



Hochschule für Angewandte Wissenschaften Hamburg
Hamburg University of Applied Sciences

Bachelorarbeit

Jakob Joachim

Methodology for Debugging Flink Applications

*Fakultät Technik und Informatik
Studiendepartment Informatik*

*Faculty of Engineering and Computer Science
Department of Computer Science*

Jakob Joachim

Methodology for Debugging Flink Applications

Bachelorarbeit eingereicht im Rahmen der Bachelorprüfung

im Studiengang Bachelor of Science Angewandte Informatik
am Department Informatik
der Fakultät Technik und Informatik
der Hochschule für Angewandte Wissenschaften Hamburg

Betreuender Prüfer: Prof. Dr. Olaf Zukunft
Zweitgutachter: Prof. Dr. Klaus-Peter Kossakowski

Eingereicht am: 2. February 2018

Jakob Joachim

Thema der Arbeit

Methodology for Debugging Flink Applications

Stichworte

Flink, Debugging, Stream Processing, Methodik zum Debuggen, Java Debuggen, Flink Debuggen

Kurzzusammenfassung

Das Debuggen von Programmen ist kompliziert, wird aber in einem verteilten System noch komplizierter. Je größer die täglich erzeugten Datenmengen werden, desto wichtiger wird die Nutzung verteilter Anwendung. Flink ist ein Framework zur verteilten Verarbeitung von Daten aus einem Stream, wobei die Daten zwischen den einzelnen Verarbeitungsschritten ebenfalls durch Streams transportiert werden. Flink erleichtert es dem Entwickler, verteilte Anwendung, die auf der Datenverarbeitung basieren zu entwickeln. Diese Thesis stellt eine Methodik zum Debuggen von eben diesen Flink Anwendungen dar.

Jakob Joachim

Title of the paper

Methodology for Debugging Flink Applications

Keywords

Flink, Debugging, Stream Processing, Methodology for Debugging, Debugging Java, Debugging Flink

Abstract

Debugging Applications is hard, debugging distributed applications is harder and the more massive the amount of data that gets generated every day gets, the higher the need for these distributed applications. Flink is a distributed stream processing framework; it processes data while only using streams. It makes it easy for developers to write these distributed applications to handle the ever-growing amounts of data in the world. This thesis explains how to debug these Flink applications by providing a debugging methodology.

Contents

1	Introduction	1
1.1	Objective of the Thesis	1
1.2	Structure of the Thesis	2
2	The Art of Debugging	3
2.1	Word Count Application	3
2.2	The Java Debugger	3
2.2.1	Basics of the Java Debugger	3
2.2.2	Debugging Principles	4
2.2.3	Conclusion	5
2.3	Why Programs Fail	6
2.3.1	Track the problem in the Database	6
2.3.2	Reproduce the failure	7
2.3.3	Automate and simplify the test case	8
2.3.4	Find possible infection origins	10
2.3.5	Focus on the most likely origins	10
2.3.6	Isolate the infection in the chain	11
2.3.7	Correct the defect	12
3	Related Work	13
3.1	Flink Framework	13
3.1.1	Basics	13
3.1.2	Stream Processing and Batch Processing	14
3.1.3	Debugging and Checkpoints	15
3.2	Similar work	17
3.2.1	BigDebug: Debugging Primitives for Interactive Big Data Processing in Spark	17
3.2.2	Debugging Distributed Systems	19
4	Debugging Flink Applications	21
4.1	Better Developing	21
4.1.1	Building Tasks	22
4.1.2	Metrics	24
4.1.3	Logging	24
4.2	The Debugging Process	25
4.2.1	Common Problems	25

4.2.2	Track the Problem in the database	26
4.2.3	Reproduce the failure	27
4.2.4	Automate and simplify the test case	31
4.2.5	Find possible infection origins	32
4.2.6	Isolate the infection in the chain	37
4.2.7	Correct the defect	39
5	Flink Backtracker	42
5.1	User manual	42
5.1.1	Debugging a Program with the Backtracker	43
5.2	Backend architecture	44
5.3	Frontend architecture	45
5.4	State of the program	46
6	Conclusion and Outlook	47
6.1	Lessons Learned	47
6.2	Outlook	48
6.3	Further Reading	48

Listings

2.1	WordCount Result	3
2.2	Word Count Example	9
4.1	FlinkSpector Example: FlinkSpector	31
4.2	Word Count Example Flink	33
4.3	Connection Exception	34
4.4	Exception in Transformation	35
4.5	Create Stacktrace	36
5.1	How to use the Backtracker	43
5.2	Save Watermark	45

List of Figures

3.1	Simple Dataflow in Flink Flink (c)	14
3.2	Distributed Dataflow in Flink Flink (b)	15
4.1	Debugging faulty data	28
5.1	Backtracking Example	44
5.2	Dataflow without the Backtracker	44
5.3	Dataflow with the Backtracker enabled	45

1 Introduction

Debugging software systems is a difficult task, to begin with. Debugging a distributed application makes this process even harder. Since the beginning of computer science, developers have always seen debugging as an unfortunate and tedious process. Always trying to minimise the time spent doing it by developing better programming styles and techniques. Unfortunately, not even the best programmer using the best possible method for his current project can write bug-free code all the time. A lot of programmers only learn how to debug by doing it; almost no one reads a book or scientific paper about how to improve one's debug technique. Barry W. Boehm estimates that reworking defects in requirements, design, and code consumes 40-50% of the total cost of software development [Boehm \(1987\)](#). It makes sense to learn how to properly debug as a better debugging understanding leads to less time spent debugging and more time developing new features. As well as improving productivity, good knowledge of debugging also increases one's awareness of potential issues while writing code. So learning how to debug correctly not only reduces the time spent debugging but also enhances the software quality written by the developer.

1.1 Objective of the Thesis

Flink is a stream processing framework designed to make it easier for developers to write distributed applications that have an continuous input (stream) of data. Even though applications for Flink are much easier to understand, write and debug it is still far from easy. This thesis tries to make it simpler for Flink developers to find bugs in their code. This is done by providing a methodology for debugging Flink and a debugging tool to the developer as well as some general recommendations for building Flink applications. When finished with this thesis the developer should have a good understanding as to why an error might occur even if the error message itself is not helpful.

1.2 Structure of the Thesis

In *chapter two*, debugging techniques are discussed and a methodology for debugging applications of all kinds is outlined

In *chapter three*, the Flink framework is explained, and similar work is analysed.

In *chapter four*, a debugging methodology for Flink is explained and the relevant steps shown

In *chapter five*, the Flink Backtracker tool is shown, and its internals discussed.

The final *chapter six* sums up the thesis, explains what lessons were learned and future work for debugging Flink applications is presented.

2 The Art of Debugging

As debugging of traditional programs is in many ways similar to debugging Flink applications it is necessary to explore what good debugging is. This chapter is an accumulation of techniques and methods that many people consider essential to debug effectively.

2.1 Word Count Application

This section quickly outlines the "WordCount Application" that is used in the thesis as an example. The Application counts how often each word is in a specific text and returns the results to the command line. An example run of the program would be:

```
1 $ java -jar wordCountSimple.jar "Hello hello how do you do"
2 how: 1
3 hello: 2
4 do: 2
5 you: 1
```

Listing 2.1: WordCount Result

Note that the application ignores capital letters in the input. This will be relevant later on.

2.2 The Java Debugger

Flink applications are very complex Java or Scala programs that still share a lot of the characteristics of typical Java programs. As such most methods of debugging Java still apply to Flink and should be used. This chapter will outline some of the ways that help Java developers to find their bugs faster and with less of a hassle.

2.2.1 Basics of the Java Debugger

To understand how the java debugger works we have to first look at how Java runs applications. Java applications run on the JVM, the Java virtual machine. To gain access to the information of the running program, the virtual machine needs to be reachable. Java provides interfaces for

these communications, in general, the IDE implements the Java Debug Interface (JDI), and the JVM Software implements the JVM Tool Interface (JVM TI). To allow communications between the two interfaces, a protocol is necessary as the two interfaces are not running on the same (virtual) system. Java uses the Java Debug Wire Protocol (JDWP) for that, which specifies which byte holds which information on the byte stream. When launching an application in the "debug" mode, the IDE also starts a JVM TI alongside the application in the JVM and provides the JVM TI with the breakpoint locations. These breakpoints can either be line numbers or method heads that only trigger when reaching that precise method. Once the program hits the particular method (or line), the JVM TI stops the application and notifies the JDI. In general, IDEs provide information about the state of relevant variables; this is done by the IDE itself as it has to send a request to the JVM TI for each variable the user might want to see.

2.2.2 Debugging Principles

There are a few simple principles that help not only to find errors in the code easier, but also to make sure that similar errors won't occur in the future. This section will lay out some of these principles.

Debugging Mindset

The mindset plays a significant role in debugging; often developers see debugging as an annoyance and try to get away from it as fast as possible. Although it is the goal of good debugging to minimise the time a developer spends on it, it should be seen as a way to learn something about why the error occurred in the first place. This learning reduces the time he has to spend searching for a similar mistake later on.

The Programmer is at Fault

Another important step is to understand that it is most probably the developer's fault when an error occurs. Thus the programmer should always first look into his code to find the bug, only if that fails he should start looking into other possible reasons. It is a common mistake not to recognise one's fault and should always be considered when debugging.

Last Change

Often the error occurs at a piece of code that is not responsible for the error itself. [Zeller \(2009\)](#) describes program errors as follows:

1. A programmer creates a defect in the code.
2. The defect causes an infection.
3. The infection spreads
4. The infection causes a failure

It is often best to first look at the last change that was made in the code, as the failure probably originated there. Obviously, that doesn't always have to be true, as there are multiple other reasons why the failure only just now got discovered.

2.2.3 Conclusion

It is important to note that this section was merely a refresher on debugging and is not enough to understand it in its completeness. For further information on the topic consider reading [Zeller \(2009\)](#). The book covers the subject of debugging much more in depth.

2.3 Why Programs Fail

This section will examine the "TRAFFIC" method of debugging that is published in Zeller (2009)'s book "Why Programs Fail". It consists of seven steps that lead to a quick removal of the failure while still preserving the information on how the failure came to be, as a big part of debugging is not to fix a failure but to make sure that the same or a similar failure won't happen again. The seven steps are:

1. Track the problem in the Database
2. Reproduce the failure
3. Automate and simplify the test case
4. Find possible infection origins
5. Focus on the most likely origins
6. Isolate the infection in the chain
7. Correct the defect

The following subsection explains each step of the method along the example program outlines above 2.1

2.3.1 Track the problem in the Database

Tracking is not the first step, but a good method at each step, logging what happened so that anybody involved knows how far each problem was investigated. It is started once someone finds a problem. Put simply tracking is holding the information on what the problem is and how close the developer is to fix it. This is mostly done on a platform that is accessible for all involved parties to make the communication between the user and the developer easier. This platform is only useful when used permanently as out of date information is more harmful than useful. Most software projects have multiple people working on them, not all of them know every bit of the program. This adds complexity to an already complicated process. It is necessary to split the work to different developers and this process has to be logged, otherwise problem reports might get lost. A developer might have 20 problem reports on his desk of which he only knows how to solve a few. Other reported problems might not even be problems, but a wanted state. For example, someone might report that a password field is showing only stars instead of the letters he put in. As this (for the developer) is a security feature, it will

not be changed and has to be logged or noted so that the same problem will hopefully not be reported again. To solve these and other issues, tracking should be used. Depending on the application these metrics are useful:

1. The State of the Problem - Is the problem new, assigned to a developer, resolved, closed, etc. This is useful for the developer as he can easily see which problems he has to work on and which are already solved. It is also beneficial for the user as he can easily see when his problem is resolved.
2. The Resolution - Is the problem fixed, invalid, won't be fixed (as the example above), a duplicate, etc. This is useful as it lets the user see to which conclusion the developer came.
3. Assigned Developer - Which developer is assigned to the problem. Makes it easy to communicate with the correct person and lets people know that the problem is worked on.
4. Severity - Is the problem crucial or is it only a minor inconvenience. Helps the developer prioritise which problems to solve first.

These are just the most important ones, depending on the project more should be added.

2.3.2 Reproduce the failure

The first real step in any debugging activity is to consistently reproduce the problem described in the problem report. This has two important reasons:

1. To observe the problem - The developer has to be able to reproduce the problem to fix it, the developer could also check the source code at a position he thinks could be responsible for the problem without reproducing it, but that makes it unnecessarily hard for the developer as a good problem report should be re-creatable.
2. To check whether the problem is fixed - It is incredibly hard to tell in most situations if a problem is solved or not without being able to rerun the problem without it happening.

Reproducing a problem can be incredibly hard, as the problem is rarely found by the person that has to fix it, but by an individual who doesn't understand how the program works. This makes it difficult for both the finder and the fixer of the problem, as the finder doesn't know which information the fixer might need, and for the fixer as he can only reproduce it with the corresponding information.

Reproducing is done by going through the following three steps until the problem can be reproduced:

1. Reproduce the problem locally - In the best case scenario, the problem can be replicated on the local machine of the developer fixing it with the information provided by the problem report. This is most successful when the problem state is not connected with many other choices in the program. For example, a button might not work (as reported in the problem report) that starts some function no matter what else happened before the button was pressed.
2. Adopt more circumstances of the problem environment - Sometimes the problem can not be reproduced by only following the steps provided by the problem report. In that case, it is necessary to check what else is known about the environment the problem was found in. This means installing the same version of depending software, using the same configuration file, using the same hardware or anything else that could influence the described problem.
3. Contact the problem finder or declare the problem invalid - If the reproduction failed even with all circumstances applied as outlined in the problem report, there are only two options left. Either the problem is not real, or the finder has not provided the necessary information. Depending on the bug this can be solved by contacting the finder and asking him for more details. It is beneficial for the developer to see the problem live, so it would be a good option to remotely gain access to the computer of the finder or get a video meeting with him. This is obviously not possible or feasible in many situations but extremely helpful if it can be arranged.

At this point, it is important to mark what information were necessary, so that another developer can reproduce the problem easier as well as saving the information as it is not guaranteed that the problem will be solved immediately. This is done with tracking: see [2.3.1](#)

2.3.3 Automate and simplify the test case

Once the problem is reproduced, it is desired to simplify the problem so that it can be replicated by a test case that as well as confirming that the problem is solved, helps avoid building a similar problem at a later time.

Simplification is done by understanding what the root of the problem is. When the problem occurs when pressing a particular button, but the button has no context to other set states of the program it is safe to assume that actions done before the button click can be omitted. Thus

leaving only the test case: After button press, function x started? This statement can now be easily modelled by a test case.

The simplification is not only essential for building test cases but also helpful for the developer when trying to locate the problem in the code as it is much easier to find the error when only looking at the relevant pieces of source code.

Why write tests? It is understandable to ask why to write a test when it is already known where the bug is originating from. The developer only has to fix the problem and manually test the program once, what good is a test case here? It is important to understand that finding a problem is giving a lot of valuable information to the developer he might not even realise. We can assume that if the problem got into the code once it might get into the code again. The only way to make sure it does not is to give the developer an immediate feedback when the problem occurs again. As well as helping the Developer in the future it can assist the developer while writing the fix to quickly test if the fix worked or not without having to manually start the application and reproduce the problem all over again.

```
1 import java.util.HashMap;
2 import java.util.Map;
3
4 public class WordCount {
5     public static void main(String[] args) {
6         Map<String, Integer> result = new HashMap<>();
7         String[] words = args[0].split("\\s+");
8         int numberOfWords = words.length;
9         System.out.println("There_are_" + numberOfWords
10            + "_words_in_total.");
11        for (String word : words) {
12            int count = result.getDefault(word, 1);
13            result.put(word.toLowerCase(), count + 1);
14        }
15        int uniqueWords = result.size();
16        System.out.println("Of_these_" + numberOfWords + "_words_"
17            + uniqueWords + "_are_unique.");
18        System.out.println("The_unique_words_are:_");
19        for (String word : result.keySet()) {
20            System.out.println(word + ":_ " + result.get(word));
21        }
22    }
```


Listing 2.2: Word Count Example

2.3.4 Find possible infection origins

Once the failure can be reproduced easily, it is time to search for what part of the code is responsible for the failure. This is often the hardest part of debugging. To make finding easier for the developer the traffic approach suggests to use backtracking to find the relevant piece of code. Backtracking is done by starting at the manifestation of the error, meaning if the program fails with an exception, the line noted in the exception. If the program just produces an unwanted or wrong value, the line is used where that value was returned. Once a starting point is found the next step is to backtrack all active and passive usages of the variable. As failures can quickly propagate through the application it is unfortunately not enough to only look at the variable that was causing the exception. We have to also look at all other variables that interact with the first variable and so on. Lines that don't use the variable can be omitted as they cannot be responsible for the failure. The information gained by this technique should be displayed by a control flow graph to make them easier to understand. The graph can often be omitted but should be used once the problem space gets too big to handle without a graph.

2.3.5 Focus on the most likely origins

Once the control flow graph is done, the developer should look at the most likely sources. This can be quite difficult as it is subjective to the developer what the most likely origins are. Some places to check can be:

1. Last Change in the Application - Often problems occur after a change in the software thus it is very likely that problem is a result of the modification. It is, of course, possible that the problem was there before but got propagated through the new piece of code. In that case, it still makes sense to look into the new code.
2. Check for common problems in the framework - when working with a framework or any other software it is a good idea to check if the problem is related to the framework and has been solved by other people already.

If both suggestions don't work out, it should be thought about where the program is the most complex as a developer is more likely to make a mistake in a complicated part than a simple one.

2.3.6 Isolate the infection in the chain

Isolating is done by creating a hypothesis and checking if it is true, then repeating the process until the problem is solved. A sample run would be:

1. Hypothesis - The application creates a faulty value.
2. Prediction - The faulty value is set to the particular variable in the expected line.
3. Experiment - Using the debugger the prediction is confirmed.
4. Observation - The faulty value is set to the variable.
5. Conclusion - The hypothesis is confirmed.

As the problem is not located yet a new hypothesis needs to be thought of. It is common to check if the failure is present at an earlier stage.

1. Hypothesis - The infection does not occur until function Y is called
2. Prediction - The variable should hold a sane value before Y is called
3. Experiment - Using the debugger the prediction is confirmed
4. Observation - The variable is already wrong
5. Conclusion - The hypothesis is rejected

As the infection is already present at the function call, the function should be examined.

1. Hypothesis - Invocation of function Y with the faulty value causes the problem
2. Prediction - If the function is called with the correct value the program runs correctly
3. Experiment - Using the debugger the function call is called with a correct value
4. Observation - The program runs correctly
5. Conclusion - The hypothesis is confirmed

This method has to be repeated until the root of the problem is discovered. The advantage of this method is that the problem will be identified eventually.

2.3.7 Correct the defect

Once the defect has been found, it can be corrected. But solving the defect is not enough as it is also necessary to check if the problem is now solved. There are two reasons why this is necessary:

1. The defect was not responsible for the error - It is common to find other bugs while investigating a problem as the code is scrutinized. Another option is that the problem consists of two defects, not one. For both these reasons, it is necessary to check whether the defect was solved or not.
2. Is the cause an error? - Sometimes the deducted cause of the problem is in fact not the cause, and the correction is just another defect that fixes the problem in the particular problem case. To make sure that the cause really was the cause it is recommended to think about the correction and if it solves the problem for all cases without compromising all other working cases.

Once all of these steps are done the only thing left to do is to mark the problem as solved in the problem report and add it to the next patch of the software.

3 Related Work

The goal of this section is to give the necessary information to understand the rest of the thesis. Debugging is a complicated process that is made much harder by distributing the system to more than one physical location. To better understand the methods and principles, this chapter will use a simple example application to comprehend the differences in each method easily.

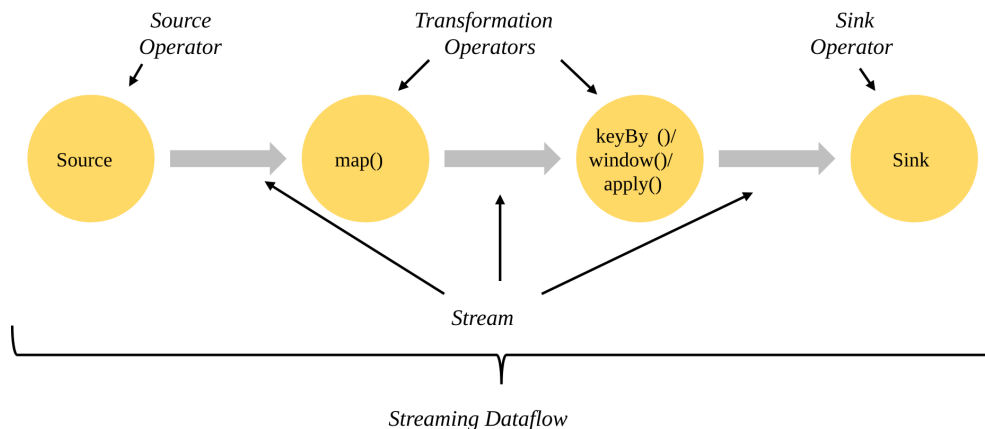
3.1 Flink Framework

The Flink Framework is a stream processing framework for the JVM. It is written in java and scala and can be used with a variety of languages and technologies. As is to be expected the most supported languages are java and scala, although there is also a python wrapper and a few more available. The following section will clarify how exactly a Flink application works and what kind of debugging and logging tools it already has.

3.1.1 Basics

Flink is a Framework for processing applications that are distributed across multiple computers. This part will lay out the foundations of the Flink Framework. Flink applications have a simple base structure that is used by the developer. Each application defines tasks which have a particular structure. Each task has an input and output stream. These are called source and sink. Because both the source and sink of these tasks are streams, they can be attached to each other thus creating a data flow from task to task until the wanted end state is reached.

Figure 3.1 shows a basic program flow. On the left side, it starts with a source. The source is then transferred over a data stream to the first task. The map operation is executed, and the result is sent over a data stream to the next task where the same procedure will start again. The resulting data in this example is the sink, meaning we reached the end of the program. The sink will typically be connected to a database or some other kind of technology for preserving the data. The most basic option would be to write the sink onto the standard output on the console.

Figure 3.1: Simple Dataflow in Flink [Flink \(c\)](#)

Until now the system can only be distributed by having the different tasks on different physical computers. This distribution process will not suffice when the amount of input data gets too high. To achieve real distribution in the application, each task can be run multiple times as so called subtasks. These subtasks run as a single thread in the JVM and are managed by a task manager. Task Managers connect to a Job Manager which coordinates the distributed execution. The data flow from one subtask to another can either be one to one or as a redistributing data flow. A redistributing data flow is necessary to achieve an even distribution of data at the next subtasks.

Figure 3.2 presents the same example as figure 3.1 just with the subtasks shown as well. The map task is a vital step in each application as it connects the data from each subtask with each other. It is easily understood by a simple example. In an application that counts how often each word is in a text, it would only split the text at each space symbol. The redistributing data flow would then create an even distribution in the next node which specifies what to count by (id task) and applies a particular aggregation function to the "apply" operation. The result is then sent to the sink operator. This explanation is a simplification of the actual MapReduce model. For a complete overview see [Dean und Ghemawat \(2008\)](#).

3.1.2 Stream Processing and Batch Processing

Apache Flink is primarily a stream processing framework, although it can also be used for batch processing. As it makes a big difference in the infrastructure of the framework which process the primary one is, it is important to lay out the differences between the two.

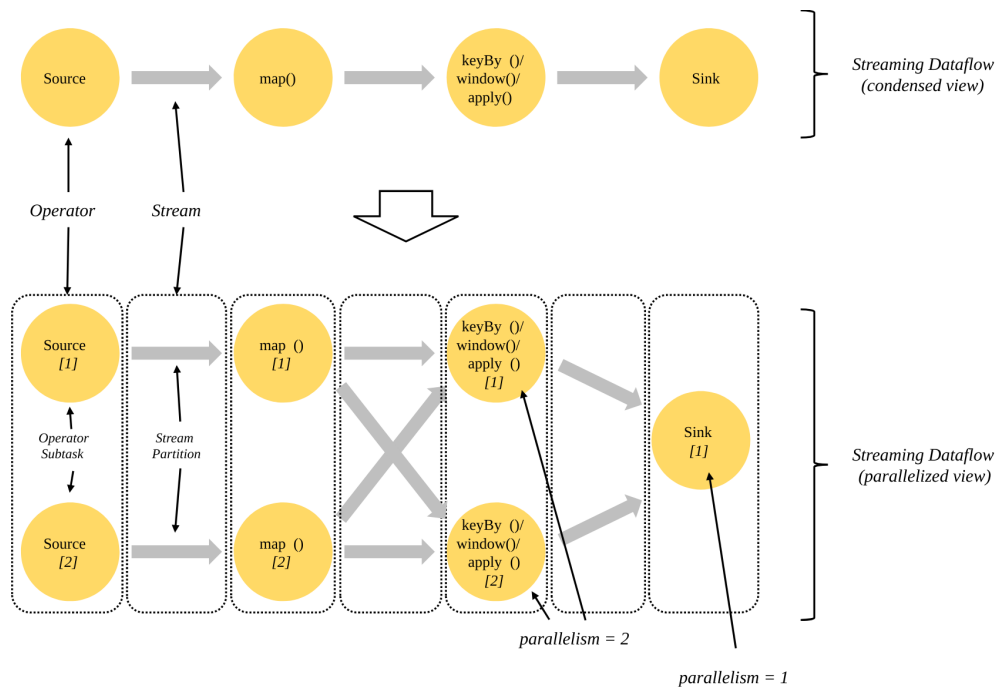


Figure 3.2: Distributed Dataflow in Flink **Flink (b)**

Stream Processing uses as the name suggests a stream to acquire the data. A stream is an endless sequence of data that can be utilised both as an input and output. As there is no end, it makes it tough to use some algorithms on it. For example, an algorithm that finds the highest number from the input. Flink uses "Windows" to solve this problem. These windows offer a frame for the operations to only use data that was received in each window. Windows can be programmed and provide various ways to define the frame. The simplest way would be just to give a time frame, but more complex structures could also be built.

Batch Processing on the other hand has set boundaries, and it is evident how much data is sent and where it ends. To support batch processing as well in the Flink Framework the windows can be programmed to have the same size as the data in the batch. That way even though Flink uses streams it still supports batch processing.

3.1.3 Debugging and Checkpoints

Flink comes with some tools that help the programmer debug his application and help prevent errors stopping the application. This section will highlight what Flink does different or on top of the regular java/scala debugging features.

Metrics

In a distributed application it can be hard to find out why something is not working properly or why some function performs worse than expected. To help the programmer understand these problems Flink provides metrics, these count or measure throughput on specific points in the application and send this information to a reporter. The reporter provides the information to external applications so that the metrics can be analysed as needed. There are four different metric classes available, a counter that can be in- or decremented, a gauge that can provide the value of a particular variable, a histogram which measures the distribution of long values and lastly a meter that measures the average throughput.

As applications could have a huge number of metrics which would make it very hard to find anything, it is important to have some grouping mechanism. Flink provides scopes to solve this. There are two types of scopes, user-scopes, defined by the user, and system-scopes that hold current information about the system state like the task in which the metric was saved. When a metric is registered, an identifier and a system scope have to be specified, a user-scope can optionally be added.

In addition to the metrics above, Flink automatically collects system information like RAM usage, the number of threads, network usage and much more.

Checkpoints

In distributed systems it is quite common that parts of the system crash, this can have many reasons, the incoming data could be formatted wrongly, a data transmission could break away before it is finished, and so forth. Normally the application would log what went wrong and stop or restart the whole application. Depending on how large the application is, this could cost a lot of time. Flink creates checkpoints which it automatically falls back to if the application crashes. The way these checkpoints work is by periodically injecting barriers at the source. After a task receives a barrier at one of its inputs, it blocks that input until it received a barrier at all of its inputs. Once that happens it takes a snapshot of all the data it received since the last snapshot. This way it is guaranteed that every piece of information is part of a snapshot all the time. To keep these snapshots from taking up too much space on the hard drive flink only stores one snapshot for every task and overrides it with the next. For a more in-depth explanation of this algorithm see: [Carbone u. a. \(2015\)](#)

As well as these automated checkpoints Flink provides user-created checkpoints called savepoints. These can be set by the user and are not getting deleted every time a new one is created. They are used to pause the application for example.

3.2 Similar work

Other people did similar work on another framework and distributed systems in general. This section will explore to what conclusions these people came and what can be applied to Flink as well.

3.2.1 BigDebug: Debugging Primitives for Interactive Big Data Processing in Spark

Gulzar u. a. (2016). The University of California built a debugger for "Apache Spark" which focuses on similar points as this thesis. The motivating intention behind their debugger is to find the cause of the problem easily without having to go through millions of logs, and without stopping the system itself. BigDebug uses five different methods to solve this predicament. This section will not only lay these methods out but also show if they are already present in Flink.

Simulated Breakpoints are used to debug parts of the application without stopping the running system in every node. This is done by spawning a new process with the same beginning states as the remote one. After that the newly spawned process can be debugged without intervening in the running system.

On-Demand Watchpoints with Guard are user-defined methods in the system that can be set to inspect the value of a certain variable, check that value and store it in case the value of the variable fails the check. This method can be useful in certain situations, for example when checking a variable that is supposed to hold a zip code to see all values that are malformed. This feature can easily be used in the Flink Framework with the use of metrics.

Crash Culprit and Remediation focuses on two things, first collecting interesting data in case of a program error so that the user can easily find out why the crash happened and secondly avoid rerunning the whole application when a crash happens. These two are grouped together as they are solved together as well. After a crash, the user will get the value that was responsible for the error, for example, the input could have been "23s" instead of 23. The rest of the application will continue to work and once the user corrects the value to "23" the program reruns this task. In Flink, this is partly implemented, as Flink creates automated checkpoints the application will not completely stop when it crashes as it would be in Spark. Flink offers no way to modify values that lead to a program error though.

Forward and Backwards Tracing is done to make it possible to track where a piece of data came from or where it will end up. This tracing is quite complicated as a piece of data in task X is a sum of all modifications that occurred before that task. The way this is implemented is by tagging each incoming piece of data with a unique identifier and adding all the identifiers that are used to create a new piece of data together. On the other hand, if data is split into multiple pieces, like splitting a sentence into words, every new piece gets a new ID, and the relation between the old and the new ID is saved. That way a programmer can easily backtrack to where a malfunctioning piece of data originated.

Fine-Grained Latency Alert is used to identify which records are causing delay. Baseline Spark can already measure the time between each task. BigDebug builds on this feature and extends it by making it available for each operator. This part is quite different in Flink as Flink offers no direct way to apply a metric to each task or operator.

Conclusion

BigDebug solves some challenges that a Flink developer has as well. It solves them only in the spark framework and suggests no methodology or procedure for debugging a spark application, it just provides some tools for the developer to use.

3.2.2 Debugging Distributed Systems

[Beschastnikh u. a. \(2016\)](#) tries to lay out fundamental challenges developers face when building distributed systems and how to solve them.

They give an overview of seven approaches that can help to build, validate and debug distributed systems. What follows is a summary of the seven methods.

Testing is crucial but can only help to reveal some errors as testing every piece of the application is impossible.

Model Checking is a form of testing that automates the testing process somewhat by checking every possible input up to some upper bound in a predefined way. The way the model checker works depends completely on the model checker itself, as there are a lot of different options available. There are for example symbolic model checkers available that explore possible executions mathematically or black-box checkers that just run the application with various inputs. Model checking can be helpful as it can cover much more ground than manually written tests, as well as check parts of the application developers, might not have thought of.

Theorem proving is a mathematical method of proving that the distributed system is free of defects even before writing a single line of code. It is mostly used for proving that the core of a new application is bug-free before spending lots of time building an application that might not work properly. Amazon is one of the companies using it and released a paper about it as well: [Newcombe u. a. \(2015\)](#)

Record and Replay is a method for analysing a single execution of the application to gain insight into why an error occurs in a system that has non-deterministic events, as these change every time the program is run even if the same input values are used.

Tracing is the method of following data through a system. It has the same goal as Record and Replay although it is easier to understand what is happening as only a subset of the data is shown.

Log analysis is a method for debugging black-box systems, but can be used on every application. It is done by applying algorithms to the logs to find bugs that are not easy to spot.

Visualisation is used to make distributed systems more transparent. If a developer has a good visual representation of his work, it is much easier for him to find bad design choices that could lead to bugs.

Conclusion

The article provides some basic methods for debugging distributed systems, the analysis of which method is useful is short though.

4 Debugging Flink Applications

This chapter is the centre part of the thesis. It explains in detail how the method works. It consists of two main sections, the first of which explores in detail what a developer should do while writing the program to minimise the work when debugging. This step is crucial as it is much harder to recreate problems locally as opposed to regular Java applications. The second part focuses on the debugging itself once a developer is informed of an error and has to figure out what is causing it.

4.1 Better Developing

Flink applications are run remotely and without an active user providing input as it is typical for a regular application. Not having a user makes it much harder to recreate the problem as we only have the log files and the stack trace as information. As such it is crucial to have all the information at hand when the program fails, or we notice discrepancies in the resulting data. The only way to make sure that the information is accessible once a problem is reported is to think about what data is necessary for the debugging developer while writing the program. Another issue is that some problems are unique to a distributed environment and won't happen when running on a local host. This section will explain what can be done while writing the program to make the debugging process much easier.

4.1.1 Building Tasks

The Flink framework uses a few features that set it apart from standard Java applications. These are referenced in [3.1](#). The one that sets these applications apart from standard Java ones while writing the application is the distribution. Applications should be built in a modular way to take full advantage of these features and also make the program easily debuggable. The development of these tasks should be done using a very similar set of rules as the Unix philosophy states. In fact, Flink modules are not too different to Unix command line tools, they both provide a service while taking an input/output in a predefined way. For the Unix command line, this is the standard output, for Flink programs, these are input and output streams.

Unix Philosophy

The Unix philosophy is defined by Doug McIlroy, the inventor of the Unix pipe as follows: [McIlroy u. a. \(1978\)](#)

1. Make each program do one thing well. To do a new job, build a fresh rather than complicate old programs by adding new features.
2. Expect the output of every program to become the input to another, as yet unknown, program. Don't clutter output with extraneous information. Avoid stringently columnar or binary input formats. Don't insist on interactive input.
3. Design and build software, even operating systems, to be tried early, ideally within weeks. Don't hesitate to throw away the clumsy parts and rebuild them.
4. Use tools in preference to unskilled help to lighten a programming task, even if you have to detour to build the tools and expect to throw some of them out after you've finished using them.

Most of the points mentioned here are of some relevance for Flink programming as well. Each task should do one job and do it well. The next paragraph will look into why that is. The second point is necessary by default in Flink; each task has to use the provided streams to work. The third point is equally important if not more important in Flink applications as it is in Unix programs. Always test each task individually to make sure it works as designed and has no flaws on its own, only then can the whole application work without any problems.

Dividing the program into various modules has a lot of advantages:

1. Easier to develop - It is much simpler to develop a smaller application as it is much harder to lose track of what each piece of code should do. This reduced complexity in return reduces the likelihood of mistakes. Once the program is completed, it results in a more stable program that can be debugged easier.
2. Better distribution - As every Flink task can run on a different computer the smaller the tasks are, the better the Flink Job Manager can distribute the load evenly among the available resources.
3. Checkpoints are easy to find - Checkpoints are a core piece of Flink technology. It allows the framework not only to jump back and repeat a failed run without having to restart the whole application but also provides information about which state the application is currently in. This is extremely helpful as a lot of Flink applications only end when the program is cancelled by the user.

The modulation of the program not only helps to achieve the advantages of the framework but also supports with debugging later as a lot of the information needed are gathered at the checkpoints.

Where to split the program

It should now be understandable that the programs should be divided into multiple tasks, the next question now is how to break the program to get a simple program where there are enough tasks but not too many, as too many would lead to the opposite effect we want to achieve.

Why are too many tasks bad? When there are too many tasks, it gets even more complicated than when everything would be in its task as basically every line of code would be in a different place. Additionally, it wouldn't increase the performance as each task has some initialisation work that would diminish the performance gain achieved by distributing it perfectly.

It makes sense to use the Unix philosophy of having one task do one thing. In most cases there are some obvious logical places to split the program. The same is true for Flink applications. Data is modified or analysed and each of these tasks could be one transformation of the data.

4.1.2 Metrics

Now that the architecture of the program is done the next question is what metrics to use in which positions to achieve the optimal security.

Once the program architecture is finished, it makes sense to think about what kind of metrics can be used where. Metrics are used to monitor the program without having to debug it and are crucial in notifying the developer if something looks wrong. The first step is figuring out where to use metrics. A good start is to look at the application in the worst possible way; what is the most likely area that an error will occur in, after that it makes sense to surround the area with metrics that will catch and log the gathered data. Another great location for metrics is at positions where the incoming data is simple, and metrics can easily be implemented. This should be the case in between tasks. As optimally each task only does one thing it should be easy to check whether the starting and ending assertions are valid.

4.1.3 Logging

Logging in Flink is straightforward and is used the same as in every other Java program that uses log4j. As Flink already provides the necessary libraries to use log4j all a developer has to do is to write the logging config file. Although the logging process itself is the same as every other Java application, it should not be forgotten to use the different log levels that log4j provides. There is a lot of logging happening out of the box just by the Flink process itself. The six logging levels, from highest to lowest are:

1. FATAL - the highest logging level, should only be used when the application cannot continue to work because of an unexpected error.
2. ERROR - whenever an unexpected exception is thrown it should be logged.
3. WARN - warnings that could lead to errors later on. These are difficult to think of beforehand but if used correctly are very valuable for the debugging developer.
4. INFO - relevant information like successful database connections and other milestones in the application to let the reader of a log file understand at which point in an application the program currently is.
5. DEBUG - should be used to record relevant information along the way that could be useful to a programmer when debugging. This could, for example, contain values of variables like database connection strings.

6. TRACE - is used to let a developer searching for a bug understand the path the application took. Should be logged into a unique log file as it would flood every other one.

4.2 The Debugging Process

There are multiple reasons how an error might be discovered; the most common one is an exception in a log file. Another option is that the end user of the results finds that some of the results are incorrect. Both of these cases require slightly different handling. An excellent way to start the debug process is by using a modified version of the Traffic approach introduced in [2.3](#).

4.2.1 Common Problems

Like with all frameworks and technologies there are some problems that often occur in Flink applications. This section will outline a few to shorten the debug process if one these is the issue:

- `NotSerializableException` - because everything is sent over the network almost everything that is written has to be serializable. This includes all transitions. Because serialization is not often on a developers mind and does not show up as an error in the IDE, it is often forgotten. There can be two reasons for this error, either some class is not implementing the "serializable" interface or local fields in a class are not serializable. If a field does not need to be serialized, because it does not need to be sent over the network, it can be marked with the "transient" keyword.
- Network Exceptions - each system has to be able to communicate with the other systems. See [4.2.7](#) for a list of possible reasons.
- `ClassCastException` - Sometimes multiple versions of the same class are loaded, this can happen when class loaders of different tasks and are later assigned to each other. Either a library that caches loaded classes is used in which case one should try to remove them, or the Flink itself is loading classes in the wrong order. This can be solved by adding "classloader.resolve-order: parent-first" to the configuration file.

The Flink website features a complete list that is regularly updated: [Flink Common Error Messages](#)

4.2.2 Track the Problem in the database

Tracking a problem is essential in every development cycle no matter in which language or with what framework and Flink is no exception. It is crucial for every developer to track the status of problems in Flink applications as it helps to minimise work. Along the already mentioned advantages in chapter 2.3.1, like having an easily accessible database of open problems and knowing which problems are more important than others, tracking the problems of Flink applications offers some other advantages as well.

1. Having access to relevant log files.
2. Knowing how past problems were solved.
3. What relevant metrics were when the error occurred.
4. Which subtask of what task manager failed which allows seeing if problems only occur on one machine.

To achieve these advantages, additionally to the already mentioned fields the tracking database needs a few additional columns. As soon as a problem is experienced the current log files should be saved so that it is easy for the developer to find the relevant lines in the log file even if he starts debugging months later. Secondly, for the same reason, all appropriate metrics should be saved as well.

The resulting columns now are:

1. Description
2. State
3. Resolution
4. Assigned Developer
5. Severity
6. Link to logs
7. Steps that were taken to resolve the problem
8. Relevant metrics
9. Task manager that was used

4.2.3 Reproduce the failure

Reproducing the problem is probably the most challenging part of debugging Flink applications. As there is no user to report the problem the only help the debugging developer has is information provided by the problem report from the last chapter.

There are two possibilities how a problem is discovered, and both require different steps to reproduce the problem. The first and more difficult one is that an error in the resulting data is discovered without an exception being recorded in the log files. This means that the program is doing something different than what the developer expected when writing it. The second option to discover a problem is by having an exception showing up in the log file.

Faulty resulting data

This section will use the word count application as an example program to debug. Notice that the exact implementation of it is irrelevant at this point. The application is a black box as only the incoming and resulting data are known. To always have the same expected result the following sentence will be used as an input each time: "Hello hello flink flink flink one two". The expected result would be:

```
1 hello - 2, flink - 3, one - 1, two - 1
```

The following text explains the process that is used in this case alongside a diagram: [4.1](#)

The first question that has to be answered is "how much data is affected". Are only a few pieces incorrect or is everything faulty? Imagine the result of the application would be:

```
1 hello - 3, flink - 1, one - 1, two - 2
```

Each word has the count of the next word. So this would be considered as the second case "everything is faulty" even though the word "one" has the correct answer. This is important to note as sometimes a fault in a program can still result in a correct result. This first question just differentiates between a few faults and a majority of faults, so the debugging developer has to look at the whole picture and see if the majority of data is corrupt. An example of only a few faults would be:

```
1 Hello - 1, hello - 1, flink - 3, one - 1, two - 1
```

Here only one additional word was counted ("Hello").

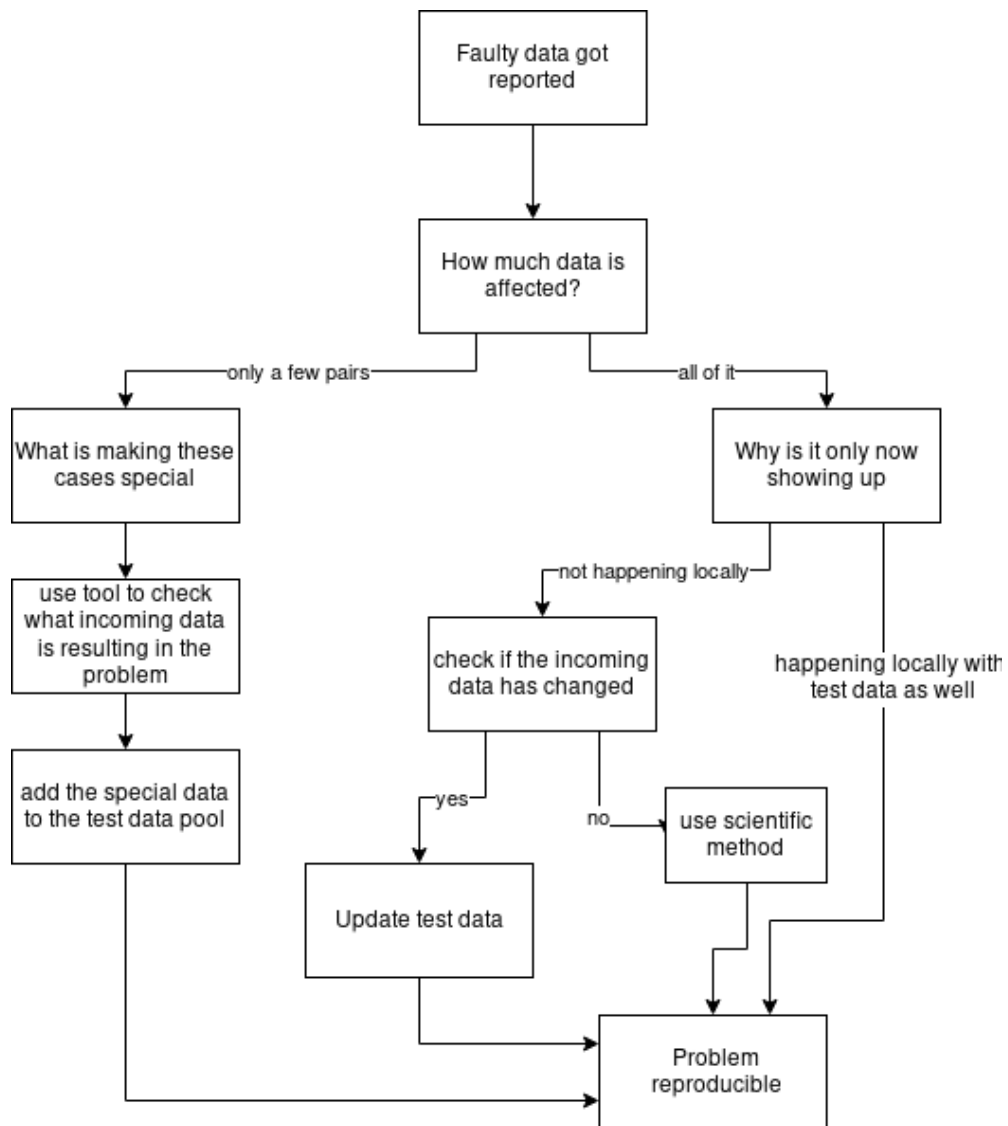


Figure 4.1: Debugging faulty data

If the first case is valid (all data is faulty) the next question that should be asked is why the problem is only now showing up. If the problem is observable for almost all the resulting data, surely it should have been noticed while testing the application. In most cases, the problem was either found during testing in which case the problem is already reproducible or was not observable on the local test machine. That means that either the incoming data is different to the local one or that something is being executed differently on the remote network than on the local machine. As Flink is responsible for the distribution and everything is running in a

JVM, it is implausible that Flink is to blame. In most cases, the data on the server will differ from the local one. If that can be confirmed the only thing left to do is to update the local test data so that the problem can be reproduced locally.

The other option was that only some pieces of data were wrong. In that case, the process of reproducing the fault is entirely different. There are two options available. First, figure out what makes the faulty data unique in comparison to the other data. In the word count example, this would be the capital "H" at the beginning of the first "Hello". If this option is successful, the unique case can be added to the test cases, and the reproduction was successful. If on the other hand, the developer can't figure out why the one failing case is different to the others the tool that is written alongside this thesis can be used, it will be explained in detail later on. It can show which incoming data was leading to which result. In the example above it could show that the first "hello" was the result of the original sentence. As there is only one sentence in this example that is not very helpful, but in a more realistic use case, there could be millions of sentences where just a few have capital letters in them. Once the starting sentence is discovered, it can easily be reproduced.

The last step before moving on to the next section is to make sure that the reproduction works. There is little to no use in finding something in the code that is supposedly the problem only to find out later that the problem is something else.

The fault should now be reproducible on a local machine as the affecting test data was found. The only fault that remains are problems that only occur on the remote network and are not happening because of differing incoming data. In almost all cases the root of these problems is either incompatible Flink components or network issues. So before starting a more complex search for the problem the developer should check the Flink documentation if all used components, like data sources, the chosen backend etc. are compatible with each other. Furthermore, it makes sense to check if all of the running systems can send packages to each other and the database as firewalls and other network related issues could be the issue. If none of these are quickly found it is recommended to use the scientific method described in: [2.3.6](#)

Exception in log file

The most common way to discover a failure in a program is by discovering an exception. It makes no difference if this exception was found in a log file or directly in the developer's console. The steps that are necessary to reproduce these errors are often easier as well.

1. What kind of error was thrown? - The developer should always keep in mind what type of exception was thrown as it is much easier to find the origin of the failure when knowing what to look for.
2. Where was the failure thrown? - This information is the starting point for the search of how to reproduce the failure. It is included in the stack trace of the exception so it shouldn't be a problem finding it.
3. Isolate the conditions in the method - Knowing under which conditions the failure occurs is crucial. Without this information, it is almost impossible to reproduce the failure reliably as only a few errors occur under all circumstances. The developer should start by just focusing on the method the exception is thrown in and finding the conditions by looking at the "if statements" that preceded the line in question. Note in which state each variable has to be for the exception to be thrown. This can include the variable that is throwing the exception as well. In most cases, the developer should now already have a pretty good idea as to when the error occurs and if that is the case can skip the next step to save time.
4. Repeat the same step for the method - Once it is understood in which conditions the failure occurs in the given method, the developer has to find out in which cases the method is being called with the conditions that were deducted in the last step. This can be done by repeating the step above only for the line where the method was called, which is included in the stack trace of the exception as well. This step has to be repeated until the start point of the application has been reached. At this point, the developer should have a good understanding when the exception is thrown and should be able to reproduce it.

Conclusion

In most cases, the developer should have no problem reproducing the error easily without going to much in depth into the steps presented in this chapter. Although these steps help to guide a developer that might not even be deeply familiar with the code through the process with ease sometimes problems occur that can not be found using these steps as they are too unique and would not appear for anybody else. If a developer stumbles over such a problem, it is suggested to apply the same method that is explained in [2.3.6](#) by creating a hypothesis checking the outcome and repeating the process until the failure can be reproduced.

4.2.4 Automate and simplify the test case

In comparison to typical Java applications where automated testing can be quite tricky as a graphical user interface, and user interaction is involved, automating the test cases for Flink applications is quite easy assuming that the problem is reproducible on a local machine. All that the developer has to do is take the result of the last section and write a test with the specific start parameters. The "Traffic Approach" also suggests simplifying the test case to be as easy to implement as possible so that only the error and nothing else gets checked. This can be useful to make it easier to help spot the problem in the code later on but is not needed for Flink applications. For example, if the word count application were to crash on "-" characters and the input test case would be an entire book page of words it makes it harder later on to find the actual character that is causing the fault but makes it easier to create the test case. As understanding the problem and simplifying the test case is mostly the same it makes sense to leave the test case long and rebuild it once the failure is understood. This way if the same mistake gets built into the code again it is easily spottable as the test case already has a name, description as well as a ticket number associated with it. Another developer could then easily reproduce the steps that were taken to solve the problem the last time.

Flink Spector The problem with running the whole application in a test case is that there is a lot of overhead. The whole application has to be built the different task manager are started all to run the program with one input. The worst part about this process is that it does not even eliminate all the failures as the execution environment differs (because it is executed locally) as well as the source of the application (because data is provided by the test case instead of the data store) and lastly the sink (because the data is not actually saved). The only information that is gained this way is that the transformations are doing what they are supposed to do. This could be much easier done if data streams could be tested individually without having all the overhead of starting Flink. The Otto Group had this issue and decided to write a plugin for Flink that allows data streams to be unit tested called the **FlinkSpector**. It can be used in the following way once the plugin is enabled:

```
1 @org.junit.Test
2 public void testWindowing() {
3
4 // Define the input DataStream:
5 DataStream<Tuple2<Integer, String>> testStream =
6     createTimedTestStreamWith(Tuple2.of(1, "fritz"))
7     .emit(Tuple2.of(1, "hans"))
8     .emit(Tuple2.of(1, "heidi"), intoWindow(30, seconds))
```

```
9     .emit(Tuple2.of(3, "peter"), intoWindow(1, minutes)
10     .repeatAll(times(2))
11     .close());
12
13 // Lets you query the output tuples like a table:
14 OutputMatcher<Tuple2<Integer, String>> matcher =
15     //define keys for the values in your tuple:
16     new MatchTuples<Tuple2<Integer, String>>("value", "name")
17         .assertThat("value", is(3))
18         .assertThat("name", either(is("fritz")).or(is("peter")))
19         .onEachRecord();
20
21 assertStream(someWindowAggregation(testStream), matcher);
22 }
```

Listing 4.1: FlinkSpector Example: **FlinkSpector**

Put simply the developer has to create a test data stream of incoming data and the similar thing for the data returned by the transformation. Once those are created the unit test is the same and can be used with the "assertStream" method call (Line 22).

4.2.5 Find possible infection origins

This step in the debugging process is just a collection of possible points to start the debug process from. It is always possible to change this point later if the chosen point turns out not to be the origin of the infection. Flink applications are very different to other Java applications in this section. Flink uses a special way to write applications that make it possible to run on multiple machines without having to specify how and more importantly which information are transferred between different machines. Because of this complication, code that is written in Flink is being executed before a single piece of information is read. At the end of each application is an execute-statement which then starts the application. When debugging the code written by the developer with a debugger the only information that can be gathered is whether or not Flink can successfully build the execution tree or not. The only relevant information this provides to the developer is that all checks that Flink itself runs before executing are true. At the time of writing this thesis (Flink 1.5), it ensures successful connections to the databases or other data stores used by the application and compatibility of the used data sources and backend. As an example, if a backend that supports snapshots is chosen, Flink will check if the first data source supports rollbacks as that is a requirement of snapshots. Other problems like incompatible datatypes can still occur afterwards.

This section includes the next step of the traffic approach as Flink applications are structured in a way that makes it very easy to find the possible origins of the information given here. To make it even easier the word count example from the beginning of the thesis is used again, this time as an actual Flink application:

```
1 public class SocketTextStreamWordCount {
2 // configuration lines removed for clarity
3 public static void main(String[] args) throws Exception {
4 // configuration lines removed for clarity
5   DataStream<String> text = env
6     .addSource(new FlinkKafkaConsumer011<>("test",
7       new SimpleStringSchema(), properties));
8   DataStream<Tuple2<String, Integer>> counts = text
9     .flatMap(new LineSplitter()).keyBy(0).sum(1);
10  counts.print();
11  env.execute("Java_WordCount_from_SocketTextStream_Example");
12 }
13
14 public static final class LineSplitter
15   implements FlatMapFunction<String, Tuple2<String, Integer>> {
16   @Override
17   public void flatMap(String value,
18     Collector<Tuple2<String, Integer>> out) throws Exception {
19     // normalize and split the line
20     String[] tokens = value.toLowerCase().split("\\W+");
21     // emit the pairs
22     for (String token : tokens) {
23       if (token.equals("test")) {
24         throw new Exception();
25       }
26       if (token.length() > 0) {
27         out.collect(new Tuple2<String, Integer>(token, 1));
28       }
29     }
30   }
31 }
32 }
```

Listing 4.2: Word Count Example Flink

What this program does is the same as the earlier word count example, it counts how often each word is in a specified input. As this is written in the Flink framework, it uses streams to store the data and transformation methods to connect the streams in a way that splits and then counts the words. The first stream is used to gather the incoming data from an Apache Kafka cluster (Line 5). The data stream is called "text" and is used in the next stream as a source (Line 8). This data stream is called "counts" and is doing the actual work of the program. It uses a flat map function to split the string at each space symbol and uses the newly created "words" to add them to a new String, Integer tuple. Flink then adds merely all these tuples that have the same key together and adds up the Integers. This can be done as the whole transformation method is an extension of the flatMap method that is eventually adding all created tuples together.

Exception before application is executed

The first case that we are going to look at is that the application throws an exception before it is being executed. This is the case if any of the lines before the execute line appear in the stack trace. In our example the execute line is line 11. so the possible exception that would fit in this category would be:

```
1 java.net.ConnectException: Connection refused (Connection refused)
2   at com.jakobjoachim.WordCountFlink
3   $WordCountFlink.main(WordCountFlink.java:5)
```

Listing 4.3: Connection Exception

As stated above connection problems are not the only issues that can cause an exception before the execute line is reached. It is also possible for some of the Flink components to be incompatible with others. No matter which of these two is the case the procedure is the same. As the failure happens before any of the Flink specific issues are encountered the method described in the standard "Traffic Approach" [2.3.6](#) is enough to find the root of the problem. This means going to the specified line in the stack trace and checking if the line could be responsible. As Flink applications are only a set of transformations on data streams, this search is relatively easy as it is most probably found either at one of the transformations or the initial connection to the database.

Because of how Flink works in this step there are only two possible roots for an issue in this category. Either the line that throws the exception or the other side. In case of Flink component issues, this means the line that sets up the data stream or environment that is being used in the line that is throwing the exception. Because of this, it is not necessary to follow the next step of the Traffic Approach (focus on the most likely origins).

Exception in a transformation method

Another possible case is that an exception was thrown while the application is running. On contrary to exceptions that occur in the data streams and will not show the line written by the developer in the stack trace, exceptions in transformation methods will show up as the source in the stack trace. If we look at line 23 and 24 in the example application, we will see that an error is thrown whenever the word "test" is read. Of course, this is set up to throw an error, and a real application would be more likely to throw a null pointer exception or similar when transforming the data (line 27). When "test" used as part of the input the following stack trace is thrown:

```
1 java.lang.Exception
2   at com.jakobjoachim.WordCountFlink
3     $LineSplitter.flatMap(WordCountFlink.java:24)
4   at com.jakobjoachim.WordCountFlink
5     $LineSplitter.flatMap(WordCountFlink.java:14)
6   at org.apache.flink.streaming.api.operators.StreamFlatMap
7     .processElement(StreamFlatMap.java:50)
8   at org.apache.flink.streaming.runtime.tasks
9     .OperatorChain$CopyingChainingOutput
10    .pushToOperator(OperatorChain.java:549)
11  . . .
```

Listing 4.4: Exception in Transformation

The only lines that are part of our code in this stack trace are the lines in the flatMap method. The stack trace does not show which of our data streams called the method. This is because the data stream was converted to a Flink internal data stream before the application was run. As such the next step is to find out which of our data streams called the transformation method in question. Possible infection origins are now:

- Bug in the code of the flatMap.
- Bug in transfer to the data stream that called the flatMap method.
- Bug in any data stream or transformation of data streams before that.

Failures in transfers between data streams are highly unlikely when the program can be reproduced locally as any incompatible data types and components are already checked. Although this is only true if Java Generics were used with all appropriate variables. These possible infection origins are enough to uses as a basis for the next section called "Isolate the defect in the chain" [4.2.6](#).

Exception in internal Flink code

The last possible location to throw an exception is the internal Flink code itself. This can have a lot of possible reasons and the least amount of helpful information in the stack trace. One possible reason could be a timeout while sending data from one machine to the other. The exception will show that a network timeout was the cause, but there is not much else in the stack trace that can help the developer narrow down the search. This is why logging correctly is crucial in this section as the most relevant pieces of information are which data stream was responsible. The log file contains such information and can be seen either directly by reading the log or by using the Flink web UI if it is used.

We should now know which data stream is causing the exception but possibly, depending on the thrown exception, not why. This step of the debugging process requires a lot of understanding of how Flink and the program that is being debugged works. There are a lot of different reasons why an internal exception might be thrown. The good news is that it is improbable that the fault is somewhere else in the code and that the data stream that is throwing the exception is working fine. In almost all cases it is either the data stream throwing the exception or the transfer to that data stream. Because of that, we can use the data stream line as our infection root for now.

No Exception but resulting data is faulty

If there is no failure and the resulting data is wrong, the developer can use the line where the false data got stored as the "exception throwing line". At that moment the failure is already present in the application so the line can serve the same way an exception would. If the program is too complex to overview how it got to that point, a stack trace can be printed to the console by adding the following line (remove this afterwards):

```
1 System.out.println(new Throwable().getStackTrace())
```

Listing 4.5: Create Stacktrace

Brief introduction to the Flink Backtracking tool To narrow down the search to a line that might be responsible, the developer can use the tool that was designed alongside this thesis and is explained in detail in chapter 5. The tool allows the developer to see the content of the data streams in real time and also shows the relation between the data and multiple data streams. This is done by using the Flink watchpoint mechanism to split the inputs into groups. In the WordCount example which has two data streams, this would result in the following with an example input of "this is a test this is a test is this a test?":

- **Data stream 1 - Watchpoint 2321:** "this is a find this is a find", **Watchpoint 2346:** "is this a find?"
- **Data stream 2 - Watchpoint 2321:** "this - 2, is - 2, a - 2, find - 2", **Watchpoint 2346:** "is - 3, this - 3, a - 3, find? - 1"

If this was part of a more significant input in a more complex program, the developer might not have seen directly why the word "find" was only counted two times when it was clearly in the input three times. Because he now has a direct overview of what data transformation happened at which point he can easily spot at which transformation the failure is firstly observable and use that as the infection origin.

4.2.6 Isolate the infection in the chain

Now that the root of the infection was narrowed down to a few possible lines of code it is time to isolate these lines to see if they are in fact the cause of the issue or merely another infected line.

It is yet again necessary to split this section into multiple parts depending on which kind of infection is present. We can skip this chapter altogether for exceptions before the application is executed because we already know for sure which line is causing the issue and only have to correct the defect.

Exception in a transformation method

Isolating the transformation or data stream that is causing the issue can be done by firstly looking at the exception itself. It might be obvious that the line showing up in the stack trace is responsible or that such an exception should never be allowed to crash the program. Another important step while finding the issue is to make sure that exceptions that crash the application are only doing so when there is no other way. It might be possible to log the exception and ignore that part of the input as long as the program can continue running. Most

Flink applications run on huge server farms where each hour costs the company a lot of money so full application crashes should be minimised.

If the issue is not apparent, it makes sense to use the tool that was developed alongside this thesis that was already briefly mentioned earlier: [4.2.5](#). In the same way that the tool is useful when finding corrupted data in the streams it can help us here to gain a deeper understanding of what is happening in each stream. Together with the reproducible test and tool, we should see how the data that is causing the issue got to the transformation in question. We can then correct the defect as we isolated the faulty transformation.

Exception in internal Flink code

No matter where in the Flink internal code an exception is thrown we can use the information from the log file to find out which data stream was responsible. Then run the program a few more times to make sure the exception is always thrown at the same data stream. The next step is to make sure that the incoming data is correct and there are no errors present beforehand. Also, it makes sense to check if the incoming data can be manually added to the test case of the transformation at the faulty data stream. If that also works, we know that the root of the infection is, in fact, the data stream. This is enough to correct the defect. If, on the other hand, one of these checks failed, we have to go back a few steps and include the option that a transformation might be faulty depending on which of these checks failed. If the exception is thrown at different places all the time, it makes sense to google the exception to make sure that it is not a common problem with Flink. If it does not use the scientific method to narrow down the search until the cause is found.

No Exception but resulting data is faulty

We now have a specific location at which the data gets corrupted. To solve the issue, the developer has to look at the transformation method that is used by the data stream. The method can then be debugged on its own by writing a unit test for the method and using the standard Java debugger with breakpoints to accurately see what is happening. To make matters even easier in case the method is very complicated the developer should still use the same method described in [4.2.7](#)

4.2.7 Correct the defect

With all the information gained by the previous steps, we now know where the issue originates from and we can fix the issue. There are a few possibilities:

4.2.7 Transformation method Faulty

4.2.7 Network issues

4.2.7 Data streams are causing the issue

Transformation method faulty

The first step in correcting a failure in a transformation method is to make sure that the unit test checking that method is updated so that the same error can't happen again. Another important thing to think about is if such an error would happen again would it be better if the application didn't crash but instead log the error as critical log level and continue? This could save the company a lot of money in computing cost. It is a good rule of thumb that Flink applications should never crash once they are deployed.

Correting the defect Now that all the requirements are done, it is time to start correcting the defect. Using the newly written test case and the Java debugger we can go through step by step and see at which point the failure appears. Most of this task should already be done by the actions beforehand. If for whatever reason the exact root of the failure could not yet be located, we can use the scientific method to further narrow down the possibilities until the root is found.

Network issues

Connection issues are quite common, following these steps should result in a working connection:

1. Insufficient number of network buffers: If this exception is thrown it means that the JVM size is too small for the amount of parallelism that is happening. It can be fixed by increasing buffer in the task manager settings. Refer to the Flink documentation for further details: [Flink \(a\)](#).
2. Network issues: Both firewalls need to be setup correctly to allow packets through. This can be checked by starting a simple TCP connection between the two systems with Netcat or a similar program.

3. Configuration issues: The database has to accept incoming connection requests from the corresponding IP range. This has to be configured in the database settings and varies between different databases.
4. Flink database adapter: Flink has to support the chosen database or a database connector has to be written. Consult the Flink documentation and check if the adapter version is compatible with the Flink version.
5. Configuration issue: The IP address or the username and password for the connection that is used in Flink could be wrong.

Data streams are incompatible

There are two possible reasons why data streams would cause problems. Either they are not compatible, or the data that is transferred in them isn't. Another option that should already have been ruled out is a network issue [4.2.7](#). Lastly, data could arrive at the latter data stream in the wrong order. This is only relevant for some applications as most Flink applications, like the word count example, do not care in which order the input arrives.

- **Data types** - The developer should make sure that Java Generics are used at every appropriate location. This ensures compatibility of the data types because the Java compiler checks them.
- **Data streams not compatible** - The developer should check the documentation of each of the used data stream features to make sure both are compatible. Lastly, if that failed, also check that the Backend is compatible in the same way.
- **Data in wrong order** - Applications that depend on data being in the correct order to be processed correctly need to use the "event time watermark" feature which can add a watermark to each piece of data and check if the previous one was processed or not before doing any more work. Because this adds a lot of waiting time and additional computing to the process it has to be enabled manually and also implemented at the specific transformations. Watermarks have to be injected at the first data source and can be used afterwards.

Conclusion

Although fixing the defect is the reason the whole debugging process was started it is not the end. As explained in the [2.3.7](#) other questions need to be answered before the developer can move on to something else.

Alongside the already mentioned tasks a Flink developer should also think about the following steps:

1. If there was no exception thrown and the failure was purely faulty data, it is not enough to have a test case that makes sure that the same error gets reported again, but the developer should also add a check into the program that throws an exception. This is important as it is crucial that faulty data can't be generated by the program.
2. If the test case was not simplified, now is the time to fix that. The developer should shorten the input to the bare minimum that would produce the failure, rename the test case accordingly and add the ticket number as a comment to the test case.

5 Flink Backtracker

This chapter will explain in detail how the Flink Backtracking tool written alongside this thesis works and how it is used. It is composed of two parts, the frontend part that shows the results in the IDE and the backend part that has to be included in the project that sends the information to the frontend. This chapter is thus split into multiple parts:

- 5.1 User manual
- 5.2 Backend architecture
- 5.3 Frontend architecture
- 5.4 State of the program

5.1 User manual

Introduction to the Flink Backtracking Tool The Flink Backtracking Tool allows the user to see what data was read at each data stream and what data lead to the result at the next data stream. The Flink Watchpoint mechanism is used which injects a barrier with an identifier ever so often. The data between each of these barriers at each data stream is the result of the data between the same barriers plus the transformation of the previous data stream. Thus the exact result of each data stream can be seen which is not possible with step by step debugging.

Prerequisites There are a few requirements that have to be met to make the program work appropriately. First and foremost it has to be a Flink application as it can only track Flink data streams. Also because it is using the watchpoint mechanism of Flink all requirements that go along with that have to be met. This means each data stream has to support rolling back data in cases of failure. Because of this requirement, it is currently only possible to use data that comes from some big data storage. In the following examples, Apache Kafka will be used.

Using the program is an easy process. The "FlinkBacktracking.jar" has to be integrated into the Flink application that is being developed, and the "FlinkBacktrackingIntelliJPlugin.jar" has

to be activated in IntelliJ to see the results of the Tool. There is only one line needed to enable the tool itself, although checkpointing still has to be enabled manually:

```
1 env.setStreamTimeCharacteristic(TimeCharacteristic.EventTime);  
2 env.enableCheckpointing(1000);  
3 FlinkBacktrack backtracker = new FlinkBacktrack(env);
```

Listing 5.1: How to use the Backtracker

Where "env" is the Execution Environment. Any number of watchpoint behaviour is possible, in this example it is set to inject a barrier after the specified amount of time. Once the Backtracker is initialised, it can be used by adding each data stream that the developer wants to be watched to the tool. This is done in the following way:

```
1 backtracker.track(dataStream, "Transition name");
```

Where "dataStream" is the dataStream that should be tracked and "Transition name" is an arbitrary name that is used to show the results under. This step has to be repeated for each data stream that should be watched. Once these steps are done, the developer can open the "Flink Backtracking View" by clicking on the button with the same name on the lower right side of IntelliJ and start the application.

5.1.1 Debugging a Program with the Backtracker

Once the plugin has been configured the application can be started. Each piece of data that is being processed by the application will show up in the "Flink Backtracker" tab in IntelliJ as well. The first tabs that are shown are the different data streams. Each watched data stream has its own tab. Each of these tabs contains a list of watchpoints that passed since the program was started and each of these includes the data that was processed by that data stream at the time of the watchpoint.

The example 5.1 shows a good use-case for the tool. The developer can easily see that the input of the program was successfully received by the first data stream. He can also see that the application is removing special characters and converting all capital letters. Because the content of the second data stream (result) is not only depending on the data stream that came before (input) but also on all the data that was processed by the stream previously. Thus the "hello,2" in the result transition at watchpoint 4316.

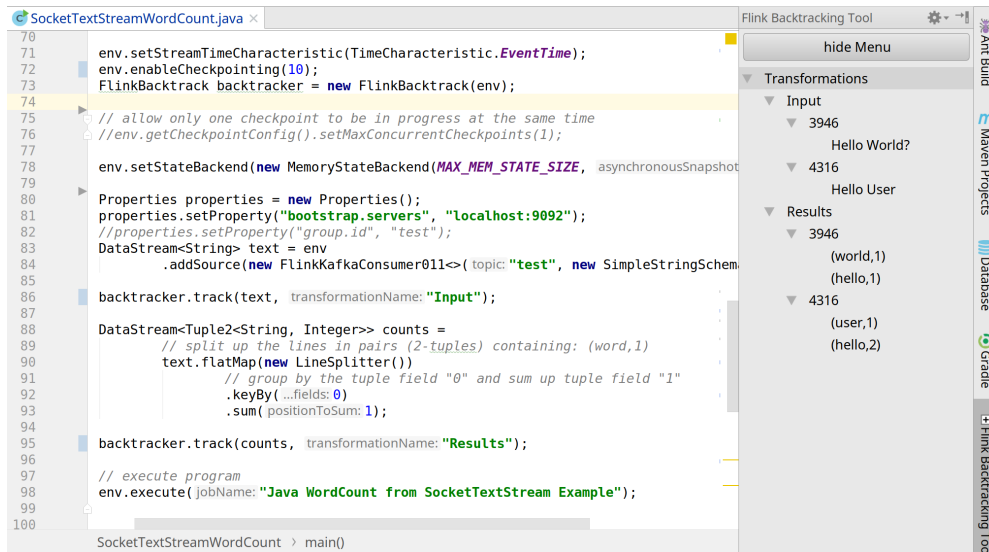


Figure 5.1: Backtracking Example

5.2 Backend architecture

The Backend is composed of two parts: One is the function that is collecting the data and sending it to the frontend the other is the main class that is used to set up the function in the correct way to ensure that the function can gather all the necessary information.

The following example (5.2) helps to underline how the Flink application is set up so that the data can be sent to the frontend.



Figure 5.2: Dataflow without the Backtracker

5.2 is a simple application with two transformations much like the WordCount application used throughout this thesis. The first transition is reading the input and the second is modifying it. At each transition, we split the data stream and send an exact copy to another transformation that is part of the Backtracker that sends it to the front end. That means that with the Backtracker enabled the earlier diagram now looks like this: 5.3

This ensures that the data is sent to the frontend but fails to show which pieces of information are related to each other. Because of that, it is necessary to add watchpoints to the system. The way that Flink handles watchpoints is simple: They are injected at the beginning of the




Figure 5.3: Dataflow with the Backtracker enabled

application as part of the data stream, and each time a transformation receives one it injects the same one back at its output at the same position. Each watchpoint is identified by a unique number and is used as a grouping method in this tool. Flink does not allow the current checkpoint to be read out because there can be multiple active at the same time in different parts of the application. The only way to get the current checkpoint for a given transition is by saving the last checkpoint that was read by the transition. The `snapshotState` method is called by the Task Manager each time a new watchpoint is registered. The following line then saves the `checkpointId` as a field of the class so that the most recent watermark is always available:

```

1 @Override
2 public void snapshotState(FunctionSnapshotContext context)
3 throws Exception {
4     currentWatermark = context.getCheckpointId();
5 }
  
```

Listing 5.2: Save Watermark

When sending messages to the frontend, the current watchpoint can then be added.

Sending Data Another important step along the way is sending the data to the frontend. This is done by establishing a socket connection to the frontend (the frontend is the server) and sending serialised Java objects over it.

5.3 Frontend architecture

The front-end architecture is pretty straightforward. A view is created and populated with a tree view together with a socket server that is awaiting messages from clients. Once it receives a message, it checks if the transformation and/or watchpoint is already present in the tree and fills the data to the correct location. The server can never stop IntelliJ from working as

it is running in a separate thread, thus ensuring that any problems with the connection are isolated.

5.4 State of the program

The backtracker tool is working, that being said it is not tested in the most vigorous way possible and there are still a few features that would enhance the usability like being able to save the results.

Saving results At the moment the only way the data can be displayed is by using the IntelliJ plugin. Because the messages are sent by the internal transformation, a second transformation could be written that instead of sending the data to the IntelliJ plugin saves it in a database or logs it to a file. When creating the Backtracker, a second constructor was needed to specify which Backend function to use.

Better UI The UI still lacks a few features that would drastically increase the ease of the program like clearing old results or saving the tree directly from there.

6 Conclusion and Outlook

Summing up all the pieces, this thesis outlined how to debug complex systems, in general, showed how Flink programs are different to regular Java applications and used that knowledge to provide a methodology for debugging Flink applications. The basic procedure of a Flink application was discussed, and the pitfalls of writing these programs marked, along the thesis a sample WordCount application was used to picturise the methods. A tool was written that provides the developer with more information than are available with Flink. Namely which piece of data at which transition lead to which part of data in the next transition. This allows the developer to gain a better overview of what the program is doing compared with what he thought the program was doing.

6.1 Lessons Learned

This thesis showed a few things about debugging and writing Flink applications that were not obvious before. Bad programming is made harder because the developer does not have the freedom of writing whatever they want. Instead, they have to build using the framework which limits the possible architectural failures. The developer can still make mistakes and does so easily because the framework extends Java to be more complex which in turn makes the development process more complex but as long as the developer understands the base concept of the framework and where to put which piece of code he always has a good point of entry to the debugging process especially with the methodology in this thesis.

6.2 Outlook

There are a few things that need further research in this thesis:

1. Only the data stream part of Flink was discussed in this thesis. As Flink also offers a dataset API, it would make sense to research how the dataset API might differ in the sense of debugging. Although a lot of the same principles still hold true there, it is by no means the same.
2. The Backtracker was not experimented with in a lot of different use cases. Data about additional computing power use and more importantly computing time could be useful.
3. The functionality of the debugger can be extended as explained in [5.4](#).
4. Because of the lack of articles about good debugging in general and more specifically debugging frameworks. It might make sense to use the information gained by this thesis to write a methodology for debugging frameworks in general.

6.3 Further Reading

This section will outline a few more articles or books to read that are relevant to the thesis but were not directly mentioned.

The beginning of the thesis provided some proposals on how to minimise the debugging. It was suggested to use theorem proving or model checking. Because it is not part of the debugging process, it was mostly left out, but because it can be quite helpful for some circumstances, it makes sense to read more about it. Amazon uses model checking in some departments and found some bugs that were not known before: [Newcombe u. a. \(2015\)](#). A common language for writing these is called TLA+. It is a functional language where the core function of the program is mathematically written. A good introduction can be found here: [Lampport \(2002\)](#).

If this thesis was read not specifically for gaining an understanding about debugging Flink but for debugging distributed systems, in general, consider reading [Olston und Reed \(2011\)](#) as well. It provides some tooling for monitoring and debugging distributed systems that Flink already includes. So some of the information that are crucial in finding bugs might be found using it. Another fascinating study about missing events in distributed systems is also worth reading [Wu u. a. \(2014\)](#).

Bibliography

- [Beschastnikh u. a. 2016] BESCHASTNIKH, Ivan ; WANG, Patty ; BRUN, Yuriy ; ERNST, Michael D.: Debugging Distributed Systems. In: *Queue* 14 (2016), März, Nr. 2, S. 50:91–50:110. – URL <http://doi.acm.org/10.1145/2927299.2940294>. – ISSN 1542-7730
- [Boehm 1987] BOEHM, B. W.: Improving Software Productivity. In: *Computer* 20 (1987), Sept, Nr. 9, S. 43–57. – ISSN 0018-9162
- [Carbone u. a. 2015] CARBONE, Paris ; FÓRA, Gyula ; EWEN, Stephan ; HARIDI, Seif ; TZOUMAS, Kostas: Lightweight Asynchronous Snapshots for Distributed Dataflows. In: *CoRR* abs/1506.08603 (2015). – URL <http://arxiv.org/abs/1506.08603>
- [Dean und Ghemawat 2008] DEAN, Jeffrey ; GHEMAWAT, Sanjay: MapReduce: simplified data processing on large clusters. In: *Communications of the ACM* 51 (2008), Nr. 1, S. 107–113
- [Flink a] FLINK, Apache: *Apache Flink*. – URL <https://flink.apache.org/>. – Accessed: 2017-01-20
- [Flink b] FLINK, Apache: *Apache Flink Dataflow Figure*. – URL https://flink.apache.org/img/parallel_dataflows.png. – Accessed: 2017-01-20
- [Flink c] FLINK, Apache: *Apache Flink Simple Figure*. – URL https://ci.apache.org/projects/flink/flink-docs-master/fig/program_dataflow.svg. – Accessed: 2017-01-20
- [Flink Common Error Messages] *Apache Flink FAQ*. – URL <https://flink.apache.org/faq.html#debugging-and-common-error-messages>. – Accessed: 2017-01-20
- [FlinkSpector] *Otto Group Flink Spector*. – URL <https://github.com/ottogroup/flink-spector>. – Accessed: 2017-01-20
- [Gulzar u. a. 2016] GULZAR, Muhammad A. ; INTERLANDI, Matteo ; YOO, Seunghyun ; TETALI, Sai D. ; CONDIE, Tyson ; MILLSTEIN, Todd ; KIM, Miryung: BigDebug: Debugging Primitives

- for Interactive Big Data Processing in Spark. In: *Proceedings of the 38th International Conference on Software Engineering*. New York, NY, USA : ACM, 2016 (ICSE '16), S. 784–795. – URL <http://doi.acm.org/10.1145/2884781.2884813>. – ISBN 978-1-4503-3900-1
- [Lamport 2002] LAMPORT, Leslie: *Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers*. Boston, MA, USA : Addison-Wesley Longman Publishing Co., Inc., 2002. – ISBN 032114306X
- [McIlroy u. a. 1978] MCLROY, M.D. ; PINSON, E. N. ; TAGUE, B. A.: The Bell System Technical Journal. Bell Laboratories. Unix Time-Sharing System Forward. 57 (6, part 2). (1978), S. 1902
- [Newcombe u. a. 2015] NEWCOMBE, Chris ; RATH, Tim ; ZHANG, Fan ; MUNTEANU, Bogdan ; BROOKER, Marc ; DEARDEUFF, Michael: How Amazon Web Services Uses Formal Methods. In: *Commun. ACM* 58 (2015), März, Nr. 4, S. 66–73. – URL <http://doi.acm.org/10.1145/2699417>. – ISSN 0001-0782
- [Olston und Reed 2011] OLSTON, Christopher ; REED, Benjamin: Inspector Gadget: A Framework for Custom Monitoring and Debugging of Distributed Dataflows. In: *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data*. New York, NY, USA : ACM, 2011 (SIGMOD '11), S. 1221–1224. – URL <http://doi.acm.org/10.1145/1989323.1989459>. – ISBN 978-1-4503-0661-4
- [Wu u. a. 2014] WU, Yang ; ZHAO, Mingchen ; HAEBERLEN, Andreas ; ZHOU, Wenchao ; LOO, Boon T.: Diagnosing Missing Events in Distributed Systems with Negative Provenance. In: *SIGCOMM Comput. Commun. Rev.* 44 (2014), August, Nr. 4, S. 383–394. – URL <http://doi.acm.org/10.1145/2740070.2626335>. – ISSN 0146-4833
- [Zeller 2009] ZELLER, Andreas: *Why Programs Fail, Second Edition: A Guide to Systematic Debugging*. 2nd. San Francisco, CA, USA : Morgan Kaufmann Publishers Inc., 2009. – ISBN 0123745152, 9780123745156

Hiermit versichere ich, dass ich die vorliegende Arbeit ohne fremde Hilfe selbständig verfasst und nur die angegebenen Hilfsmittel benutzt habe.

Hamburg, 2. February 2018

 Jakob Joachim