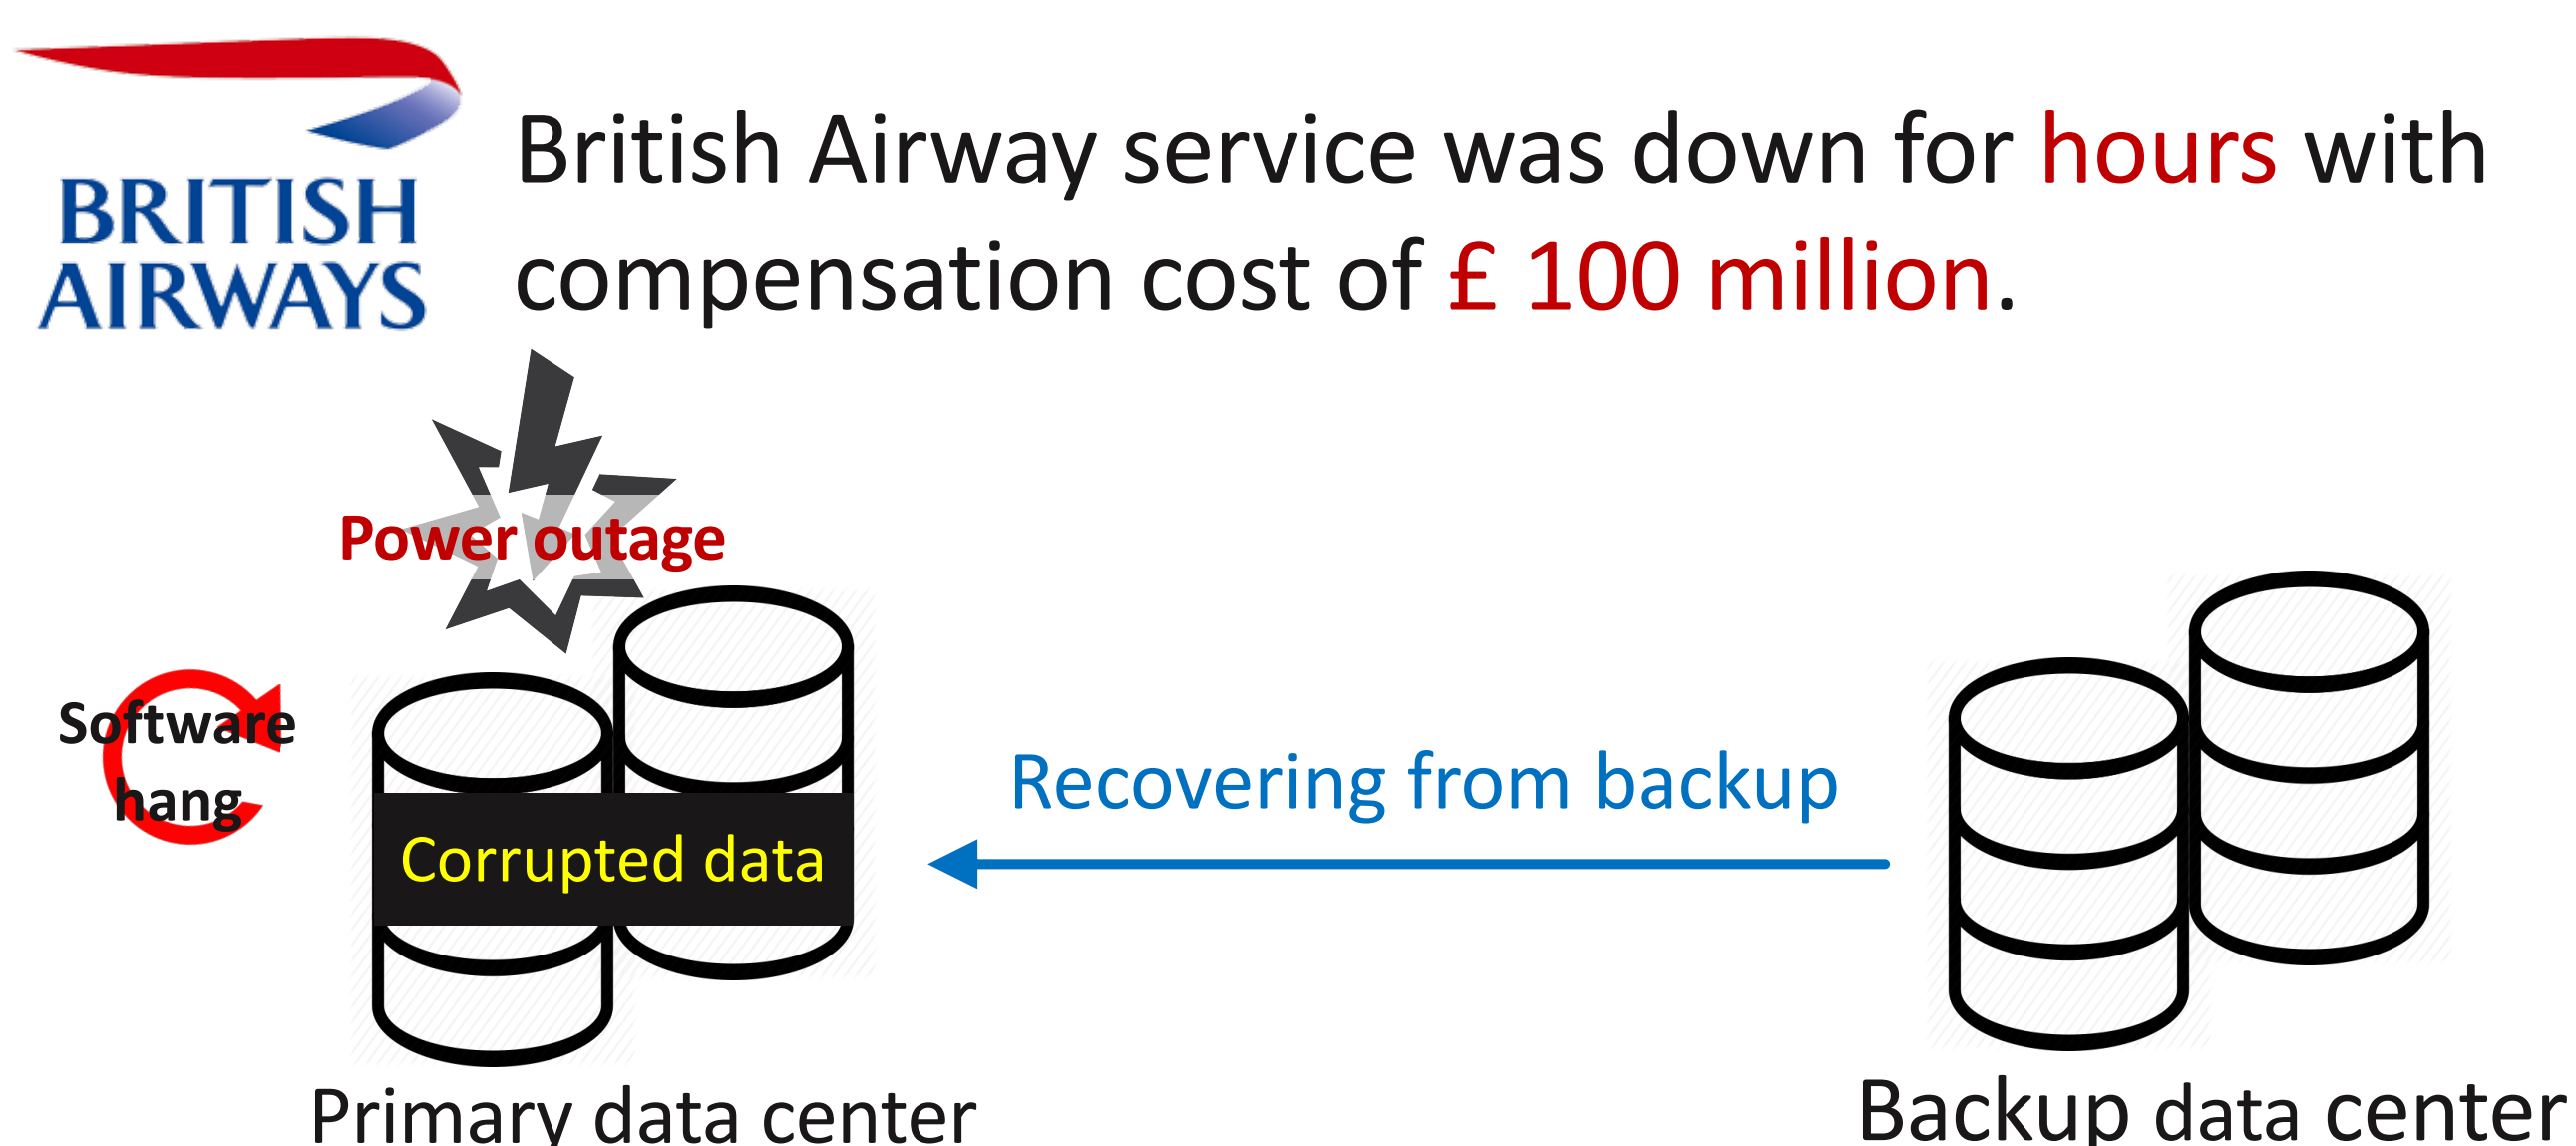


Motivation



Contributions

- DScope**: a static analysis tool to detect data-corruption related hang bugs.
 - Detecting potential infinite loops caused by data corruptions;
 - Pruning always-exit loops through loop stride and bound analysis;
 - Categorizing real-world data-corruption hang bugs into four common types.

Overview

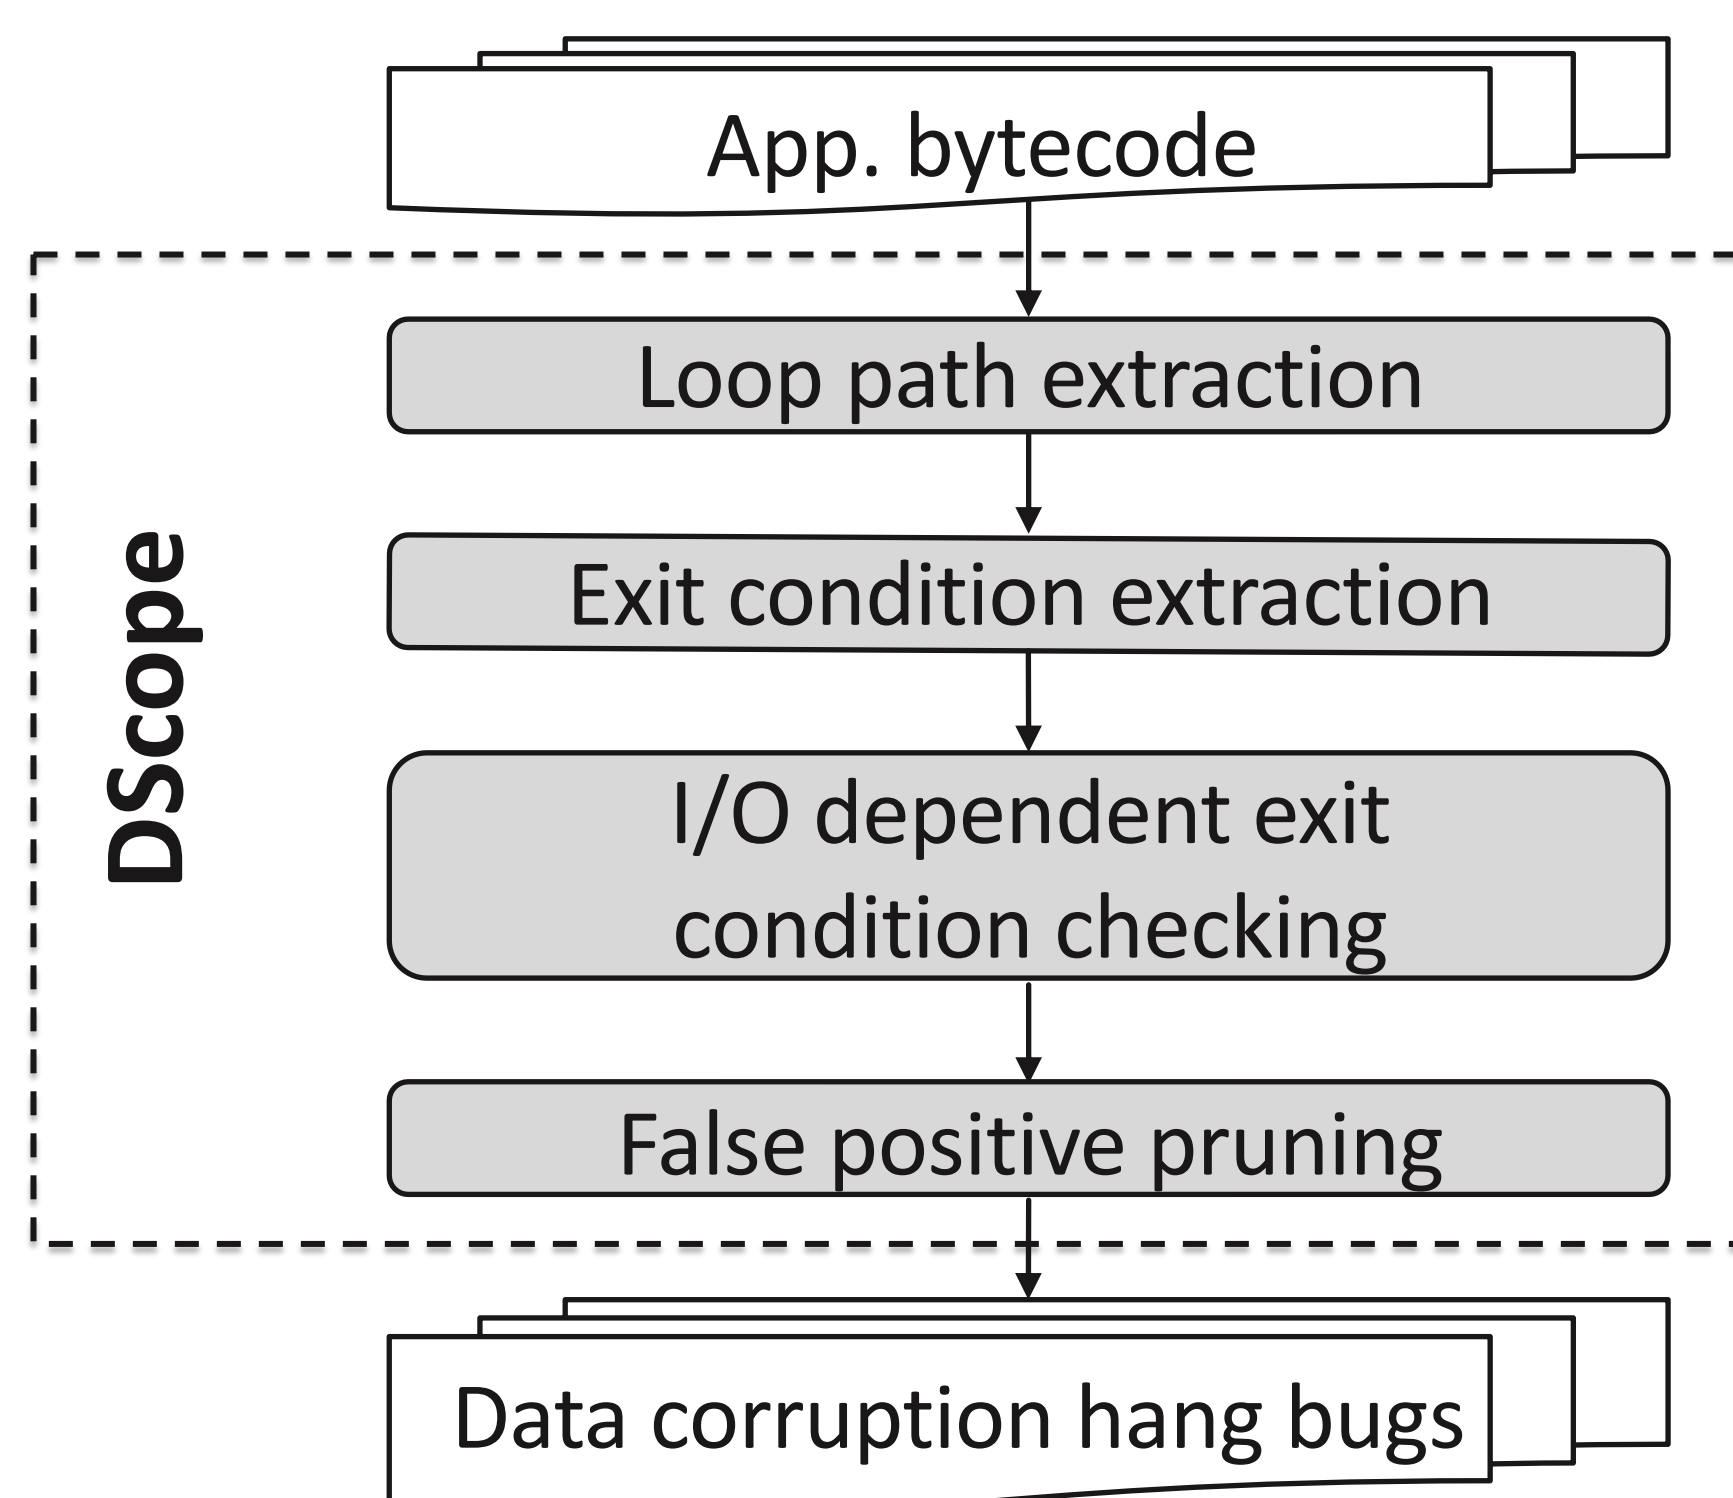
Example: Hadoop-8614

```

183 public static void skipFully(InputStream in, long len) throws IOException {
184     while (len > 0) {
185         long ret = in.skip(len);
186         ...
189         len -= ret;
    }
}

```

The loop stride (ret) is always 0 when in is corrupted.



False Positive Filtering

```

307 public static long readVLong(DataInput stream) throws IOException {
308     byte firstByte = stream.readByte();
309     int len = decodeVIntSize(firstByte);
310     ...
314     for (int idx = 0; idx < len-1; idx++) {
315         ...
318     }
319 }

```

len is I/O dependent

It's a FP b/c the loop stride is always 1 and the upper bound (len-1) is fixed.

False positive conditions:

- The loop stride is always positive when the loop index has a fixed upper bound;
- The loop stride is always negative when the loop index has a fixed lower bound.

Data Corruption Hang Bug Types

1 Error codes returned by I/O operations directly affect loop strides.
e.g., Hadoop-8614

2 Corrupted data content indirectly affects loop strides.

Example: HDFS-13514

```

194 BUFFER_SIZE = conf.getInt();
78 private void readLocalFile(Path path, ...) throws IOException {
79     ...
84     byte[] data = new byte[BUFFER_SIZE];
85     long size = 0;
86     while (size >= 0) {
87         size = in.read(data);
88     }
}

```

Corrupted configuration file

empty array

The loop stride (size) is always 0 when conducting read op on an empty array.

3 Improper exception handling directly affects loop strides.

Example: Yarn-6991

```

64 try {
73     //create pidFile
79 } catch (IOException ioe) {
80     LOG.info("...failed...");
81 }
88 File f = new File(pidFile);
89 while (!f.exists()) {
91     Thread.sleep(500);
99 }

```

throw exception

silent failure

The silent failure makes the loop exit condition always be false.

Thread #1

Thread #2

4 Improper exception handling indirectly affects loop strides.

Example: Cassandra-9881

```

120 while (!dataFile.isEOF()) {
129     try {
130         decorateKey(ByteBufferUtil.readWithShortLength(dataFile));
134         dataSize = dataFile.readLong();
139     } catch (Throwable th) {
140         //ignore exception
    }
}

```

throw exception

The improper exception handling makes the loop stride always be 0 (readLong API is skipped).

Detection Results

#	Bug name	System version	Bug type	Known or new	Detected			#	Bug name	System version	Bug type	Known or new	Detected		
					DScope	Findbugs	Infer						DScope	Findbugs	Infer
1	Cassandra-7330	2.0.8	#1	known	✓	x	x	22	Mapreduce-7088	2.5	#1	new	✓	x	x
2	Cassandra-9881	2.0.8	#3	known	✓	x	✓	23	Mapreduce-7089	2.5	#1	new	✓	x	x
3	Compress-87	1.0	#1	known	✓	x	x	24	Yarn-163	0.23	#1	known	✓	✓	x
4	Compress-451	1.0	#2	new	✓	x	x	25	Yarn-2905	2.5	#1	known	✓	x	x
5	Hadoop-8614	0.23	#1	known	✓	x	x	26	Yarn-6991	0.23	#4	new	✓	x	x
6	Hadoop-15088	2.5	#1	new	✓	x	x	27	Yarn-6991	2.5	#4	new	✓	x	x
7	Hadoop-15415	0.23	#2	new	✓	x	x	28	Hive-5235	1.0	#1	known	✓	x	x
8	Hadoop-15417	2.5	#2	new	✓	x	x	29	Hive-13397	1.0	#2	known	✓	x	x
9	Hadoop-15424	0.23	#2	new	✓	x	x	30	Hive-18142	1.0	#2	new	✓	x	x
10	Hadoop-15424	2.5	#1	new	✓	x	x	31	Hive-18216	2.3.2	#1	new	✓	x	x
11	Hadoop-15425	2.5	#1	new	✓	x	x	32	Hive-18217	2.3.2	#1	new	✓	x	x
12	Hadoop-15429	0.23	#2	new	✓	x	x	33	Hive-18219	1.0	#2	new	✓	x	x
13	HDFS-4882	0.23	#3	known	✓	x	x	34	Hive-18219	2.3.2	#2	new	✓	x	x
14	HDFS-5892	2.5	#2	known	✓	✓	x	35	Hive-19391	1.0	#2	new	✓	x	x
15	HDFS-13513	2.5	#2	new	✓	x	x	36	Hive-19392	1.0	#2	new	✓	x	x
16	HDFS-13514	2.5	#2	new	✓	x	x	37	Hive-19392	2.3.2	#2	new	✓	x	x
17	Mapreduce-2185	0.23	#2	known	✓	x	x	38	Hive-19395	1.0	#1	new	✓	x	x
18	Mapreduce-2862	0.23	#2	known	✓	x	x	39	Hive-19406	2.3.2	#2	new	✓	x	x
19	Mapreduce-6990	0.23	#1	new	✓	x	x	40	Kafka-6271	0.10	#1	new	✓	x	x
20	Mapreduce-6990	0.23	#1	new	✓	x	x	41	Lucene-772	2.1	#2	new	✓	x	x
21	Mapreduce-6990	0.23	#1	new	✓	x	x	42	Lucene-8294	2.1	#2	known	✓	x	x

- DScope identifies 42 true data corruption hang bugs including 29 new bugs.
- Findbugs identifies 2 true data corruption hang bugs.
- Infer identifies 1 true data corruption hang bug.

Conclusion

- DScope is a new data corruption hang bug detection tool for cloud server systems.
 - Combines candidate bug discovery and false positive pattern filtering.
 - Evaluated over 9 cloud server systems and detects 42 true corruption hang bugs including 29 new bugs.