

Assignment 2: MapReduce

1 Set-Up

To obtain large text files to test the program with, I downloaded the first 10 long books I could think of from archive.org in txt file form. These were:

1. The Bible;
2. *War & Peace* by Leo Tolstoy;
3. Plutarch's *Lives*;
4. Herodotus' *Histories*;
5. *City of God* by Augustine of Hippo;
6. *Faust* by Goethe;
7. *Wealth of Nations* by Adam Smith;
8. *Capital* by Karl Marx;
9. The complete works of William Shakespeare;
10. *Structure & Interpretation of Computer Programs* by Harold Abelson & Gerald Jay Sussman.

2 Baseline Results

I modified the code to measure & output the time taken by each approach, in milliseconds. I also added timing for the different phases of the two MapReduce implementations, timing the map time, group time, and reduce time separately.

```
[andrew@arch] ~/currsem/CT414/assignments/assignment2/code P [root@arch] +  
% javac *.java && java MapReduceFiles ../data/augustine_city_of_god.txt ../data/Faust-Goethe.txt ../data/herodotus_histories.txt ../data/holy_bible.txt ../data/marx_capital.txt ../data/plutarch_lives.txt ../da  
ta/shakespeare_complete_works.txt ../data/structure_and_interpretation_of_computer_programs.txt ../data/war_and_peace.txt ../data/wealth_of_nations.txt  
Brute Force Results:  
  Total Time: 1295  
MapReduce Results:  
  Map Time: 1387  
  Group Time: 682  
  Reduce Times: 414  
  Total Time: 2483  
Distributed MapReduce Results:  
  Map Time: 726  
  Group Time: 726  
  Reduce Time: 31813  
  Total Time: 31813  
[andrew@arch] ~/currsem/CT414/assignments/assignment2/code P [root@arch] +  
% □
```

Figure 1: Baseline results for my list of files (in milliseconds)

As can be seen from the above terminal screenshot, the brute force approach performed best with no modifications, followed by the non-distributed MapReduce, followed by the distributed MapReduce; this is to be expected, as the brute force approach is the simplest & requires the fewest iterations over the data and no complex data structures. The non-distributed MapReduce requires more intermediate data structure and more iterations over the data. Finally, the non-optimised version of the distributed MapReduce is the slowest because it spawns a thread for each word in the dataset, causing massive stress on the CPU and memory.

I also updated the code to use ArrayLists rather than LinkedLists to reduce memory overhead and have faster traversal.

```
[andrew@arch] ~/currsem/CT414/assignments/assignment2/code P [root@arch] +  
% javac *.java && java MapReduceFiles ../data/augustine_city_of_god.txt ../data/Faust-Goethe.txt ../data/herodotus_histories.txt ../data/holy_bible.txt ../data/marx_capital.txt ../data/plutarch_lives.txt ../da  
ta/shakespeare_complete_works.txt ../data/structure_and_interpretation_of_computer_programs.txt ../data/war_and_peace.txt ../data/wealth_of_nations.txt  
Brute Force Results:  
  Total Time: 1387  
MapReduce Results:  
  Map Time: 623  
  Group Time: 547  
  Reduce Times: 202  
  Total Time: 1372  
Distributed MapReduce Results:  
  Map Time: 374  
  Group Time: 388  
  Reduce Time: 31874  
  Total Time: 31874  
[andrew@arch] ~/currsem/CT414/assignments/assignment2/code P [root@arch] +  
% □
```

Figure 2: Baseline results with ArrayList update (in milliseconds)

As can be seen from the above terminal screenshot, this has no affect on the brute force results (besides slight variance due to background processes running on my laptop) as this approach did not use LinkedLists anyway. The non-distributed MapReduce approach was significantly faster due to the faster iteration and lower memory overhead. The distributed MapReduce saw significant improvements in the map & group phases, but these were dwarfed by the still greatly inefficient reduce phase.

3 Testing the Updated Code

After implementing the requested changes in steps 2–6 of the assignment specification, I then implemented a grid-search function which tested a range of values for the number of lines of text per map thread and the number of words per reduce thread. The results of this grid-search were exported to a CSV file for analysis. I then wrote a Python script to visualise the parameter combinations using heatmaps. Heatmaps that contain results pertaining only to the map phase and the reduce phase, as well as a table of the results from the CSV file can be found in the Appendix.

```

~/Downloads/CS141/assignments/assignment3/code/3
$ java -jar MapReduceF115 .../data/jeanpierre_city_of_pod.txt .../data/faust-Goethe.txt .../data/herodotus_histories.txt .../data/holy_bible.txt .../data/marx_capital.txt .../data/plutarch_lives.txt .../data/shakespeare_complete_works.txt .../data/structure_and_interpretation_of_computer_programs.txt .../data/war_and_peace.txt .../data/wealth_of_nations.txt && python3 plots.py
INFO: Starting Grid Search - test:
MapLines: 1000, ReduceWords: 100
Map Time: 1534 ms
Group Time: 253 ms
Reduce Time: 230 ms
Total Time: 2033 ms
-----
MapLines: 1000, ReduceWords: 150
Map Time: 1422 ms
Group Time: 352 ms
Reduce Time: 145 ms
Total Time: 1689 ms
-----
MapLines: 1000, ReduceWords: 200
Map Time: 1129 ms
Group Time: 243 ms
Reduce Time: 69 ms
Total Time: 1421 ms
-----
MapLines: 1000, ReduceWords: 500
Map Time: 1103 ms
Group Time: 281 ms
Reduce Time: 56 ms
Total Time: 1485 ms
-----
MapLines: 1000, ReduceWords: 1000
Map Time: 1425 ms
Group Time: 227 ms
Reduce Time: 69 ms
Total Time: 1731 ms
-----
MapLines: 2000, ReduceWords: 100
Map Time: 1184 ms
Group Time: 290 ms
Reduce Time: 150 ms
Total Time: 1573 ms
-----
MapLines: 2000, ReduceWords: 150
Map Time: 1158 ms
Group Time: 303 ms
Reduce Time: 59 ms
Total Time: 1520 ms
-----
MapLines: 2000, ReduceWords: 200
Map Time: 1216 ms
Group Time: 209 ms
Reduce Time: 46 ms
Total Time: 1531 ms
-----
MapLines: 2000, ReduceWords: 500
Map Time: 1292 ms

```

Figure 3: Running the grid-search and plotting the results

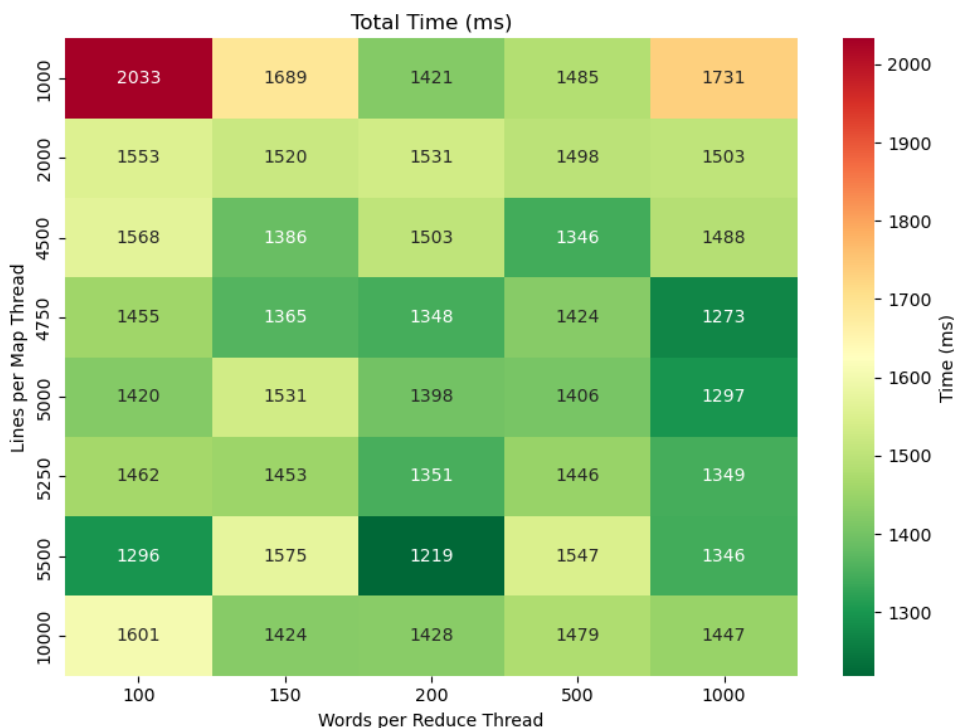


Figure 4: Heatmap of total time taken by each parameter combination

As can be seen from the heatmap above, the best result achieved by the distributed MapReduce approach was 1,219 milliseconds

onds with the parameter combination of 200 words per reduce thread and 5,500 lines per map thread. Not only is this an improvement of over 2600% ($\frac{31874}{1219} \times 100 = 2614.766$), but it beats the original brute force result of 1,307 milliseconds and the non-distributed MapReduce result of 1,372 milliseconds. We can tell that this isn't a fluke, as this is consistent with other neighbouring values in the heatmap, with similar times of 1,296 milliseconds, 1,273 milliseconds, etc.

What I find most interesting about these results is that they are so similar to the brute force results; while the fully-optimised & tuned distributed MapReduce does beat brute force, it only does so by a narrow margin. I have a few ideas as to why this might be the case: it may be due to the brute force approach being more suitable to CPU caching due to its sequential approach, or due to the high thread overhead running it on the CPU of a single laptop, but I think that the main reason is that the dataset isn't actually very big. While 10 large books may seem to be quite a lot of data, MapReduce was created to deal with petabytes of data, not a few megabytes. I would hypothesise that the real performance benefits of distributed MapReduce would only become clear if the testing was repeated on at least a couple gigabytes of data.

4 Appendix: Source Code

```
1 import java.util.*;
2 import java.io.*;
3
4 public class MapReduceFiles {
5
6     private static final String CSV_FILE = "performance_results.csv";
7
8     public static void main(String[] args) {
9         if (args.length < 1) {
10             System.err.println("Usage: java MapReduceFiles file1.txt file2.txt ... fileN.txt");
11             return;
12         }
13
14         Map<String, String> input = new HashMap<>();
15         try {
16             for (String filename : args) {
17                 input.put(filename, readFile(filename));
18             }
19         } catch (IOException ex) {
20             System.err.println("Error reading files: " + ex.getMessage());
21             ex.printStackTrace();
22             return;
23         }
24
25         int[] mapSizes = {1000, 2000, 4500, 4750, 5000, 5250, 5500, 10000};
26         int[] reduceSizes = {100, 150, 200, 500, 1000};
27
28         System.out.println("==== Starting Grid Search =====");
29
30         try (PrintWriter writer = new PrintWriter(new FileWriter(CSV_FILE))) {
31             writer.println("MapLines,ReduceWords,MapTime,GroupTime,ReduceTime,TotalTime");
32
33             for (int mapSize : mapSizes) {
34                 for (int reduceSize : reduceSizes) {
35                     runDistributedMapReduce(input, mapSize, reduceSize, writer);
36                 }
37             }
38
39         } catch (IOException e) {
40             System.err.println("Error writing to CSV file: " + e.getMessage());
41         }
42
43         System.out.println("==== Grid Search Complete =====");
44         System.out.println("Results saved to: " + CSV_FILE);
45     }
46 }
```

```

47 public static void runDistributedMapReduce(Map<String, String> input, int linesPerMapThread, int
↪ wordsPerReduceThread, PrintWriter csvWriter) {
48     final Map<String, Map<String, Integer>> output = new HashMap<>();
49
50     // MAP Phase
51     Long mapStartTime = System.currentTimeMillis();
52     List<MappedItem> mappedItems = Collections.synchronizedList(new ArrayList<>());
53
54     final MapCallback<String, MappedItem> mapCallback = new MapCallback<>() {
55         public synchronized void mapDone(String file, List<MappedItem> results) {
56             mappedItems.addAll(results);
57         }
58     };
59
60     List<Thread> mapCluster = new ArrayList<>();
61     for (Map.Entry<String, String> entry : input.entrySet()) {
62         final String file = entry.getKey();
63         final String[] lines = entry.getValue().split("\\r?\\n");
64
65         for (int i = 0; i < lines.length; i += linesPerMapThread) {
66             int end = Math.min(i + linesPerMapThread, lines.length);
67             final List<String> chunk = new ArrayList<>();
68             for (int j = i; j < end; j++) {
69                 chunk.addAll(splitLongLine(lines[j]));
70             }
71
72             Thread t = new Thread(() -> map(file, chunk, mapCallback));
73             mapCluster.add(t);
74             t.start();
75         }
76     }
77
78     for (Thread t : mapCluster) {
79         try {
80             t.join();
81         } catch (InterruptedException e) {
82             throw new RuntimeException(e);
83         }
84     }
85
86     Long mapTotalTime = System.currentTimeMillis() - mapStartTime;
87
88     // GROUP Phase
89     Long groupStartTime = System.currentTimeMillis();
90     Map<String, List<String>> groupedItems = new HashMap<>();
91     for (MappedItem item : mappedItems) {
92         groupedItems.computeIfAbsent(item.getWord(), k -> new ArrayList<>()).add(item.getFile());
93     }
94     Long groupTotalTime = System.currentTimeMillis() - groupStartTime;
95
96     // REDUCE Phase
97     Long reduceStartTime = System.currentTimeMillis();
98     final ReduceCallback<String, String, Integer> reduceCallback = (word, result) -> {
99         synchronized (output) {
100             output.put(word, result);
101         }
102     };
103
104     List<Thread> reduceCluster = new ArrayList<>();
105     List<Map<String, List<String>>> reduceChunks = new ArrayList<>();
106     Map<String, List<String>> currentChunk = new HashMap<>();

```

```

107     int count = 0;
108
109     for (Map.Entry<String, List<String>> entry : groupedItems.entrySet()) {
110         currentChunk.put(entry.getKey(), entry.getValue());
111         count++;
112         if (count >= wordsPerReduceThread) {
113             reduceChunks.add(currentChunk);
114             currentChunk = new HashMap<>();
115             count = 0;
116         }
117     }
118     if (!currentChunk.isEmpty()) reduceChunks.add(currentChunk);
119
120     for (final Map<String, List<String>> chunk : reduceChunks) {
121         Thread t = new Thread(() -> {
122             for (Map.Entry<String, List<String>> entry : chunk.entrySet()) {
123                 reduce(entry.getKey(), entry.getValue(), reduceCallback);
124             }
125         });
126         reduceCluster.add(t);
127         t.start();
128     }
129
130     for (Thread t : reduceCluster) {
131         try {
132             t.join();
133         } catch (InterruptedException e) {
134             throw new RuntimeException(e);
135         }
136     }
137
138     long reduceTotalTime = System.currentTimeMillis() - reduceStartTime;
139     long totalTime = mapTotalTime + groupTotalTime + reduceTotalTime;
140
141     // Print & Log
142     System.out.println("MapLines: " + linesPerMapThread + ", ReduceWords: " + wordsPerReduceThread);
143     System.out.println("\tMap Time: " + mapTotalTime + " ms");
144     System.out.println("\tGroup Time: " + groupTotalTime + " ms");
145     System.out.println("\tReduce Time: " + reduceTotalTime + " ms");
146     System.out.println("\tTotal Time: " + totalTime + " ms");
147     System.out.println("-----");
148
149     csvWriter.printf("%d,%d,%d,%d,%d,%d%n",
150         linesPerMapThread, wordsPerReduceThread,
151         mapTotalTime, groupTotalTime, reduceTotalTime, totalTime);
152     csvWriter.flush();
153 }
154
155 public static void map(String file, List<String> lines, MapCallback<String, MappedItem> callback) {
156     List<MappedItem> results = new ArrayList<>();
157     for (String line : lines) {
158         String[] words = line.trim().split("\\s+");
159         for (String word : words) {
160             word = word.replaceAll("[^a-zA-Z]", "").toLowerCase();
161             if (!word.isEmpty()) {
162                 results.add(new MappedItem(word, file));
163             }
164         }
165     }
166     callback.mapDone(file, results);
167 }

```

```

168
169 public static void reduce(String word, List<String> list, ReduceCallback<String, String, Integer>
↪ callback) {
170     Map<String, Integer> reducedList = new HashMap<>();
171     for (String file : list) {
172         reducedList.put(file, reducedList.getOrDefault(file, 0) + 1);
173     }
174     callback.reduceDone(word, reducedList);
175 }
176
177 public interface MapCallback<E, V> {
178     void mapDone(E key, List<V> values);
179 }
180
181 public interface ReduceCallback<E, K, V> {
182     void reduceDone(E e, Map<K, V> results);
183 }
184
185 private static class MappedItem {
186     private final String word;
187     private final String file;
188
189     public MappedItem(String word, String file) {
190         this.word = word;
191         this.file = file;
192     }
193
194     public String getWord() {
195         return word;
196     }
197
198     public String getFile() {
199         return file;
200     }
201
202     @Override
203     public String toString() {
204         return "[" + word + ", " + file + "]";
205     }
206 }
207
208 private static String readFile(String pathname) throws IOException {
209     File file = new File(pathname);
210     StringBuilder fileContents = new StringBuilder((int) file.length());
211     Scanner scanner = new Scanner(new BufferedReader(new FileReader(file)));
212     String lineSeparator = System.getProperty("line.separator");
213
214     try {
215         while (scanner.hasNextLine()) {
216             fileContents.append(scanner.nextLine()).append(lineSeparator);
217         }
218         return fileContents.toString();
219     } finally {
220         scanner.close();
221     }
222 }
223
224 private static List<String> splitLongLine(String line) {
225     List<String> result = new ArrayList<>();
226     while (line.length() > 80) {
227         int splitAt = line.lastIndexOf(' ', 80);

```

```

228     if (splitAt <= 0) splitAt = 80;
229     result.add(line.substring(0, splitAt));
230     line = line.substring(splitAt).trim();
231 }
232 if (!line.isEmpty()) result.add(line);
233 return result;
234 }
235 }

```

Listing 1: MapReduceFiles.java

```

1 import pandas as pd
2 import seaborn as sns
3 import matplotlib.pyplot as plt
4
5 df = pd.read_csv('performance_results.csv')
6
7 def save_heatmap(metric, title, filename):
8     pivot = df.pivot(index='MapLines', columns='ReduceWords', values=metric)
9     plt.figure(figsize=(8, 6))
10    sns.heatmap(
11        pivot,
12        annot=True,
13        fmt="d",
14        cmap="RdYlGn_r",
15        cbar_kws={'label': 'Time (ms)'}
16    )
17    plt.title(title)
18    plt.ylabel("Lines per Map Thread")
19    plt.xlabel("Words per Reduce Thread")
20    plt.tight_layout()
21    plt.savefig(filename)
22    plt.close()
23    print(f"Saved: {filename}")
24
25 save_heatmap('TotalTime', 'Total Time (ms)', '../latex/images/total_time_heatmap.png')
26 save_heatmap('MapTime', 'Map Time (ms)', '../latex/images/map_time_heatmap.png')
27 save_heatmap('ReduceTime', 'Reduce Time (ms)', '../latex/images/reduce_time_heatmap.png')

```

Listing 2: plots.py

Map Lines	Reduce Words	Map Time (ms)	Group Time (ms)	Reduce Time (ms)	Total Time (ms)
1000	100	1534	269	230	2033
1000	150	1192	352	145	1689
1000	200	1129	243	49	1421
1000	500	1168	281	36	1485
1000	1000	1425	237	69	1731
2000	100	1104	290	159	1553
2000	150	1158	303	59	1520
2000	200	1216	269	46	1531
2000	500	1202	260	36	1498
2000	1000	1202	264	37	1503
4500	100	1145	252	171	1568
4500	150	1089	238	59	1386
4500	200	995	294	214	1503
4500	500	863	295	188	1346
4500	1000	1183	250	55	1488
4750	100	1016	267	172	1455
4750	150	1078	228	59	1365
4750	200	1041	260	47	1348
4750	500	1110	278	36	1424
4750	1000	975	263	35	1273
5000	100	1069	277	74	1420
5000	150	1253	224	54	1531
5000	200	874	306	218	1398
5000	500	1118	252	36	1406
5000	1000	1006	244	47	1297
5250	100	1165	225	72	1462
5250	150	1008	272	173	1453
5250	200	1054	250	47	1351
5250	500	1134	275	37	1446
5250	1000	1027	280	42	1349
5500	100	976	249	71	1296
5500	150	1062	285	228	1575
5500	200	882	290	47	1219
5500	500	1254	257	36	1547
5500	1000	999	308	39	1346
10000	100	1167	360	74	1601
10000	150	1093	270	61	1424
10000	200	1075	296	57	1428
10000	500	1161	279	39	1479
10000	1000	1154	255	38	1447

Table 1: Results written to performance_results.csv

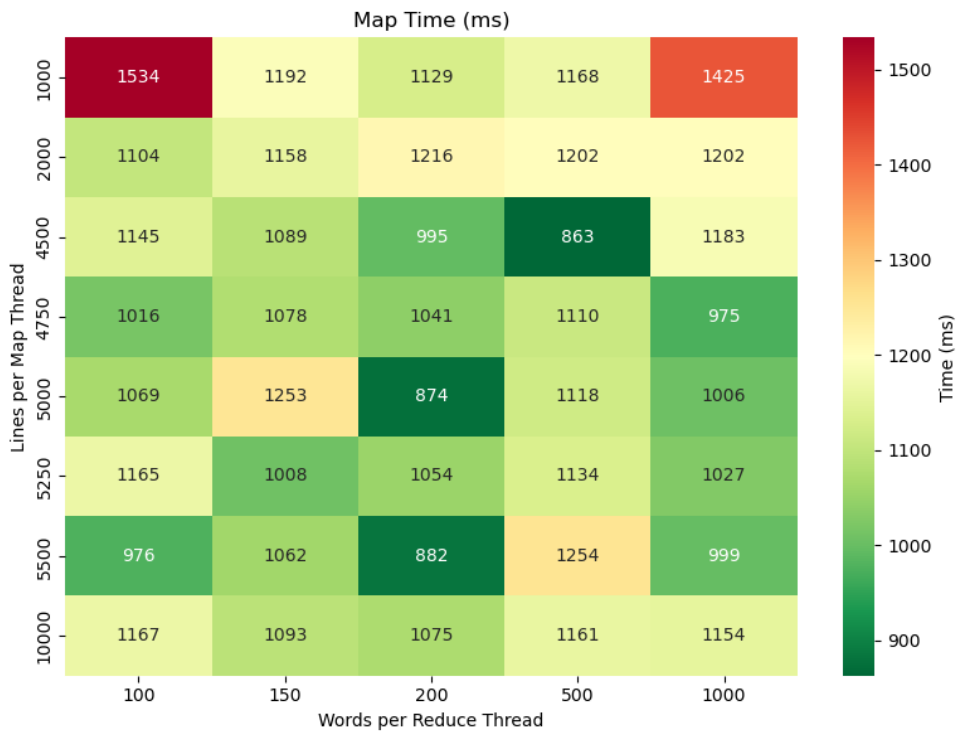


Figure 5: Heatmap of time taken during the map phase by each parameter combination

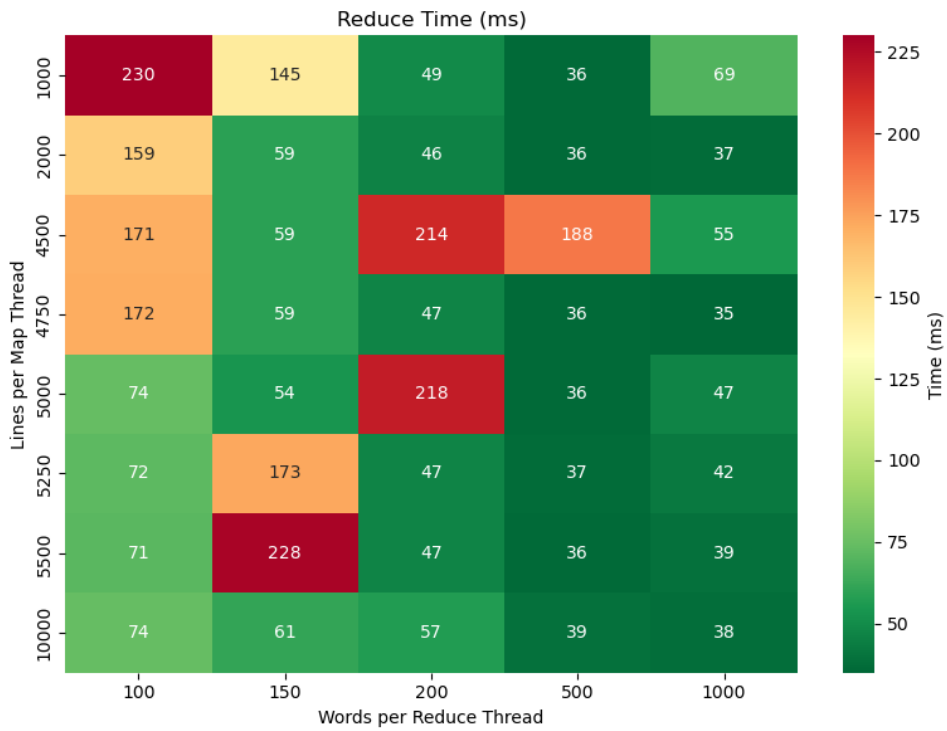


Figure 6: Heatmap of time taken during the reduce phase by each parameter combination