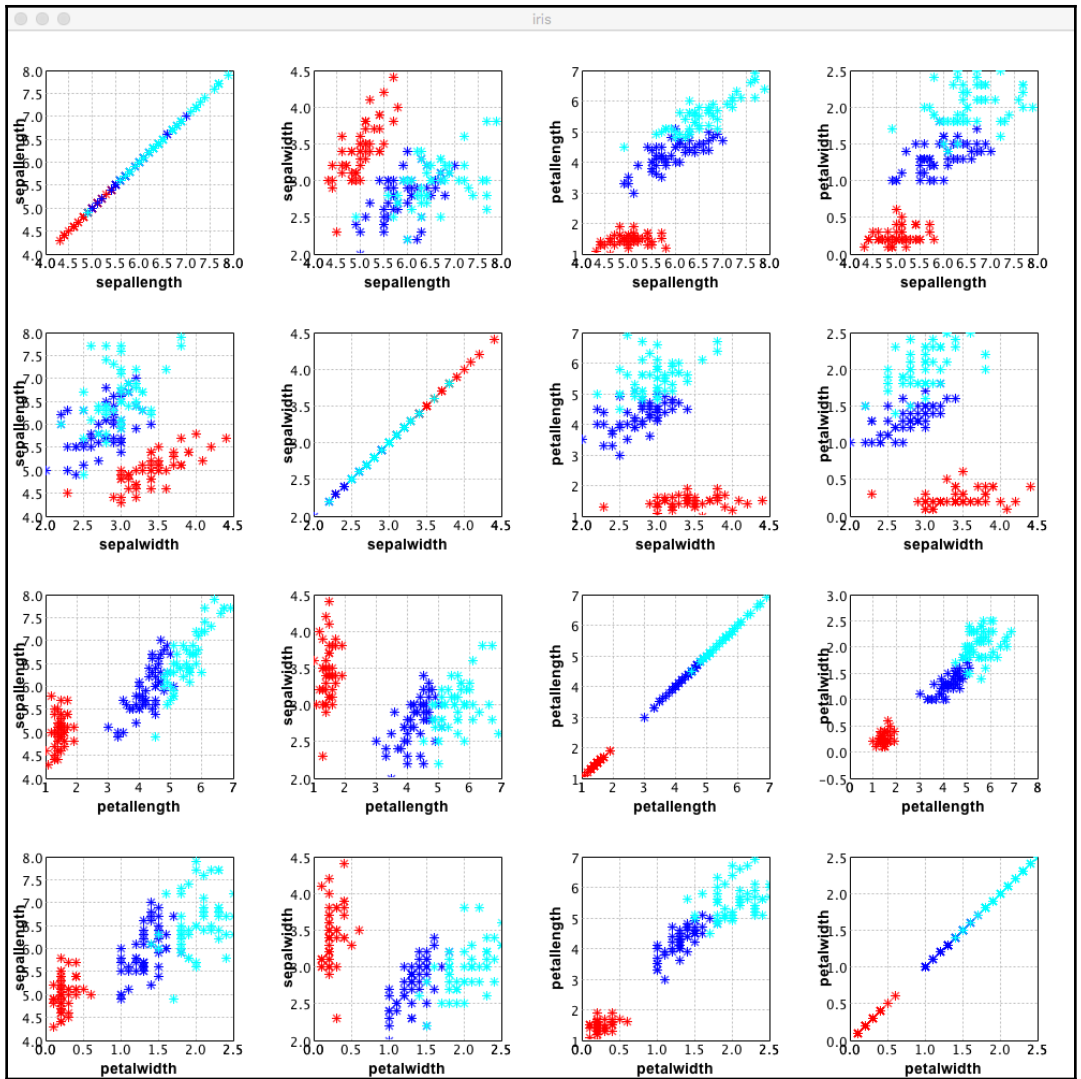
```

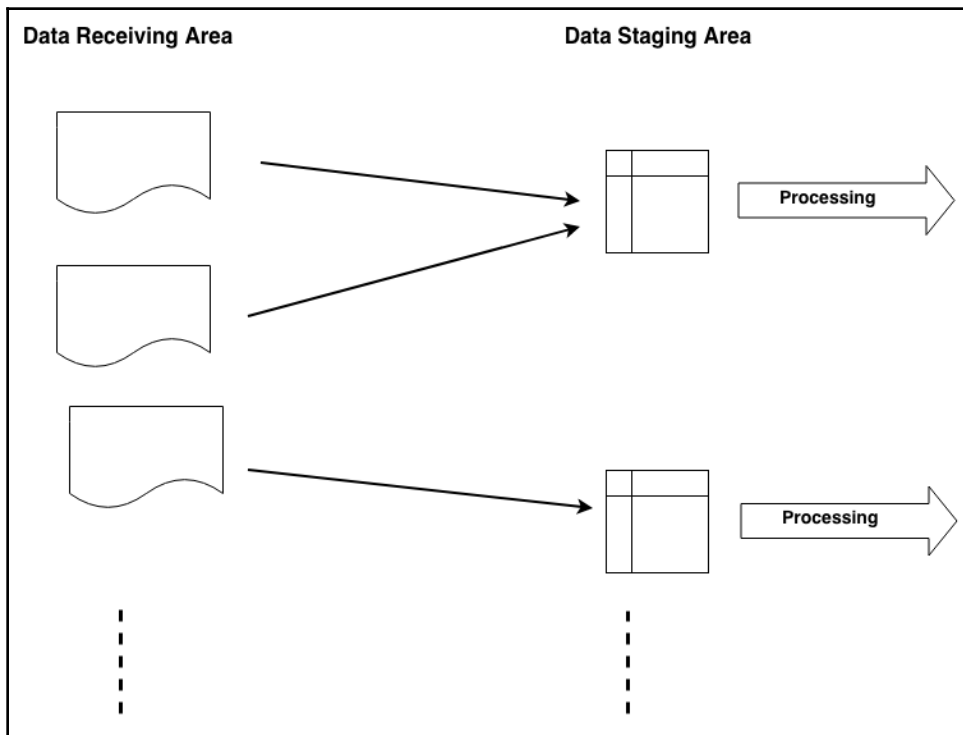


Chapter 2: Data Analysis Life Cycle

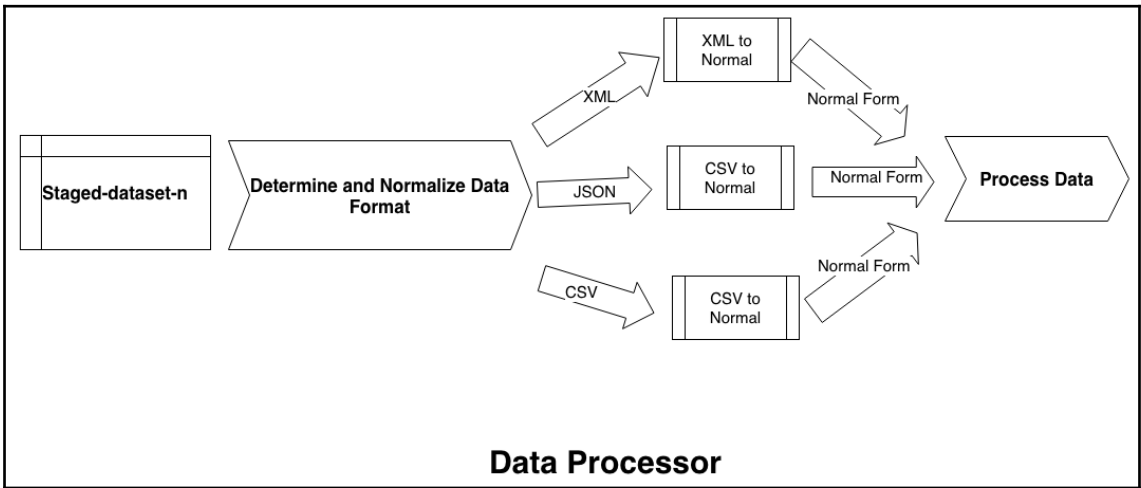



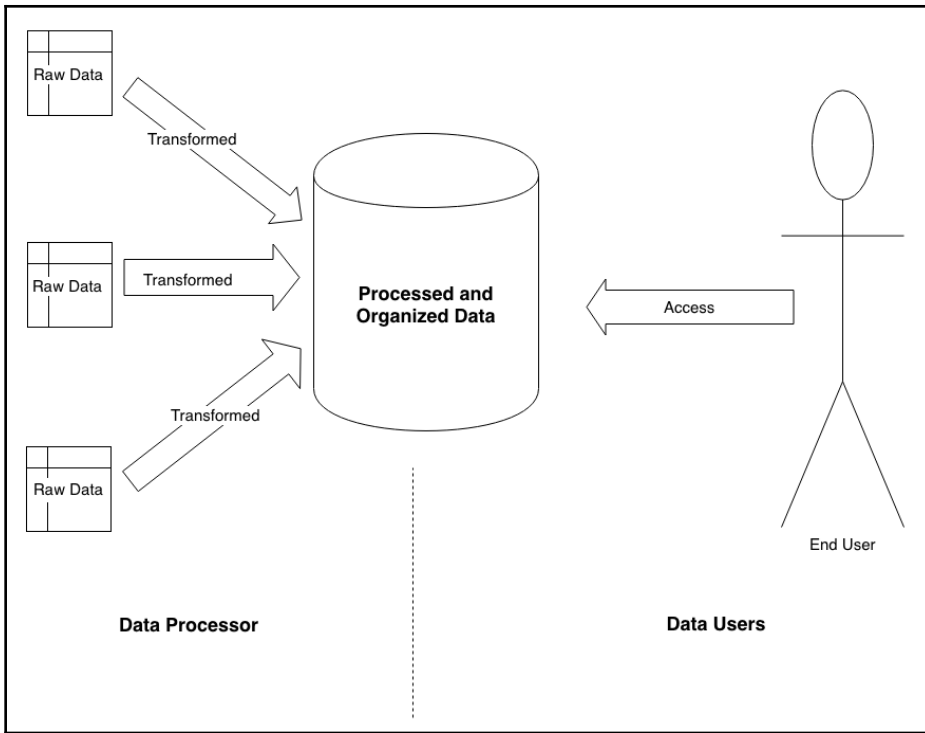




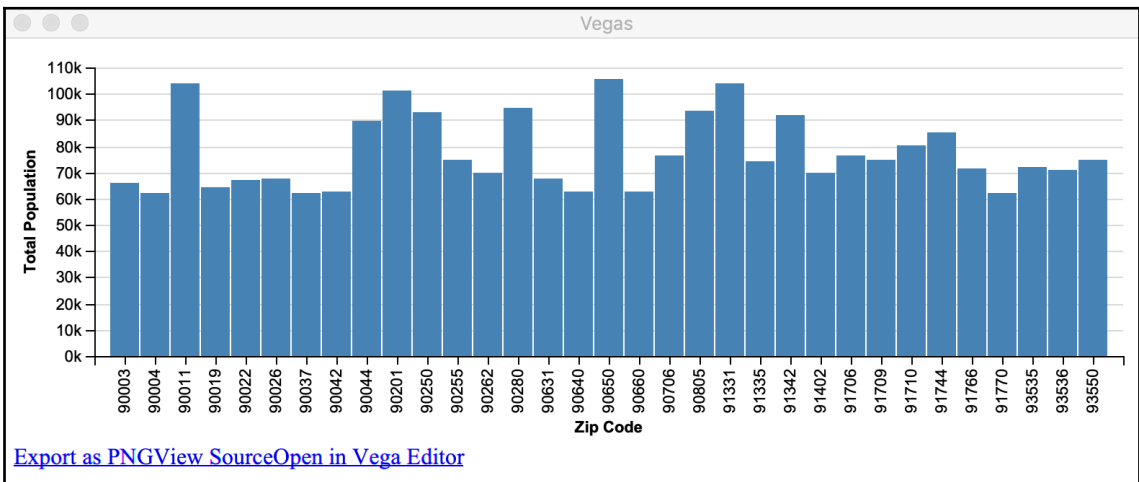
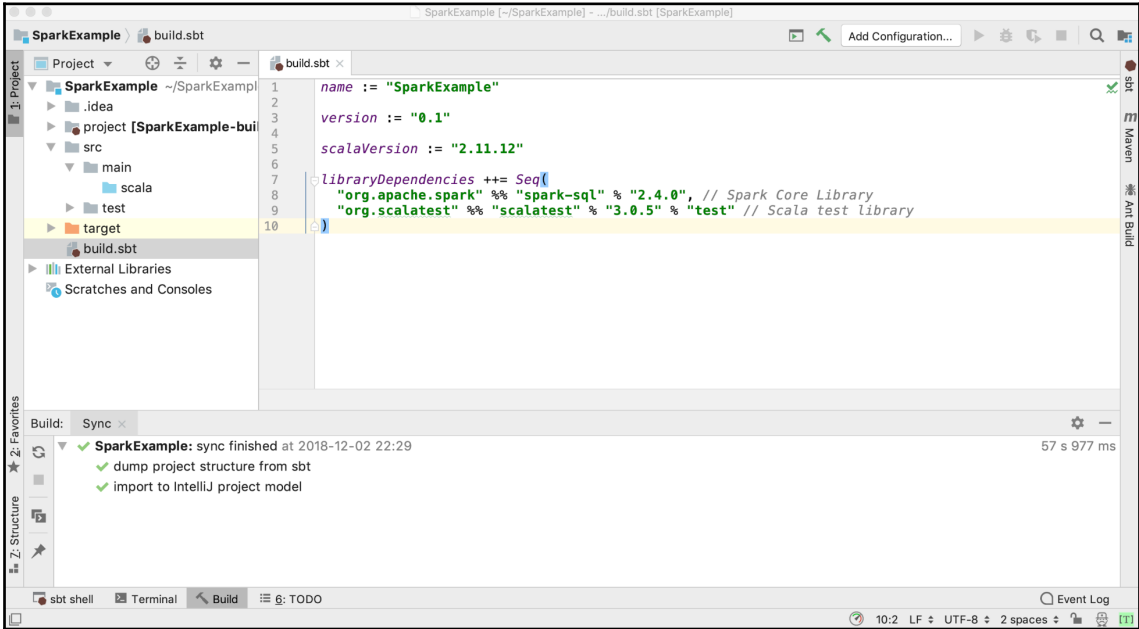
Chapter 3: Data Ingestion





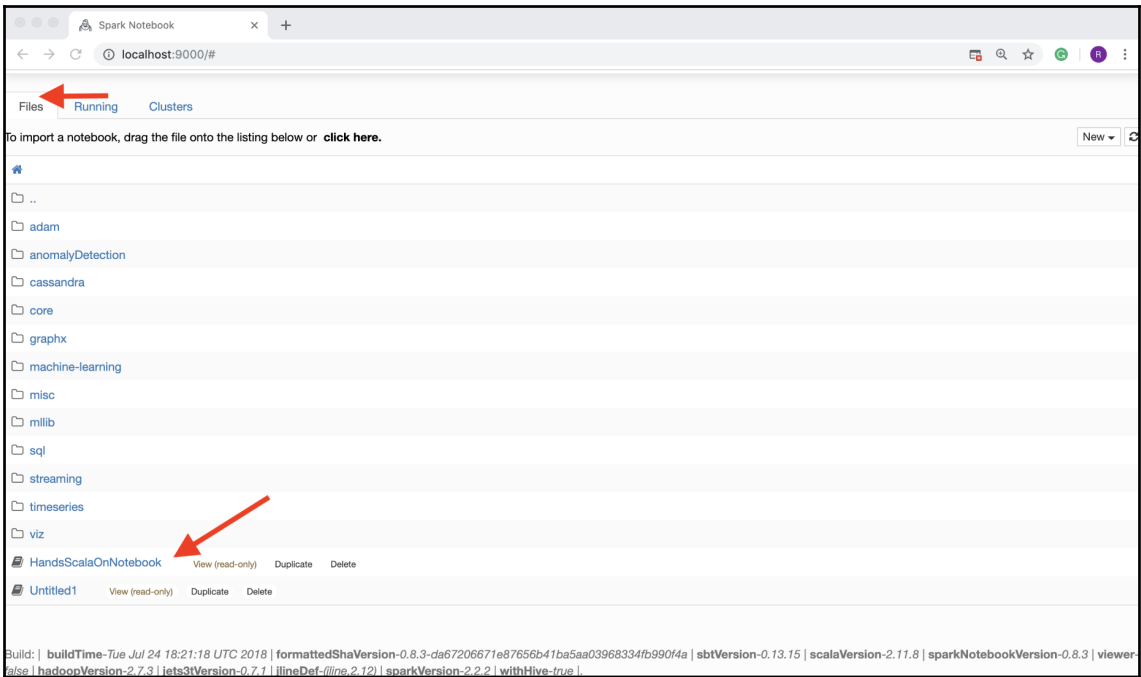
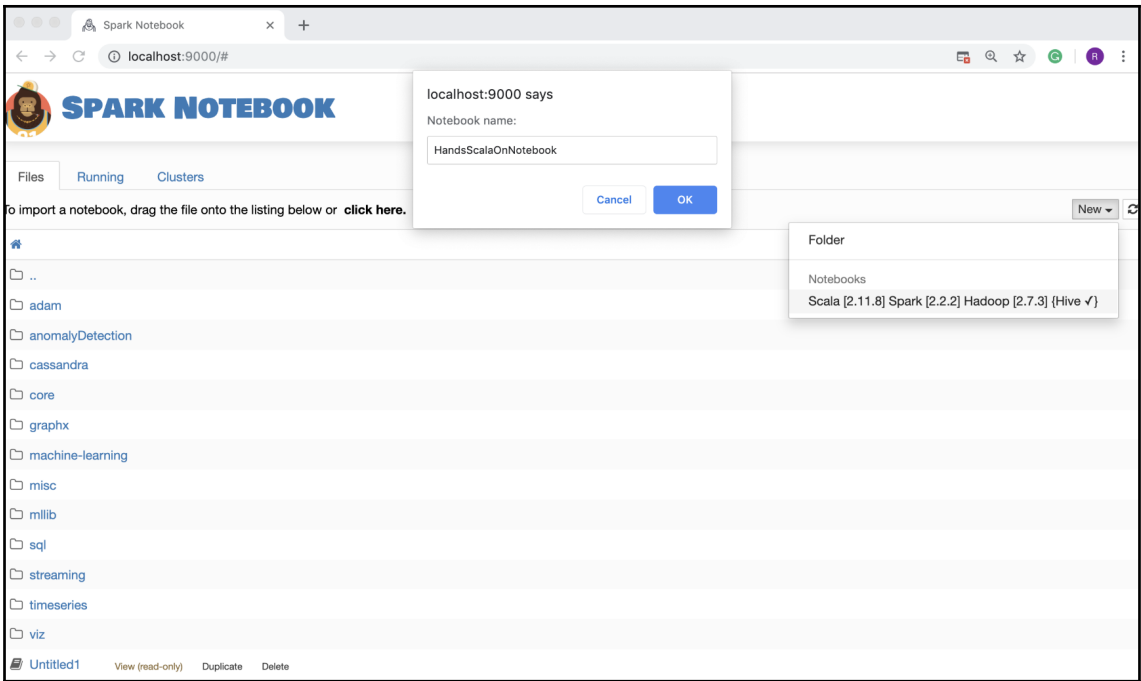


Chapter 4: Data Exploration and Visualization



tgz		deb		zip		docker	
Notebook 0.8.3		Notebook 0.7.0-pre2		Notebook 0.8.3		Notebook 0.8.3	
Scala 2.11	Scala 2.10	Scala 2.11	Scala 2.10	Scala 2.11	Scala 2.10	Scala 2.11	Scala 2.11
Spark 2.2.2 and Hadoop	Spark 2.2.2 and Hadoop	Spark 1.6.3 and Hadoop	Spark 1.6.0 and Hadoop	Spark 2.2.2 and Hadoop	Spark 2.2.2 and Hadoop	Spark 2.2.2 and Hadoop	Spark 2.2.2 and Hadoop
<ul style="list-style-type: none"> 2.7.2 parquet 	<ul style="list-style-type: none"> 2.7.2 parquet 	<ul style="list-style-type: none"> 2.7.2 	<ul style="list-style-type: none"> 2.6.0 hive parquet 	<ul style="list-style-type: none"> 2.6.0 hive parquet 2.7.2 hive parquet 2.7.2 parquet 2.7.3 hive parquet 	<ul style="list-style-type: none"> 1.0.3 parquet 2.7.2 parquet 	<ul style="list-style-type: none"> 2.6.0 parquet 2.6.0-cdh5.10.1 hive parquet 2.6.0 hive parquet 2.7.2 parquet 	<ul style="list-style-type: none"> 2.6.0 parquet 2.6.0-cdh5.10.1 hive parquet 2.6.0 hive parquet 2.7.2 parquet

The screenshot shows a web browser window with the URL `localhost:9000`. The page title is "SPARK NOTEBOOK". Below the title, there are tabs for "Files", "Running", and "Clusters". A message says "To import a notebook, drag the file onto the listing below or [click here](#)." Below this is a list of folders: `..`, `adam`, `anomalyDetection`, `cassandra`, `core`, `graphx`, `machine-learning`, `misc`, `mllib`, `sql`, `streaming`, `timeseries`, and `viz`. A "New" button is in the top right corner. A dropdown menu is open, showing "Folder", "Notebooks", and "Scala [2.11.8] Spark [2.2.2] Hadoop [2.7.3] {Hive ✓}" with a checkmark. At the bottom, there is a footer with build information: "Build: | buildTime-Tue Jul 24 18:21:18 UTC 2018 | formattedShaVersion-0.8.3-da67206671e87656b41ba5aa03968334fb990f4a | sbtVersion-0.13.15 | scalaVersion-2.11.8 | sparkNotebookVersion-0.8.3 | viewer-alse | hadoopVersion-2.7.3 | jets3tVersion-0.7.1 | jlineDef-((line,2,12)) | sparkVersion-2.2.2 | withHive=true |".



HandsScalaOnNotebook

localhost:9000/notebooks/HandsScalaOnNotebook.snb

SPARK NOTEBOOK HandsScalaOnNotebook (saved)

File Edit View Insert Cell Kernel Help

Scala [2.11.8] Spark [2.2.2] Hadoop [2.7.3] {Hive ✓}

Code and Output Area

Variable and Errors Area

Vars Errors

Terms defined

Name	Type

Build: | buildTime-Tue Jul 24 18:21:18 UTC 2018 | formattedShaVersion-0.8.3-da67206671e87656b41ba5aa03968334fb990f4a | sbtVersion-0.13.15 | scalaVersion-2.11.8 | sparkNotebookVersion-0.8.3 | viewer-false | hadoopVersion-2.7.3 | jets3tVersion-0.7.1 | jlineDef-(jline,2.12) | sparkVersion-2.2.2 | withHive=true |

HandsScalaOnNotebook

localhost:9000/notebooks/HandsScalaOnNotebook.snb#

SPARK NOTEBOOK HandsScalaOnNotebook (saved)

File Edit View Insert Cell Kernel Help

Scala [2.11.8] Spark [2.2.2] Hadoop [2.7.3] {Hive ✓}

```

Define a Scala Case Class
case class Person(name: String, age: Int)
defined class Person
Took: 2.770s, at 2019-04-23 22:31

CustomPlotlyChart
Seq(Person("James Bond", 50), Person("Jon Doe", 25), Person("Mickey Mouse", 18), Person("Foo", 33)),
dataOptions="type: 'bar'",
dataSources="x: 'name', y: 'age'"
res2: notebook.front.widgets.charts.CustomPlotlyChart[Seq[Person]] = <CustomPlotlyChart widget>
4 entries total
Took: 2.444s, at 2019-04-23 22:32

```

name as x-axis

age as y-axis

bar

Name	Age
James Bond	50
Jon Doe	25
Mickey Mouse	18
Foo	33

Terms defined

Name	Type
Person	type
res2	notebook.front.widgets.charts.CustomPlotlyChart

HandsScalaOnNotebook x +

localhost:9000/notebooks/HandsScalaOnNotebook.snb#

SPARK NOTEBOOK HandsScalaOnNotebook (unsaved changes)

File Edit View Insert Cell Kernel Help

Scala [2.11.8] Spark [2.2.2] Hadoop [2.7.3] (Hive ✓)

```

case class Person(name: String, age: Int)
defined class Person
Took: 2.145s, at 2019-04-23 23:04

val persons = Seq(Person("James Bond", 50), Person("Jon Doe", 25), Person("Mickey Mouse", 18), Person("Foo", 33)).toDF
persons: org.apache.spark.sql.DataFrame = [name: string, age: int] Spark DataFrame
Took: 6.915s, at 2019-04-23 23:04

CustomPlotlyChart(
  persons,
  dataOptions="(type: 'bar')",
  dataSources="(x: 'name', y: 'age')")
res3: notebook.front.widgets.charts.CustomPlotlyChart[org.apache.spark.sql.DataFrame] = <CustomPlotlyChart widget>
4 entries total

```

Name	Age
James Bond	50
Jon Doe	25
Mickey Mouse	18
Foo	33

Vars Errors

Terms defined

Name	Type
Person	type
persons	org.apache.spark.sql.Dataset
res3	notebook.front.widgets.charts.CustomPlotlyChart

HandsScalaOnNotebook x +

localhost:9000/notebooks/HandsScalaOnNotebook.snb#

File Edit View Insert Cell Kernel Help

Scala [2.11.8] Spark [2.2.2] Hadoop [2.7.3] (Hive ✓)

```

df3: org.apache.spark.sql.DataFrame = [num: int, experiment: string]
Took: 1.651s, at 2019-04-24 00:12

val df = df1.union(df2).union(df3)
df: org.apache.spark.sql.Dataset[org.apache.spark.sql.Row] = [num: int, experiment: string]
Took: 1.812s, at 2019-04-24 00:12

CustomPlotlyChart(df, dataOptions="(type: 'box', splitBy: 'experiment')", dataSources="(y: 'num')")
res5: notebook.front.widgets.charts.CustomPlotlyChart[org.apache.spark.sql.Dataset[org.apache.spark.sql.Row]] = <CustomPlotlyChart widget>
60 entries total

```

Vars Errors

Terms defined

Name	Type
df1	org.apache.spark.sql.Dataset
df2	org.apache.spark.sql.Dataset
df3	org.apache.spark.sql.Dataset
df	org.apache.spark.sql.Dataset
res5	notebook.front.widgets.charts.CustomPlotlyChart

HandsScalaOnNotebook

localhost:9000/notebooks/HandsScalaOnNotebook.snb#

SPARK NOTEBOOK HandsScalaOnNotebook (saved)

File Edit View Insert Cell Kernel Help

Scala [2.11.8] Spark [2.2.2] Hadoop [2.7.3] (Hive ✓)

```
val df = sparkSession.range(1000).map(n => (n, scala.util.Random.nextGaussian()))
df: org.apache.spark.sql.DataFrame = [num: bigint, value: double]
```

Took: 5.878s, at 2019-04-24 01:06

```
CustomPlotlyChart(df, dataOptions="{type: 'histogram', opacity: 0.7}", dataSources="{x: 'value'}")
```

res2: notebook.front.widgets.charts.CustomPlotlyChart[org.apache.spark.sql.DataFrame] = <CustomPlotlyChart widget>

1000 or more entries total (Warning: showing only first 1000 rows)

Took: 4.829s, at 2019-04-24 01:07

Vars 2 Errors

Terms defined

Name	Type
df	org.apache.spark.sql.Dataset
res2	notebook.front.widgets.charts.CustomPlotlyChart

HandsScalaOnNotebook

localhost:9000/notebooks/HandsScalaOnNotebook.snb#

SPARK NOTEBOOK HandsScalaOnNotebook (saved)

File Edit View Insert Cell Kernel Help

Scala [2.11.8] Spark [2.2.2] Hadoop [2.7.3] (Hive ✓)

```
val df = sparkSession.range(1000).map(n => (n, scala.util.Random.nextGaussian()))
df: org.apache.spark.sql.DataFrame = [num: bigint, value: double]
```

Took: 5.878s, at 2019-04-24 01:06

```
CustomPlotlyChart(df, dataOptions="{type: 'histogram', opacity: 0.7}", dataSources="{y: 'value'}")
```

res4: notebook.front.widgets.charts.CustomPlotlyChart[org.apache.spark.sql.DataFrame] = <CustomPlotlyChart widget>

1000 or more entries total (Warning: showing only first 1000 rows)

Took: 2.362s, at 2019-04-24 01:14

Vars 3 Errors

Terms defined

Name	Type
df	org.apache.spark.sql.Dataset
res2	notebook.front.widgets.charts.CustomPlotlyChart
res4	notebook.front.widgets.charts.CustomPlotlyChart

HandsScalaOnNotebook

localhost:9000/notebooks/HandsScalaOnNotebook.snb#

```
val df = Seq((1, 10, 30, "blue"), (2, 11, 60, "green"), (3, 12, 90, "red")).toDF("x", "y", "impact", "color")
df: org.apache.spark.sql.DataFrame = [x: int, y: int ... 2 more fields]


CustomPlotlyChart(df, layout="(title: 'Impact', showlegend: false, height: 600, width: 600)",
dataOptions="(mode: 'markers', dataSources="(x: 'x', y: 'y', marker: (size: 'impact', color: 'color'))")
res2: notebook.front.widgets.charts.CustomPlotlyChart[org.apache.spark.sql.DataFrame] = <CustomPlotlyChart widget>
3 entries total
```

Name	Type
df	org.apache.spark.sql.Dataset
res2	notebook.front.widgets.charts.CustomPlotlyChart

Chapter 5: Applying Statistics and Hypothesis Testing

```
numsSorted: scala.collection.immutable.IndexedSeq[Int] = Vector(1, 2, 3, 4, 5, 6, 7, 8, 9, 10,
11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34,
35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 50, 51, 52, 53, 54, 55, 56, 57, 58,
59, 60, 61, 62, 63, 64, 65, 66, 67, 68, 69, 70, 71, 72, 73, 74, 75, 75, 76, 77, 78, 79, 80, 81, 82,
83, 84, 85, 86, 87, 88, 89, 90, 91, 92, 93, 94, 95, 96, 97, 98, 99, 100)
```

Chapter 6: Introduction to Spark for Distributed Data Analysis



Lightning-fast unified analytics engine

Download Libraries Documentation Examples Community Developers Apache Software Foundation

Download Apache Spark™

1. Choose a Spark release: ←
2. Choose a package type: ←
3. Download Spark: [spark-2.4.0-bin-hadoop2.7.tgz](#) ←
4. Verify this release using the [2.4.0 signatures and checksums](#) and [project release KEYS](#).

Link with Spark

Spark artifacts are hosted in [Maven Central](#). You can add a Maven dependency with the following coordinates:

```
groupId: org.apache.spark
artifactId: spark-core_2.11
version: 2.4.0
```

Installing with PyPi

[PySpark](#) is now available in [pypi](#). To install just run `pip install pyspark`.

Release Notes for Stable Releases

- [Spark 2.4.0](#) (Nov 02 2018)
- [Spark 2.3.2](#) (Sep 24 2018)
- [Spark 2.2.3](#) (Jan 11 2019)

Archived Releases

As new Spark releases come out for each development stream, previous ones will be archived, but they are still available at [Spark release archives](#).

APACHECON North America Sept. 9-12 Las Vegas, Nevada 2019

Download Spark

Built-in Libraries:
[SQL and DataFrames](#)
[Spark Streaming](#)
[MLlib \(machine learning\)](#)
[GraphX \(graph\)](#)

Third-Party Projects

Latest News
Spark 2.2.3 released (Jan 11, 2019)
Spark+AI Summit (April 23-25th, 2019, San Francisco) agenda posted (Dec 19, 2018)
Spark 2.4.0 released (Nov 02, 2018)
Spark 2.3.2 released (Sep 24, 2018)
[Archive](#)

Apache Spark, Spark, Apache, the Apache feather logo, and the Apache Spark project logo are either registered trademarks or trademarks of The Apache Software Foundation in the United States and other countries. See guidance on use of Apache Spark [trademarks](#). All other marks mentioned may be trademarks or registered trademarks of their respective owners. Copyright © 2018 The Apache Software Foundation, Licensed under the [Apache License, Version 2.0](#).



Google Custom

- [The Apache Way](#)
- [Contribute](#)
- [ASF Sponsors](#)

We suggest the following mirror site for your download:

<http://mirror.cogentco.com/pub/apache/spark/spark-2.4.0/spark-2.4.0-bin-hadoop2.7.tgz> 

Other mirror sites are suggested below.

It is essential that you [verify the integrity](#) of the downloaded file using the PGP signature (.asc file) or a hash (.md5 or .sha* file).

Please only use the backup mirrors to download KEYS, PGP and MD5 sigs/hashes or if no other mirrors are working.

HTTP

<http://apache.claz.org/spark/spark-2.4.0/spark-2.4.0-bin-hadoop2.7.tgz>

<http://apache.cs.utah.edu/spark/spark-2.4.0/spark-2.4.0-bin-hadoop2.7.tgz>

<http://apache.mirrors.hoobly.com/spark/spark-2.4.0/spark-2.4.0-bin-hadoop2.7.tgz>


<http://apache.mirrors.ionfish.org/spark/spark-2.4.0/spark-2.4.0-bin-hadoop2.7.tgz>

<http://apache.mirrors.lucidnetworks.net/spark/spark-2.4.0/spark-2.4.0-bin-hadoop2.7.tgz>

<http://apache.mirrors.pair.com/spark/spark-2.4.0/spark-2.4.0-bin-hadoop2.7.tgz>

<http://apache.mirrors.tds.net/spark/spark-2.4.0/spark-2.4.0-bin-hadoop2.7.tgz>

Spark 2.4.0 Jobs Stages Storage Environment Executors Spark shell application UI

Spark Jobs (1) 

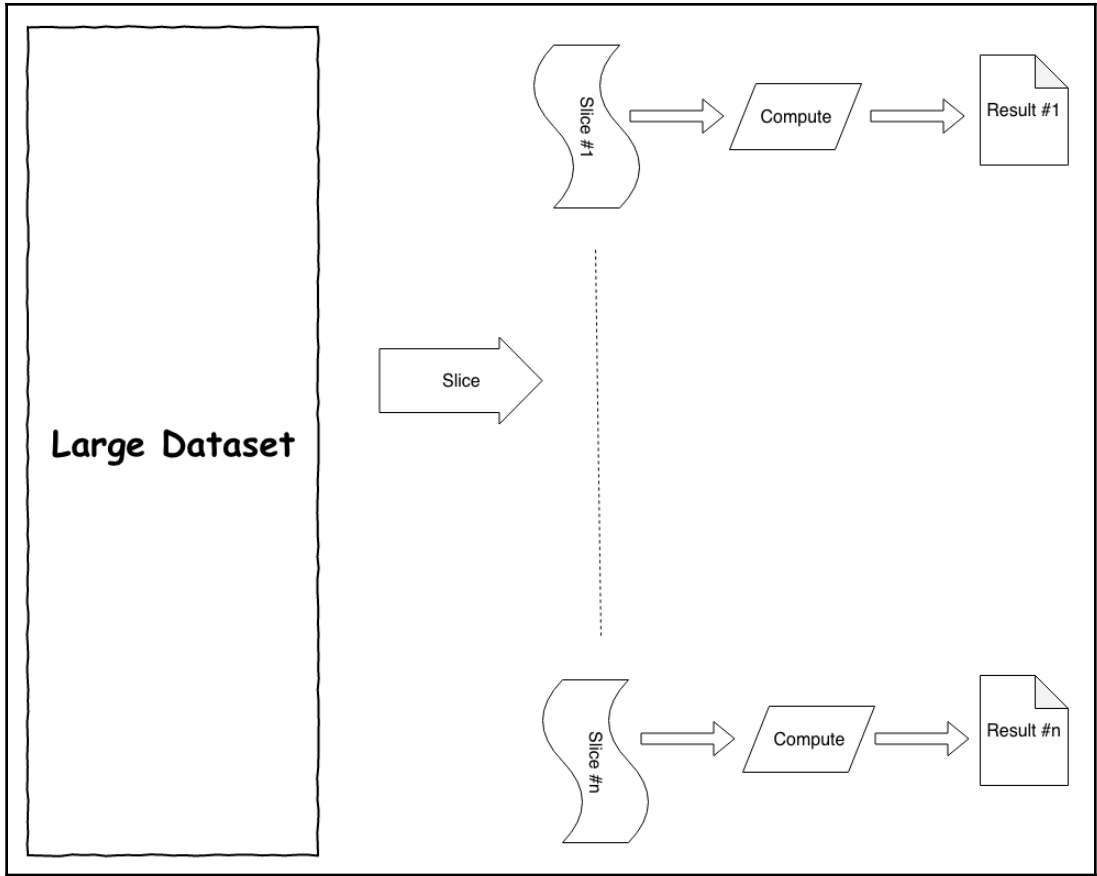
User: rasehgata
Total Uptime: 2.3 min
Scheduling Mode: FIFO

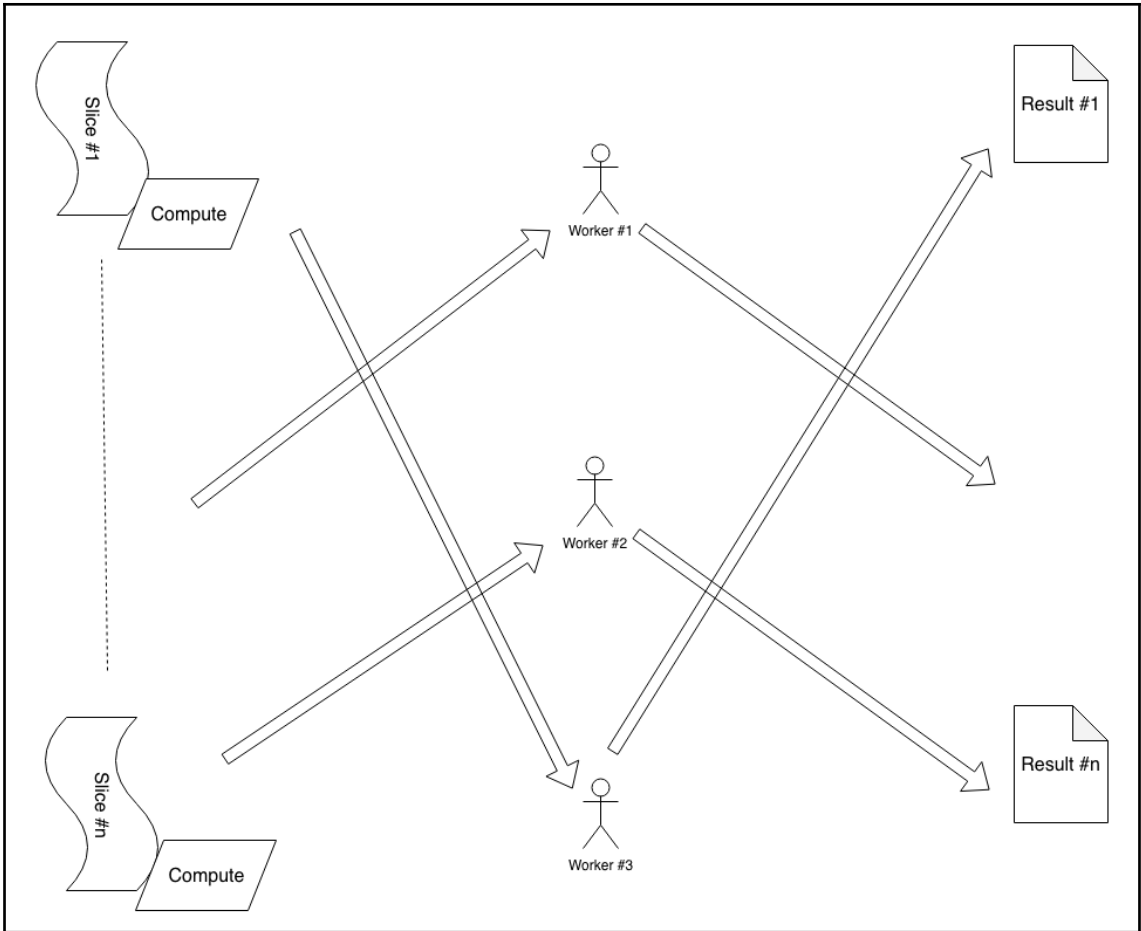
Event Timeline
 Enable zooming

Executors
Added
Removed
Executor drive added

Jobs
Succeeded
Failed
Running

January 2019 Tue 29 Wed 30 Thu 31 Feb 1 February 2019 Sat 2 Sun 3 Mon 4





spark 2.4.0 Jobs Stages Storage Environment Executors SQL Spark shell application UI

Spark Jobs (?)

User: rajeshgupta
 Total Uptime: 2.2 min
 Scheduling Mode: FIFO
 Completed Jobs: 2

Event Timeline

Completed Jobs (2)

numRDD.reduce(a, b) => if (a >= b) a else b

Job Id	Description	Submitted	Duration	Stages: Succeeded/Total	Tasks (for all stages): Succeeded/Total
1	reduce at <console>:26 reduce at <console>:26	2019/02/03 10:06:09	0.1 s	1/1	8/8
0	foreachPartition at <console>:26 foreachPartition at <console>:26	2019/02/03 10:05:59	0.4 s	1/1	8/8

numRDD.foreachPartition(p => println(p.size))

spark 2.4.0
Jobs Stages Storage Environment Executors SQL
Spark shell application UI

Details for Job 1

Status: SUCCEEDED
Completed Stages: 1

Event Timeline

Executors

- Added
- Removed

Stages

- Completed
- Failed
- Active

DAG Visualization

Completed Stages (1)

Stage Id	Description	Submitted	Duration	Tasks: Succeeded/Total	Input	Output	Shuffle Read	Shuffle Write
1	reduce at <console>:26	2019/02/03 10:06:09	0.1 s	8/8				

2.4.0

[Jobs](#)
[Stages](#)
[Storage](#)
[Environment](#)
[Executors](#)
[SQL](#)

Spark shell application UI

Details for Stage 1 (Attempt 0)

Total Time Across All Tasks: 0.7 s
 Locality Level Summary: Process local: 8

[▶ DAG Visualization](#)
[▶ Show Additional Metrics](#)
[▶ Event Timeline](#)

Summary Metrics for 8 Completed Tasks

Metric	Min	25th percentile	Median	75th percentile	Max
Duration	80 ms	89 ms	95 ms	0.1 s	0.1 s
GC Time	0 ms	0 ms	0 ms	0 ms	0 ms

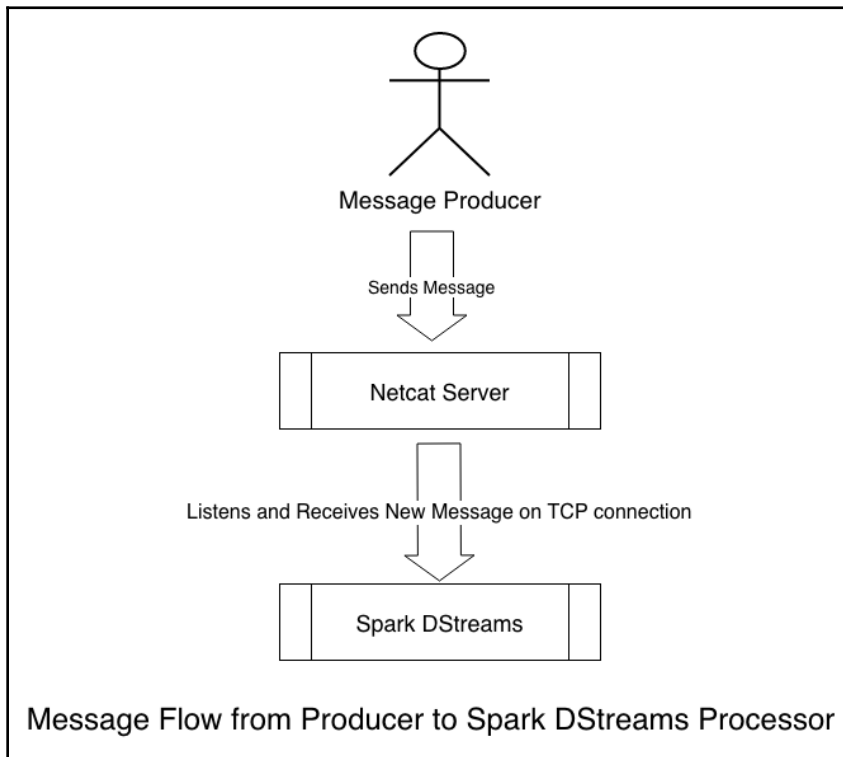
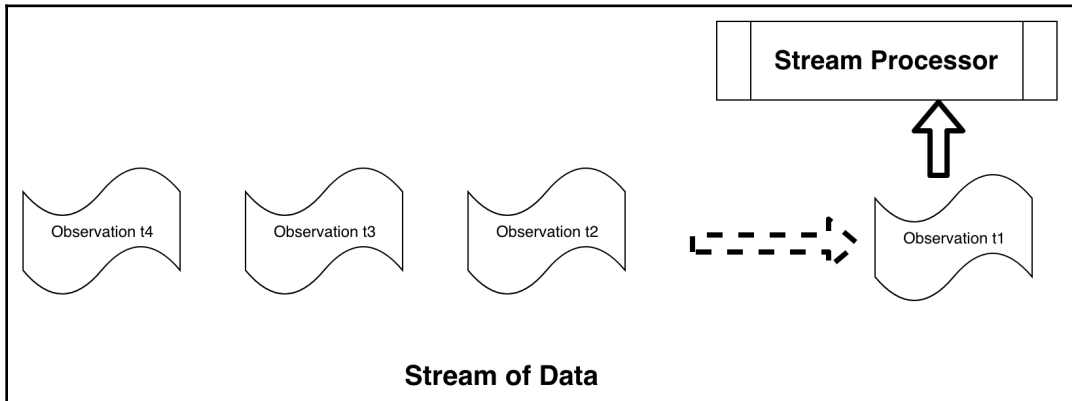
Aggregated Metrics by Executor


Executor ID	Address	Task Time	Total Tasks	Failed Tasks	Killed Tasks	Succeeded Tasks	Blacklisted
driver	192.168.1.31:62675	0.8 s	8	0	0	8	false

Tasks (8)

Index	ID	Attempt	Status	Locality Level	Executor ID	Host	Launch Time	Duration	GC Time	Errors
0	8	0	SUCCESS	PROCESS_LOCAL	driver	localhost	2019/02/03 10:06:09	0.1 s		
1	9	0	SUCCESS	PROCESS_LOCAL	driver	localhost	2019/02/03 10:06:09	0.1 s		
2	10	0	SUCCESS	PROCESS_LOCAL	driver	localhost	2019/02/03 10:06:09	0.1 s		
3	11	0	SUCCESS	PROCESS_LOCAL	driver	localhost	2019/02/03 10:06:09	0.1 s		
4	12	0	SUCCESS	PROCESS_LOCAL	driver	localhost	2019/02/03 10:06:09	97 ms		
5	13	0	SUCCESS	PROCESS_LOCAL	driver	localhost	2019/02/03 10:06:09	0.1 s		
6	14	0	SUCCESS	PROCESS_LOCAL	driver	localhost	2019/02/03 10:06:09	95 ms		
7	15	0	SUCCESS	PROCESS_LOCAL	driver	localhost	2019/02/03 10:06:09	0.1 s		

Chapter 8: Near Real-Time Data Analysis Using Streaming





Jobs
Stages
Storage
Environment
Executors
Streaming

Spark Jobs ^(?)

User: rajeshgupta
 Total Uptime: 2.8 min
 Scheduling Mode: FIFO
 Active Jobs: 1
 Completed Jobs: 52

[Event Timeline](#)

Active Jobs (1)


Job Id	Description	Submitted	Duration	Stages: Succeeded/Total	Tasks (for
0	Streaming job running receiver 0 start at <console>:32	2019/04/07 08:27:13 <small>(kill)</small>	1.8 min	0/1	

Completed Jobs (52)

Job Id	Description	Submitted	Duration	Stages: Succeeded/Total
52	Streaming job from [output operation 0, batch time 08:28:55] print at <console>:32	2019/04/07 08:28:55	7 ms	1/1 (1 skipped)
51	Streaming job from [output operation 0, batch time 08:28:55] print at <console>:32	2019/04/07 08:28:55	8 ms	1/1 (1 skipped)
50	Streaming job from [output operation 0, batch time 08:28:50] print at <console>:32	2019/04/07 08:28:50	6 ms	1/1
49	Streaming job from [output operation 0, batch time 08:28:50] print at <console>:32	2019/04/07 08:28:50	29 ms	1/1 (2 skipped)
48	Streaming job from [output operation 0, batch time 08:28:50] print at <console>:32	2019/04/07 08:28:50	9 ms	1/1 (2 skipped)
47	Streaming job from [output operation 0, batch time 08:28:45] print at <console>:32	2019/04/07 08:28:45	8 ms	1/1 (1 skipped)
46	Streaming job from [output operation 0, batch time 08:28:45] print at <console>:32	2019/04/07 08:28:45	10 ms	1/1 (1 skipped)
45	Streaming job from [output operation 0, batch time 08:28:40] print at <console>:32	2019/04/07 08:28:40	5 ms	1/1
44	Streaming job from [output operation 0, batch time 08:28:40]	2019/04/07 08:28:40	30 ms	1/1 (2 skipped)

Streaming Statistics

Running batches of **5 seconds** for **3 minutes 26 seconds** since **2019/04/07 08:27:13** (41 completed batches, 6 records)



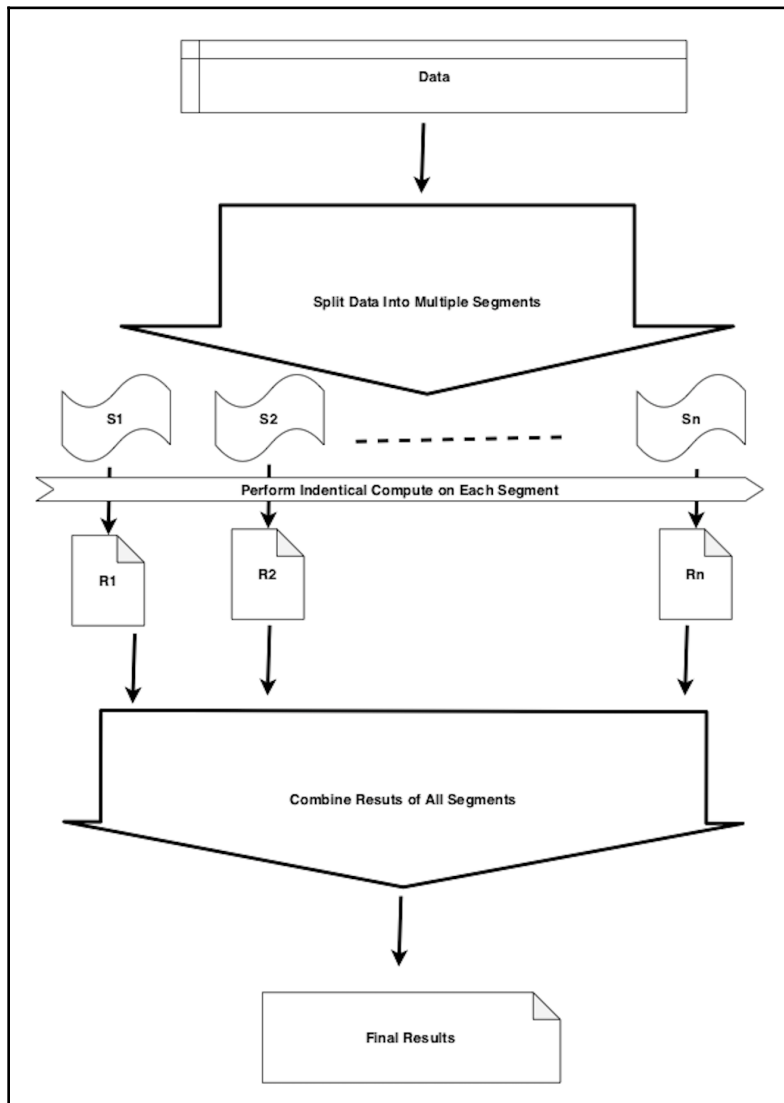
Active Batches (0)

Batch Time	Records	Scheduling Delay ^(?)	Processing Time ^(?)	Output
------------	---------	---------------------------------	--------------------------------	--------

Completed Batches (last 41 out of 41)

Batch Time	Records	Scheduling Delay ^(?)	Processing Time ^(?)	Total Del
2019/04/07 08:30:35	0 records	1 ms	23 ms	24 ms
2019/04/07 08:30:30	0 records	0 ms	0.1 s	0.1 s

Chapter 9: Working with Data at Scale



Spark shell - Details for Job 0

localhost:4040/jobs/job/?id=0

Spark 2.4.0

Status: SUCCEEDED
Completed Stages: 3

3 Stages

Completed Stages (3)

Stage Id	Description	Submitted	Duration	Tasks: Succeeded/Total	Input	Output	Shuffle Read	Shuffle Write
2	collect at <console>-25	2019/04/14 08:34:05	38 ms	2/2			200.0 B	
1	rdd at <console>-25	2019/04/14 08:34:05	0.3 s	2/2			560.0 B	200.0 B
0	rdd at <console>-25	2019/04/14 08:34:04	0.2 s	4/4				560.0 B

Spark shell - Details for Job 0

localhost:4040/jobs/job/?id=0

Spark 2.4.0

Status: SUCCEEDED
Completed Stages: 3

3 Stages

Completed Stages (3)

Stage Id	Description	Submitted	Duration	Tasks: Succeeded/Total	Input	Output	Shuffle Read	Shuffle Write
2	collect at <console>-25	2019/04/14 09:29:31	62 ms	10/10			2.3 KB	
1	rdd at <console>-25	2019/04/14 09:29:30	0.5 s	10/10			4.9 MB	2.3 KB
0	rdd at <console>-25	2019/04/14 09:29:29	0.6 s	4/4				4.9 MB

