

Performance comparison of simplification algorithms for polygons in the context of web applications



Masterarbeit
Institut für Informatik
Universität Augsburg

vorgelegt von

Alfred Melch

Matrikelnummer xxx

Augsburg, August 2019

1. Gutachter: Prof. Dr. Jörg Hähner
2. Gutachter: Prof. Dr. Sabine Timpf
Betreuer: Prof. Dr. Jörg Hähner

Abstract

Abstract goes here

Contents

1	Introduction	1
1.1	Binary instruction sets on the web platform	1
1.2	Performance as important factor for web applications	1
1.3	Topology simplification for rendering performance	2
1.4	Related work	2
1.5	Structure of this thesis	2
2	Theory	4
2.1	Generalization in cartography	4
2.1.1	Goals of reducing data	4
2.1.2	Automated generalization	4
2.2	Geodata formats on the Web	4
2.3	Polyline simplification	7
2.3.1	Summary	8
2.4	Web runtimes	8
2.4.1	Introduction to Webassembly	8
3	Methodology	11
3.1	State of the art: Simplify.js	11
3.2	The webassembly solution	12
3.3	The implementation of a web framework	16
3.3.1	External libraries	16
3.3.2	The application logic	17
3.3.3	Benchmark cases and chart types	17
3.3.4	The different benchmark types	18
3.3.5	The benchmark suite	19
3.3.6	The user interface	20
3.4	The test data	22
4	Results	24
4.1	Case 1 - Windows - wasm vs js	24
4.2	Case 2 - Windows - wasm runtime analysis	27
4.3	Case 3 - MacBook Pro - wasm vs js	29
4.4	Case 4 - Ubuntu - turf.js analysis	32
4.5	Case 5 - iPhone - mobile testing	34

5	Discussion	35
5.1	Browser differences for the JavaScript implementations	35
5.2	Browser differences for Simplify.wasm	36
5.3	Insights into Simplify.wasm	36
5.4	Comparison Simplify.wasm vs Simplify.js	37
5.5	Analysis of Turf.js implementation	37
5.6	Mobile device analysis	37
6	Conclusion	38
6.1	Enhancements	38
6.2	Future Work	38
7	Practical application	39
7.1	State of the art: psimpl	39
7.2	Compiling to WebAssembly	39
7.3	The implementation	41
7.4	The user interface	43

1 Introduction

Simplification of polygonal data structures is the task of reducing data points while preserving topological characteristics. The simplification often takes the form of removing points that make up the geometry. There are several solutions that tackle the problem in different ways. This thesis aims to compare and classify these solutions by various heuristics. Performance and compression rate are quantitative heuristic used. Positional, length and area errors will also be measured to quantify simplification errors. With the rising trend of moving desktop applications to the web platform also geographic information systems (GIS) have experienced the shift towards web browsers ¹. Performance is critical in these applications. Since simplification is an important factor to performance the solutions will be tested by constructing a web application using a technology called WebAssembly.

1.1 Binary instruction sets on the web platform

The recent development of WebAssembly allows code written in various programming languages to be run natively in web browsers. So far JavaScript was the only native programming language on the web. The goals of WebAssembly are to define a binary instruction format as a compilation target to execute code at native speed and taking advantage of common hardware capabilities ². The integration into the web platform brings portability to a wide range of platforms like mobile and internet of things (IoT). The usage of this technology promises performance gains that will be tested by this thesis. The results can give conclusions to whether WebAssembly is worth a consideration for web applications with geographic computational aspects. Web GIS is an example technology that would benefit greatly of such an advancement. Thus far WebAssembly has been shipped to the stable version of the four most used browser engines ³. The mainly targeted high-level languages for compilation are C and C++. Also a compiler for Rust and a TypeScript subset has been developed. It will be explored how existing implementations could easily be adopted when using a compiler.

1.2 Performance as important factor for web applications

There has been a rapid growth of complex applications running in web-browsers. These so called progressive web apps (PWA) combine the fast reachability of web pages with

¹<https://www.esri.com/about/newsroom/arcnews/implementing-web-gis/>

²<https://webassembly.org/>

³<https://lists.w3.org/Archives/Public/public-webassembly/2017Feb/0002.html>

the feature richness of locally installed applications. Even though these applications can grow quite complex, the requirement for fast page loads and instant user interaction still remains. One way to cope this need is the use of compression algorithms to reduce the amount of data transmitted and processed. In a way simplification is a form of data compression. Web servers use lossless compression algorithms like gzip to deflate data before transmission. Browsers that implement these algorithms can then fully restore the requested resources resulting in lower bandwidth usage. The algorithms presented here however remove information from the data in a way that cannot be restored. This is called lossy compression. The most common usage on the web is the compression of image data.

1.3 Topology simplification for rendering performance

While compression is often used to minimize bandwidth usage the compression of geospatial data can particularly influence rendering performance. The bottleneck for rendering often is the svg transformation used to display topology on the web. Implementing simplification algorithms for use on the web platform can lead to smoother user experience when working with large geodata sets.

1.4 Related work

Related Work

1.5 Structure of this thesis

This thesis is structured into a theoretical and a practical component. First the theoretical principles will be reviewed. Topology of polygonal data will be explained as how to describe geodata on the web. A number of algorithms will be introduced in this section. Each algorithm will be dissected by complexity, characteristics and the possible influence to the heuristics mentioned above. An introduction to WebAssembly will be given here.

In the next chapter the practical implementation will be presented. This section is divided in two parts since two web applications are produced in this thesis. The first one is a benchmark comparison of an algorithm implemented in JavaScript and in WebAssembly. It will be used to investigate if performance of established implementations can be improved by a new technology. The second part is about several algorithms brought to the web by compiling an existing C++ library. This application can be used for qualitative analysis of the algorithms. It will show live results to see the characteristics and influence of single parameters.

The results of the above methods will be shown in chapter 4. After discussion of the results a conclusion will finish the thesis.

2 Theory

In this chapter the theory behind polygon simplification will be explained. The simplification process is part of generalization in cartography. So first a few words about it will be dropped to give a broad overview about the topic. It will be clarified which goals drive the reducing of data quantity, especially in the context of web applications. Then the data formats will be explained that make up the data. From there a closer look can be taken how the simplification can be accomplished.

2.1 Generalization in cartography

2.1.1 Goals of reducing data

2.1.2 Automated generalization

2.2 Geodata formats on the Web

Here the data formats that are used through this theses will be explained.

The JavaScript Object Notation (JSON) Data Interchange Format was derived from the ECMAScript Programming Language Standard⁴. It is a text format for the serialization of structured data. As a text format is suitable well for the data exchange between server and client. Also it can easily be consumed by JavaScript. These characteristics are ideal for web based applications. It does however only support a limited number of data types. Four primitive ones (string, number, boolean and null) and two structured ones (objects and array). Objects are an unordered collection of name-value pairs, while arrays are simply ordered lists of values. JSON was meant as a replacement for XML as it provides a more human readable format. Through nesting complex data structures can be created.

The GeoJSON Format is a geospatial data interchange format⁵. As the name suggests it is based on JSON and deals with data representing geographic features. There are several geometry types defined to be compatible with the types in the OpenGIS Simple Features Implementation Specification for SQL⁶. These are Point, MultiPoint, LineString, MultiLineString, Polygon, Multipolygon and the heterogeneous GeometryCollection. Listing 1 shows a simple example of a GeoJSON object with one point feature.

⁴<https://tools.ietf.org/html/rfc8259>

⁵<https://tools.ietf.org/html/rfc7946>

⁶https://portal.opengeospatial.org/files/?artifact_id=829

A more complete example can be viewed in the file `./data/example-7946.geojson`.

```
1 {  
2   "type": "Feature",  
3   "geometry": {  
4     "type": "Point",  
5     "coordinates": [125.6, 10.1]  
6   },  
7   "properties": {  
8     "name": "Dinagat Islands"  
9   }  
10 }
```

Listing 1: An example for a GeoJSON object

The feature types differ in the format of their coordinates property. A position is an array of at least two elements representing longitude and latitude. An optional third element can be added to specify altitude. All cases in this thesis will only deal with two-dimensional positions. While the coordinates member of a *Point*-feature is simply a single position, a *LineString*-feature describes its geometry through an Array of at least two positions. More interesting is the specification for Polygons. It introduces the concept of the *linear ring* as a closed *LineString* with at least four positions where the first and last positions are equivalent. The *Polygon*'s coordinates member is an array of linear rings with the first one representing the exterior ring and all others interior rings, also named surface and holes respectively. At last the coordinates member of *MultiLineStrings* and *MultiPolygons* is defined as a single array of its singular feature type.

GeoJSON is mainly used for web-based mapping. Since it is based on JSON it inherits its strength. There is for one the enhanced readability through reduced markup overhead compared to XML-based data types like GML. Interoperability with web applications comes for free since the parsing of JSON-objects is integrated in JavaScript. Unlike the Esri Shapefile format a single file is sufficient to store and transmit all relevant data, including feature properties.

To its downsides count that a text based cannot store the geometries as efficiently as it would be possible with a binary format. Also only vector-based data types can be represented. Another disadvantage can be the strictly non-topologic approach. Every feature is completely described by one entry. However when there are features that share common components, like boundaries in neighboring polygons, these data points will be encoded twice in the GeoJSON object. On the one hand this further poses concerns about data size. On the other hand it is more difficult to execute topological analysis on the data set. Luckily there is a related data structure to tackle this problem.

Extract more info about topology: <https://www.esri.com/news/arcuser/0401/topo.html>

TopoJSON is an extension of GeoJSON and aims to encode datastructures into a shared topology⁷. It supports the same geometry types as GeoJSON. It differs in some additional properties to use and new object types like "Topology" and "GeometryCollection". Its main feature is that LineStrings, Polygons and their multiplicatory equivalents must define line segments in a common property called "arcs". The geometries themselves then reference the arcs from which they are made up. This reduces redundancy of data points. Another feature is the quantization of positions. To use it one can define a "transform" object which specifies a scale and translate point to encode all coordinates. Together with delta-encoding of position arrays one obtains integer values better suited for efficient serialization and reduced file size.

find better word

Explain why topology-preserving shape simplification is important

Coordinate representation Both GeoJSON and TopoJSON represent positions as an array of numbers. The elements depict longitude, latitude and optionally altitude in that order. For simplicity this thesis will deal with two-dimensional positions only. A polyline is described by creating an array of these positions as seen in listing 2.

```
1 [[102.0, 0.0], [103.0, 1.0], [104.0, 0.0], [105.0, 1.0]]
```

Listing 2: Polyline coordinates in nested-array form

There will be however one library in this thesis which expects coordinates in a different format. Listing 3 shows a polyline in the sense of this library. Here one location is represented by an object with x and y properties.

```
1 [{x: 102.0, y: 0.0}, {x: 103.0, y: 1.0}, {x: 104.0, y: 0.0}, {x: 105.0, y: 1.0}]
```

Listing 3: Polyline in array-of-objects form

To distinguish these formats in future references the first format will be called nested-array format, while the latter will be called array-of-objects format.

⁷<https://github.com/topojson/topojson-specification>

2.3 Polyline simplification

In this chapter several algorithms for polyline simplification will be explained. For each algorithm a short summary of the routine and the runtime complexity will. At the end a comparison will be drawn to determine the method used for benchmarking.

n-th point algorithm This algorithm is fairly simplistic. It was described in 1966 by Tobler. The routine is to select every n-th coordinate of the polyline to retain. The larger the value of n is, the greater the simplification will be.

The Random-point routine is derived from the n-th point algorithm. It sections the line into parts containing n consecutive positions. From each section a random point is chosen to construct the simplified line.

$O(n)$

Radial distance algorithm Another simple algorithm to reduce points clustered too closely together. The algorithm will sequentially go through the line and eliminate all points whose distance to the current key is shorter than a given tolerance limit. As soon as a point with greater distance is found, it becomes the new key.

$O(n)$

Perpendicular distance algorithm Again a tolerance limit is given. The measure to check against is the perpendicular distance of a point to the line connecting its two neighbors. All points that exceed this limit are retained.

$O(n)$

Reumann-Witkam simplification As the name implies this algorithm was developed by Reumann and Witkam. In 1974 they described the routine that constructs a "corridor/search area" by placing two parallel lines in the direction of its initial tangent. The distance from this segment is user specified. Then the successive points will be checked until a point outside of this area is found. Its predecessor becomes a key and the two points mark the new tangent for the search area. This procedure is repeated until the last point is reached.

$O(n)$

The Opheim simplification Opheim extends the Reumann-Witkam algorithm in 1982 by constraining the search area. To do that two parameters **dmin** and **dmax** are given. From the key point on the last point inside a radial distance search region defined by **dmin** is taken to form the direction of the search corridor. If there is no point inside this region the subsequent point is taken. Then the process from the Reumann-Witkam algorithm is applied with the corridor constrained to a maximum distance of **dmax**.

$O(n)$

Lang simplification Lang described this algorithm in 1969. The search area is defined by a specified number of points too look ahead of the key point. A line is constructed from the key point to the last point in the search area. If the perpendicular distance of all intermediate points to this line is below a tolerance limit, they will be removed and the last point is the new key. Otherwise the search area is shrunk by excluding this last point until the requirement is met or there are no more intermediate points.

Douglas-Peucker simplification

with reduction parameter

Visvalingam-Whyatt simplification

Zhao-Saalfeld simplification

$O(??)$

maybe

$O(m^n)$

$O(n*m)$

$O(n*m)$

2.3.1 Summary

2.4 Running the algorithms on the web platform

JavaScript has been the only native programming language of web browsers for a long time. With the development of WebAssembly there seems to be an alternative on its way with high promises. This technology and the benefits and drawbacks to it will be explained in this chapter. It will be used to execute the algorithms under inspection in this thesis.

2.4.1 Introduction to Webassembly

WebAssembly is designed by engineers from the four major browser vendors (Mozilla, Google, Apple, Microsoft). It is a portable low-level bytecode designed as target for compilation of high-level languages. By being an abstraction over modern hardware it is language-, hardware-, and platform-independent. It is intended to be run in a stack-based virtual machine. This way it is not restrained to the Web platform or a JavaScript environment. Some key concepts are the structuring into modules with exported and imported definitions and the linear memory model. Memory is represented as a large array of bytes that can be dynamically grown. Security is ensured by the linear memory being disjoint from code space, the execution stack and the engine's data structures. Another

feature of WebAssembly is the possibility of streaming compilation and the parallelization of compilation processes.⁸

Benefits of WebAssembly The goals of WebAssembly have been well defined. It's semantics are intended to be safe and fast to execute and bring portability by language-, hardware- and platform-independence. Furthermore it should be deterministic and have simple interoperability with the web platform. For its representation the following goals are declared. It shall be compact and easy to decode, validate and compile. Parallelization and streamable compilation are also mentioned.

These goals are not specific to WebAssembly. They can be seen as properties that a low-level compilation target for the web should have. In fact there have been previous attempts to run low-level code on the web. Examples are Microsoft's ActiveX, Native Client (NaCl) and Emscripten each having issues complying with the goals. Java and Flash are examples for managed runtime plugins. Their usage is declining however not at least due to falling short on the goals mentioned above.

It is often stated that WebAssembly can bring performance benefits. It makes sense that statically typed machine code beats scripting languages performance wise. It has to be observed however if the overhead of switching contexts will neglect this performance gain. JavaScript has made a lot of performance improvements over the past years. Not at least Googles development on the V8 engine has brought JavaScript to an acceptable speed for extensive calculations. The engine observes the execution of running JavaScript code and will perform optimizations that can be compared to optimizations of compilers.

The JavaScript ecosystem has rapidly evolved the past years. Thanks to package managers like bower, npm and yarn it is simple to pull code from external sources into ones codebase. Initially thought for server sided JavaScript execution the ecosystem has found its way into front-end development via module bundlers like browserify, webpack and rollup. In course of this growth many algorithms and implementations have been ported to JavaScript for use on the web. With WebAssembly this ecosystem can be broadened even further. By lifting the language barrier existing work of many more programmers can be reused on the web. Whole libraries exclusive for native development could be imported by a few simple tweaks. Codecs not supported by browsers can be made available for use in any browser supporting WebAssembly. One example could be the promising AV1 video codec. In this these the C++ library psimpl will be utilized to bring polyline simplification to the web. This library already implements various

⁸<https://people.mpi-sws.org/~rossberg/papers/Haas,%20Rossberg,%20Schuff,%20Titzer,%20Gohman,%20Wagner,%20Zakai,%20Bastien,%20Holman%20-%20Bringing%20the%20Web%20up%20to%20Speed%20with%20WebAssembly.pdf>

algorithms for this task. It will be further introduced in chapter 7.1.

Existing compilers

emscripten

assemblyscript

rust

Technical hurdles

Managing memory

passing arrays

3 Implementation of a performance benchmark

In this chapter I will explain the approach to improve the performance of a simplification algorithm in a web browser via WebAssembly. The go-to library for this kind of operation is Simplify.js. It is the JavaScript implementation of the Douglas-Peucker algorithm with optional radial distance preprocessing. The library will be rebuilt in the C programming language and compiled to WebAssembly with Emscripten. A web page is built to produce benchmarking insights to compare the two approaches performance wise.

3.1 State of the art: Simplify.js

Simplify.js calls itself a "tiny high-performance JavaScript polyline simplification library"⁹. It was extracted from Leaflet, the "leading open-source JavaScript library for mobile-friendly interactive maps"¹⁰. Due to its usage in leaflet and Turf.js, a geospatial analysis library, it is the most common used library for polyline simplification. The library itself currently has 20,066 weekly downloads while the Turf.js derivate @turf/simplify has 30,389. Turf.js maintains an unmodified fork of the library in its own repository.

The Douglas-Peucker algorithm is implemented with an optional radial distance preprocessing routine. This preprocessing trades performance for quality. Thus the mode for disabling this routine is called highest quality.

Interestingly the library expects coordinates to be a list of object with x and y properties. GeoJSON and TopoJSON however store coordinates in nested array form (see chapter 2.2). Luckily since the library is small and written in JavaScript any skilled web developer can easily fork and modify the code for his own purpose. This is even pointed out in the source code. The fact that Turf.js, which can be seen as a convenience wrapper for processing GeoJSON data, decided to keep the library as is might indicate some benefit to this format. Listing 4 shows how Turf.js calls Simplify.js. Instead of altering the source code the data is transformed back and forth between the formats on each call. It is questionable if this practice is advisable at all.

Since it is not clear which case is faster, and given how simple the required changes are, two versions of Simplify.js will be tested. The original version, which expects the coordinates to be in array-of-objects format and the altered version, which operates on nested arrays. Listing 5 shows an extract of the changes performed on the library. Instead of using properties, the coordinate values are accessed by index. Except for the removal of the licensing header the alterations are restricted to these kind of changes. The full list

⁹<https://mourner.gifformthub.io/simplify-js/>

¹⁰<https://leafletjs.com/>

leaflet
down-
loads

```

1 function simplifyLine(coordinates, tolerance, highQuality) {
2   return simplifyJS(coordinates.map(function (coord) {
3     return {x: coord[0], y: coord[1], z: coord[2]};
4   }), tolerance, highQuality).map(function (coords) {
5     return (coords.z) ? [coords.x, coords.y, coords.z] : [coords.x,
6       coords.y];
7   });
}

```

Listing 4: Turf.js usage of simplify.js

of changes can be viewed in `lib/simplify-js-alternative/simplify.diff`.

```

1 13,14c4,5
2 <     var dx = p1.x - p2.x,
3 <     dy = p1.y - p2.y;
4 ---
5 >     var dx = p1[0] - p2[0],
6 >     dy = p1[1] - p2[1];

```

Listing 5: Snippet of the difference between the original Simplify.js and alternative

3.2 The webassembly solution

In scope of this thesis a library will be created that implements the same procedure as Simplify.JS in C code. It will be made available on the web platform through WebAssembly. In the style of the model library it will be called `Simplify.wasm`. The compiler to use will be Emscripten as it is the standard for porting C code to WebAssembly.

As mentioned the first step is to port simplify.JS to the C programming language. The file `lib/simplify-wasm/simplify.c` shows the attempt. It is kept as close to the JavaScript library as possible. This may result in C-untypical coding style but prevents skewed results from unexpected optimizations to the procedure itself. The entry point is not the `main`-function but a function called `simplify`. This is specified to the compiler as can be seen in listing 6.

Furthermore the functions `malloc` and `free` from the standard library are made available for the host environment. Compiling the code through Emscripten produces a binary file in `wasm` format and the glue code as JavaScript. These files are called `simplify.wasm` and `simplify.js` respectively.

An example usage can be seen in `lib/simplify-wasm/example.html`. Even through the memory access is abstracted in this example the process is still unhandy and far

More
about
the com-
piler call


```

1 OPTIMIZE="-O3"
2
3 simplify.wasm simplify.js: simplify.c
4   emcc \
5     ${OPTIMIZE} \
6     --closure 1 \
7     -s WASM=1 \
8     -s ALLOW_MEMORY_GROWTH=1 \
9     -s MODULARIZE=1 \
10    -s EXPORT_ES6=1 \
11    -s EXPORTED_FUNCTIONS='["_simplify", "_malloc", "_free"]' \
12    -o simplify.js \
13    simplify.c

```

Listing 6: The compiler call

from a drop-in replacement of Simplify.js. Thus in `lib/simplify-wasm/index.js` a further abstraction to the Emscripten emitted code was written. The exported function `simplifyWasm` handles module instantiation, memory access and the correct call to the exported wasm function. Finding the correct path to the wasm binary is not always clear however when the code is imported from another location. The proposed solution is to leave the resolving of the code-path to an asset bundler that processes the file in a preprocessing step.

```

1 export async function simplifyWasm(coords, tolerance, highestQuality) {
2   const module = await getModule()
3   const buffer = storeCoords(module, coords)
4   const resultInfo = module._simplify(
5     buffer,
6     coords.length * 2,
7     tolerance,
8     highestQuality
9   )
10  module._free(buffer)
11  return loadResultAndFreeMemory(module, resultInfo)
12 }

```

`../lib/simplify-wasm/index.js`

Listing 3.2 shows the function `simplifyWasm`. Further explanation will follow regarding the abstractions `getModule`, `storeCoords` and `loadResultAndFreeMemory`.

Module instantiation will be done on the first call only but requires the function to be asynchronous. For a neater experience in handling Emscripten modules a util-

ity function named `initEmscripten`¹¹ was written to turn the module factory into a JavaScript Promise that resolves on finished compilation. The usage of this function can be seen in listing 7. The resulting WebAssembly module is cached in the variable `emscriptenModule`.

```

1 let emscriptenModule
2 export async function getModule() {
3   if (!emscriptenModule)
4     emscriptenModule = initEmscriptenModule(wasmModuleFactory, wasmUrl)
5   return await emscriptenModule
6 }

```

Listing 7: My Caption

Storing coordinates into the module memory is done in the function `storeCoords`. Emscripten offers multiple views on the module memory. These correspond to the available WebAssembly data types (e.g. `HEAP8`, `HEAPU8`, `HEAPF32`, `HEAPF64`, ...) ¹². As Javascript numbers are always represented as a double-precision 64-bit binary ¹³ (IEEE 754-2008) the `HEAP64`-view is the way to go to not lose precision. Accordingly the datatype `double` is used in C to work with the data. Listing 8 shows the transfer of coordinates into the module memory. In line 3 the memory is allocated using the exported `malloc`-function. A JavaScript `TypedArray` is used for accessing the buffer such that the loop for storing the values (lines 5 - 8) is trivial.

```

1 export function storeCoords(module, coords) {
2   const flatSize = coords.length * 2
3   const offset = module._malloc(flatSize * Float64Array.
4     BYTES_PER_ELEMENT)
5   const heapView = new Float64Array(module.HEAPF64.buffer, offset,
6     flatSize)
7   for (let i = 0; i < coords.length; i++) {
8     heapView[2 * i] = coords[i][0]
9     heapView[2 * i + 1] = coords[i][1]
10  }
11  return offset
12 }

```

Listing 8: The storeCoords function

¹¹[/lib/wasm-util/initEmscripten.js](#)

¹²https://emscripten.org/docs/api_reference/preamble.js.html#type-accessors-for-the-memory-model

¹³<https://www.ecma-international.org/ecma-262/6.0/#sec-4.3.20>

programmi
Check
for
coords
length <
2

To read the result back from memory we have to look at how the simplification will be returned in the C code. Listing 9 shows the entry point for the C code. This is the function that gets called from JavaScript. As expected arrays are represented as pointers with corresponding length. The first block of code (line 2 - 6) is only meant for declaring needed variables. Lines 8 to 12 mark the radial distance preprocessing. The result of this simplification is stored in an auxiliary array named `resultRdDistance`. In this case `points` will have to point to the new array and the length is adjusted. Finally the Douglas-Peucker procedure is invoked after reserving enough memory. The auxiliary array can be freed afterwards. The problem now is to return the result pointer and the array length back to the calling code. The fact that pointers in Emscripten are represented by an integer will be exploited to return a fixed size array of two containing the values. A hacky solution but it works. We can now look back at how the JavaScript code reads the result.

Fact
check.
evtl un-
signed

```

1  int* simplify(double * points, int length, double tolerance, int
    highestQuality) {
2      double sqTolerance = tolerance * tolerance;
3      double* resultRdDistance = NULL;
4      double* result = NULL;
5      int resultLength;
6
7      if (!highestQuality) {
8          resultRdDistance = malloc(length * sizeof(double));
9          length = simplifyRadialDist(points, length, sqTolerance,
resultRdDistance);
10         points = resultRdDistance;
11     }
12
13     result = malloc(length * sizeof(double));
14     resultLength = simplifyDouglasPeucker(points, length, sqTolerance,
result);
15     free(resultRdDistance);
16
17     int* resultInfo = malloc(2);
18     resultInfo[0] = (int) result;
19     resultInfo[1] = resultLength;
20     return resultInfo;
21 }

```

Listing 9: Entrypoint in the C-file

Listing 10 shows the code to read the values back from module memory. The result pointer and its length are acquired by dereferencing the `resultInfo`-array. The buffer to use is the heap for unsigned 32-bit integers. This information can then be used to align the `Float64Array`-view on the 64-bit heap. Constructing the appropriate coordinate representation by reversing the flattening can be looked up in the same file. It is realised

in the `unflattenCoords` function. At last it is important to actually free the memory reserved for both the result and the result-information. The exported method `free` is the way to go here.

```
1 export function loadResultAndFreeMemory(module, resultInfo) {  
2   const [resultPointer, resultLength] = new Uint32Array(  
3     module.HEAPU32.buffer,  
4     resultInfo,  
5     2  
6   )  
7   const simplified = new Float64Array(  
8     module.HEAPF64.buffer,  
9     resultPointer,  
10    resultLength  
11  )  
12  const coords = unflattenCoords(simplified)  
13  module._free(resultInfo)  
14  module._free(resultPointer)  
15  return coords
```

Listing 10: Loading coordinates back from module memory

3.3 The implementation of a web framework

The performance comparison of the two methods will be realized in a web page. It will be built as a front-end web-application that allows the user to specify the input parameters of the benchmark. These parameters are: The polyline to simplify, a range of tolerances to use for simplification and if the so called high quality mode shall be used. By building this application it will be possible to test a variety of use cases on multiple devices. Also the behavior of the algorithms can be researched under different preconditions. In the scope of this thesis a few cases will be investigated. The application structure will now be introduced.

3.3.1 External libraries

The dynamic aspects of the web page will be built in JavaScript to make it run in the browser. Webpack¹⁴ will be used to bundle the application code and use compilers like babel¹⁵ on the source code. As mentioned in section 3.2 the bundler is also useful for handling references to the WebAssembly binary as it resolves the filename to the correct download path to use. There will be intentionally no transpiling of the JavaScript code

¹⁴<https://webpack.js.org/>

¹⁵<https://babeljs.io/>

to older versions of the ECMA standard. This is often done to increase compatibility with older browsers. Luckily this is not a requirement in this case and by refraining from this practice there will also be no unintentional impact on the application performance. Libraries in use are Benchmark.js¹⁶ for statistically significant benchmarking results, React¹⁷ for the building the user interface and Chart.js¹⁸ for drawing graphs.

3.3.2 The application logic

The web page consist of static and dynamic content. The static parts refer to the header and footer with explanation about the project. Those are written directly into the root HTML document. The dynamic parts are injected by JavaScript. Those will be further discussed in this chapter as they are the main application logic.

The web app is built to test a variety of cases with multiple datapoints. As mentioned Benchmark.js will be used for statistically significant results. It is however rather slow as it needs about 5 to 6 seconds per datapoint. This is why multiple types of benchmarking methods are implemented. Figure 3.3.2 shows the corresponding UML diagram of the application. One can see the UI components in the top-left corner. The root component is **App**. It gathers all the internal state of its children and passes state down where it is needed.

3.3.3 Benchmark cases and chart types

In the upper right corner the different Use-Cases are listed. These cases implement a function "**fn**" to benchmark. Additional methods for setting up the function and clean up afterwards can be implemented as given by the parent class **BenchmarkCase**. Concrete cases can be created by instantiating one of the BenchmarkCases with a defined set of parameters. There are three charts that will be rendered using a subset of these cases. These are:

- **Simplify.js vs Simplify.wasm** - This Chart shows the performance of the simplification by Simplify.js, the altered version of Simplify.js and the newly developed Simplify.wasm.
- **Simplify.wasm runtime analysis** - To further gain insights to WebAssembly performance this stacked barchart shows the runtime of a call to Simplify.wasm. It is partitioned into time spent for preparing data (**storeCords**), the algorithm itself and the time it took for the coordinates being restored from memory (**loadResult**).

Cases

¹⁶<https://benchmarkjs.com/>

¹⁷<https://reactjs.org/>

¹⁸<https://www.chartjs.org/>

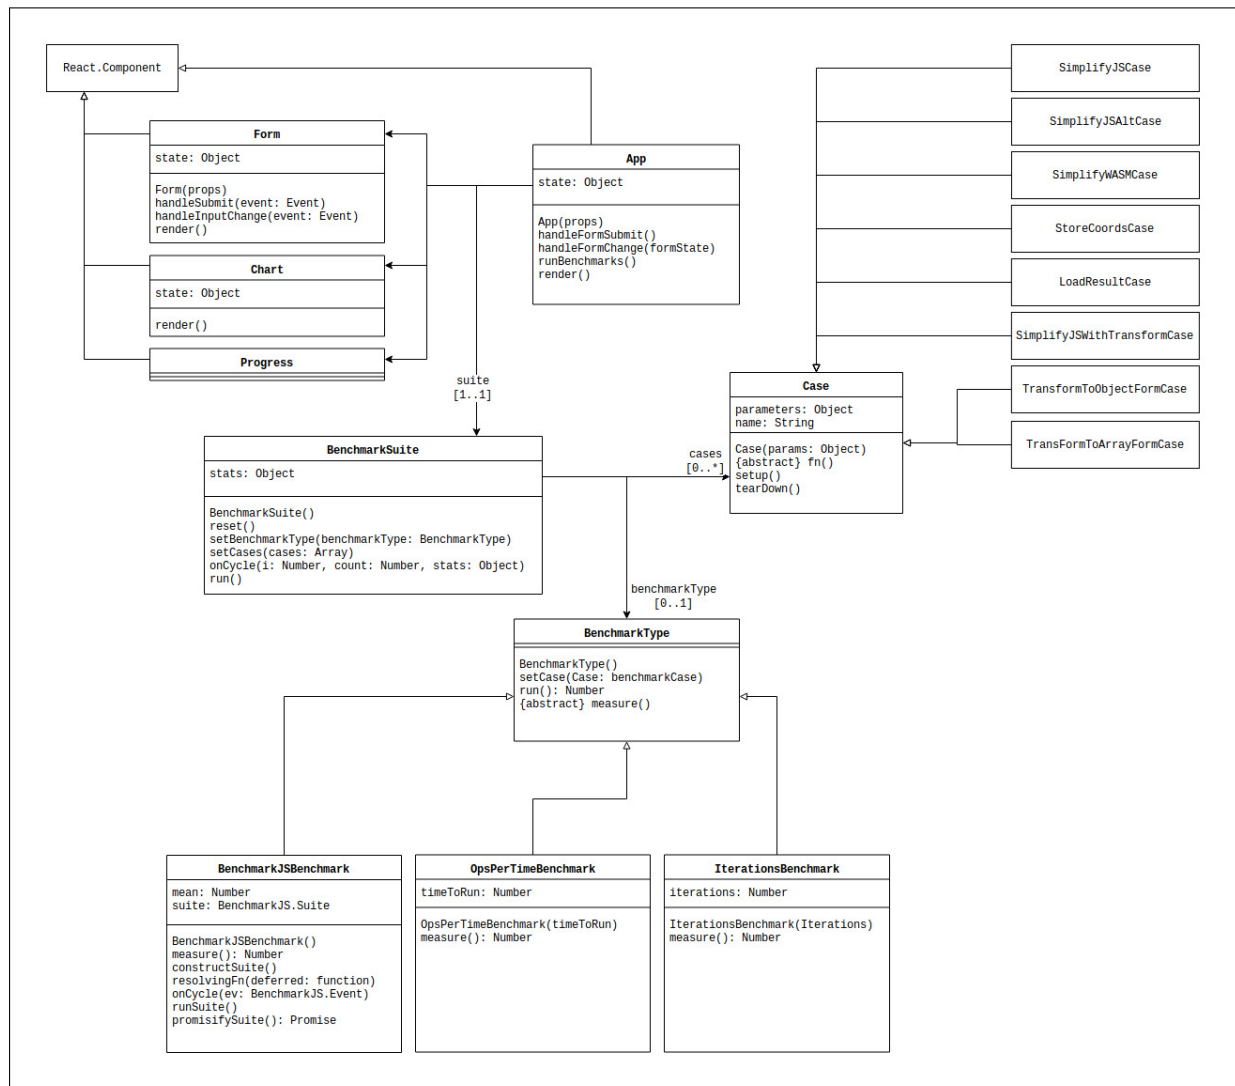


Figure 1: UML diagram of the benchmarking application

- **Turf.js method runtime analysis** - The last chart will use a similar structure. This time it analyses the performance impact of the back and forth transformation of data used in Truf.js.

Cases

3.3.4 The different benchmark types

On the bottom the different types of Benchmarks implemented can be seen. They all implement the abstract `measure` function to return the mean time to run a function specified in the given `BenchmarkCase`. The `IterationsBenchmark` runs the function a specified number of times, while the `OpsPerTimeBenchmark` always runs a certain amount of milliseconds to run as much iterations as possible. Both methods got their benefits and

drawbacks. Using the iterations approach one cannot determine the time the benchmark runs beforehand. With fast devices and a small number of iterations one can even fall in the trap of the duration falling under the accuracy of the timer used. Those results would be unusable of course. It is however a very fast way of determining the speed of a function. And it holds valuable for getting a first approximation of how the algorithms perform over the span of datapoints. The second type, the operations per time benchmark, seems to overcome this problem. It is however prone to garbage collection, engine optimizations and other background processes.¹⁹

Benchmark.js combines these approaches. In a first step it approximates the runtime in a few cycles. From this value it calculates the number of iterations to reach an uncertainty of at most 1%. Then the samples are gathered.²⁰

more
about
Bench-
mark.js

3.3.5 The benchmark suite

For running multiple benchmarks the class `BenchmarkSuite` was created. It takes a list of `BenchmarkCases` and runs them through a `BenchmarkType`. The Suite manages starting, pausing and stopping of going through list. It updates the statistics gathered on each cycle. By injecting an `onCycle` method, the `App` component can give live feedback about the progress.

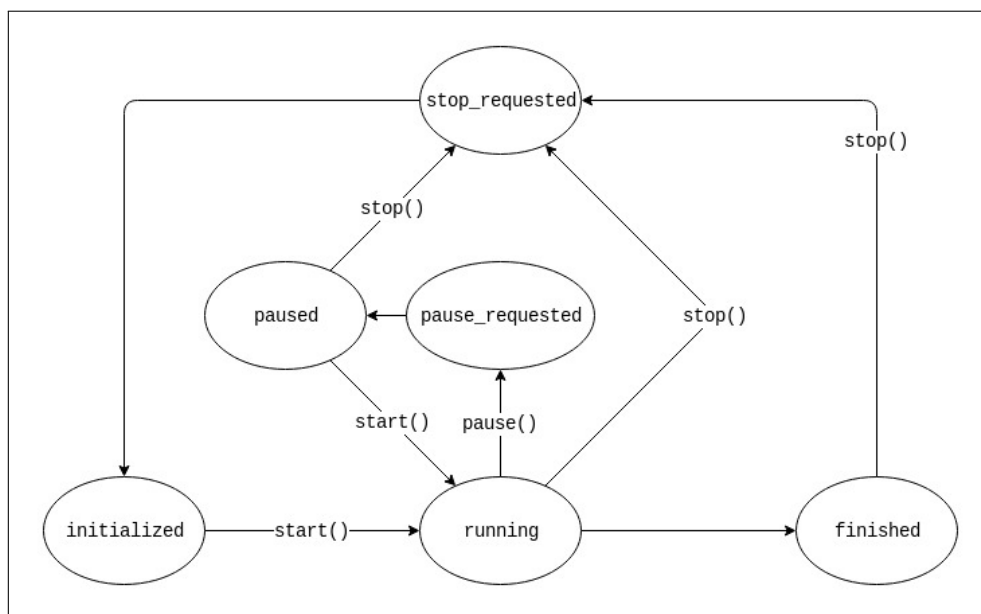


Figure 2: The state machine for the benchmark suite

Figure 3.3.5 shows the state machine of the suite. Based on this diagram the ui

¹⁹<https://calendar.perfplanet.com/2010/bulletproof-javascript-benchmarks/>

²⁰<http://monsur.hossa.in/2012/12/11/benchmarkjs.html>

component shows action buttons so the user can interact with the state. While running the suite checks if a state change was requested and acts accordingly by pausing the benchmarks or resetting all statistics gathered when stopping.

3.3.6 The user interface

The user interface has three regions. One for configuring input parameters. One for controlling the benchmark process and at last a diagram of the results. Figure 3 shows the user interface.

Settings At first the input parameters of the algorithm have to be specified. For that there are some polylines prepared to choose from. They are introduced in chapter 3.4. Instead of testing a single tolerance value the user can specify a range. This way the behavior of the algorithms can be observed in one chart. The high quality mode got its name from Simplify.js. If it is enabled there will be no radial-distance preprocessing step before applying the Douglas-Peucker routine. The next option determines which benchmarks will be run. The options are mentioned in chapter 3.3.3. One of the three benchmark methods implemented can be selected. Depending on the method chosen additional options will show to further specify the benchmark parameters. The last option deals with chart rendering. Debouncing limits the rate at which functions fire. In this case the chart will delay rendering when datapoints come in at a fast rate.

Run Benchmark This is the control that displays the status of the benchmark suite. Here benchmarks can be started, stopped, paused and resumed. It also shows the progress of the benchmarks completed in percentage and absolute numbers.

Chart The chart shows a live diagram of the results. The title represents the selected chart. The legend gives information on which benchmark cases will run. Also the algorithm parameters (dataset and high quality mode) and current platform description can be found here. The tolerance range maps over the x-Axis. On the y-Axis two scales can be seen. The left hand shows by which unit the performance is displayed. This scale corresponds to the colored lines. Every chart will show the number of positions in the result as a grey line. Its scale is displayed on the right. This information is important for selecting a proper tolerance range as it shows if a appropriate order of magnitude has been chosen. Below the chart additional control elements are placed to adjust the visualization. The first selection lets the user choose between a linear or logarithmic y-Axis. The second one changes the unit of measure for performance. The two options are the




Figure 3: The user interface for benchmarking application.

mean time in milliseconds per operation (ms) and the number of operations that can be run in one second (hz). These options are only available for the chart "Simplify.wasm vs. Simplify.js" as the other two charts are stacked bar charts where changing the default options won't make sense. Finally the result can be saved via a download button. A separate page can be fed with this file to display the diagram only.

3.4 The test data

Here the test data will be shown. There are two data sets chosen to operate on. The first is a testing sample used in Simplify.js the second one a boundary generated from the OpenStreetMap (OSM) data.

Simplify.js example This is the polyline used by Simplify.js to demonstrate its capabilities. Figure 4 shows the widget on its homepage. The user can modify the parameters with the interactive elements and view the live result. The data comes from a 10.700 mile car route from Lisboa, Portugal to Singapore and is based on OpenStreetMap data. The line is defined by 73.752 positions. Even with low tolerances this number reduces drastically. This example shows perfectly why it is important to generalize polylines before rendering them.

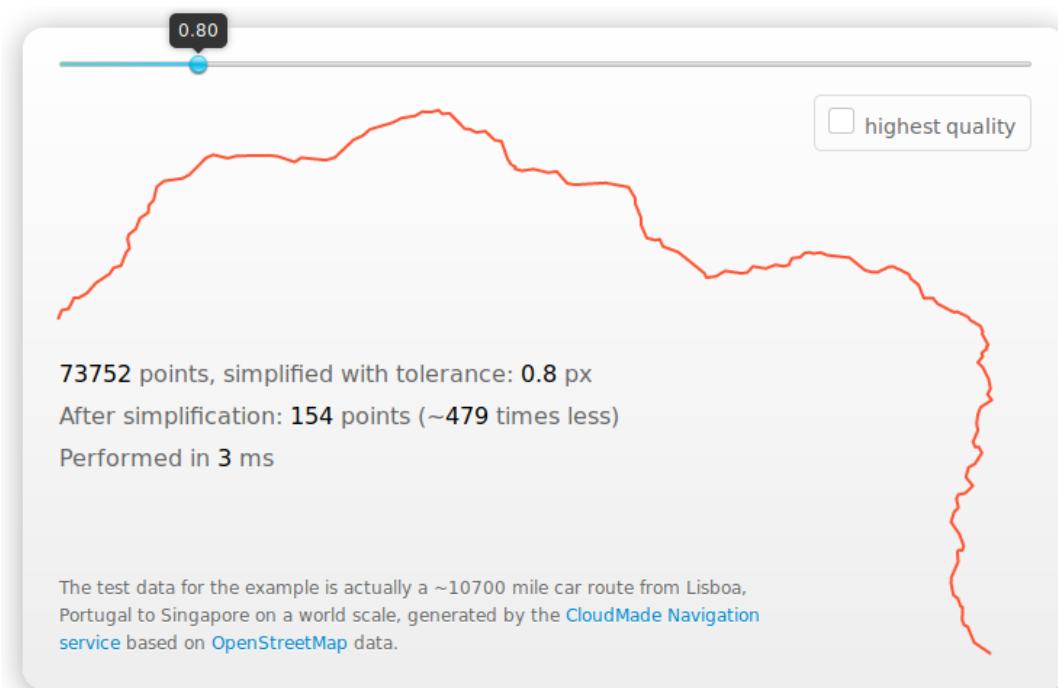


Figure 4: The Simplify.js test data visualized

Bavaria outline The second polyline used for benchmarking contains 116.829 positions. It represents the outline of a german federate state, namely bavaria. It was extracted from the OSM dataset by selecting administrative boundaries. On the contrary to the former polyline this one is a closed line, often used in polygons to represent a surface. The plotted line can be seen in figure 5.




Figure 5: The Bavaria test data visualized

Simple line There is a third line used in the application to choose from. This one is however not used for benchmarking since it contains only 8 points. It is merely a placeholder to prevent the client application to load a bigger data sets from the server on page load. This way the transmitted data size will be reduced. The larger lines will only be requested when they are actually needed.

4 Benchmark results

In this chapter the results are presented. There were a multitude of tests to make. Multiple devices were used to run several benchmarks on different browsers and under various parameters. To organize which benchmarks had to run, first all the problem dimensions were clarified. Devices will be categorized into desktop and mobile devices. The browsers to test will come from the four major browser vendors which were involved in WebAssembly development. Those are Firefox from Mozilla, Chrome from Google, Edge from Microsoft and Safari from Apple. For either of the two data sets a fixed range of tolerances is set to maintain consistency across the diagrams. The values are explained in chapter 3.4. The other parameter "high quality" can be either switched on or off. The three chart types are explained in chapter 3.3.3.

describe website for results

Each section in this chapter describes a set of benchmarks run on the same system. A table in the beginning will indicate the problem dimensions chosen to inspect. After a description of the system and a short summary of the case the results will be presented in the form of graphs. Those are the graphs produced from the application described in chapter 3.3

4.1 Case 1 - Windows - wasm vs js

Device:	Desktop		Mobile	
Browser:	Firefox	Chrome	Edge	Safari
Dataset:	Simplify.js example		Bavaria outline	
High Quality:	On		Off	
Charts:	Simplify.js vs. Simplify.wasm	Simplify.wasm runtime analysis		Turf.js runtime analysis

Table 1: Problem dimensions of Case 1

At first it will be observed how the algorithms perform under different browsers. The chart to use for this is the "Simplify.js vs Simplify.wasm" chart. For that a Windows system was chosen as it allows to run benchmarks under three of the four browsers in question. The dataset is the Simplify.js example which will be simplified with and without the high quality mode.

The device is a HP Pavilion x360 - 14-ba101ng²¹ convertible. It contains an Intel® Core™ i5-8250U Processor with 4 cores, 6MB cache. The operating system is Windows

²¹<https://support.hp.com/us-en/product/hp-pavilion-14-ba100-x360-convertible-pc/16851098/model/18280360/document/c05691748>

hp pavilion
6 charts
questions:
1, 3, 5

10 and the browsers are on their newest versions with Chrome 76, Firefox 68 and Edge 44.18362.1.0.




Figure 6: Simplify.wasm vs. Simplify.js benchmark result of Windows device with Chrome browser on dataset "Simplify.js example" without high quality mode.




Figure 7: Simplify.wasm vs. Simplify.js benchmark result of Windows device with Firefox browser on dataset "Simplify.js example" without high quality mode.




Figure 8: Simplify.wasm vs. Simplify.js benchmark result of Windows device with Edge browser on dataset "Simplify.js example" without high quality mode.




Figure 9: Simplify.wasm vs. Simplify.js benchmark result of Windows device with Chrome browser on dataset "Simplify.js example" with high quality mode.




Figure 10: Simplify.wasm vs. Simplify.js benchmark result of Windows device with Firefox browser on dataset "Simplify.js example" with high quality mode.




Figure 11: Simplify.wasm vs. Simplify.js benchmark result of Windows device with Edge browser on dataset "Simplify.js example" with high quality mode.

4.2 Case 2 - Windows - wasm runtime analysis

For this case the same device as in the former case is used. To compare the results of the two cases the same dataset is used. Under the Edge browser the Simplify.wasm

hp pavilion
2 charts
questions:
2, 3, 5

maybe
add ffox

Device:	Desktop		Mobile	
Browser:	Firefox	Chrome	Edge	Safari
Dataset:	Simplify.js example		Bavaria outline	
High Quality:	On		Off	
Charts:	Simplify.js vs. Simplify.wasm		Simplify.wasm runtime analysis	
			Turf.js runtime analysis	

Table 2: Problem dimensions of Case 2

runtime analysis was measured.




Figure 12: Simplify.wasm runtime analysis benchmark result of Windows device with Edge browser on dataset "Simplify.js example" without high quality mode.




Figure 13: Simplify.wasm runtime analysis benchmark result of Windows device with Edge browser on dataset "Simplify.js example" with high quality mode.

4.3 Case 3 - MacBook Pro - wasm vs js

A 2018 MacBook Pro 15" will be used to test the safari browser. For comparison the benchmarks will also be held under Firefox on MacOS. This time the bavarian boundary will be simplified with both preprocessing enabled and disabled.

Device:	Desktop		Mobile	
Browser:	Firefox	Chrome	Edge	Safari
Dataset:	Simplify.js example		Bavaria outline	
High Quality:	On		Off	
Charts:	Simplify.js vs. Simplify.wasm		Simplify.wasm runtime analysis	Turf.js runtime analysis

Table 3: Problem dimensions of Case 3

MacBook Pro 15 4 charts Chrome and FF comparable to results above




Figure 14: Simplify.wasm vs. Simplify.js benchmark result of MacBook Pro device with Safari browser on dataset "Bavaria outline" without high quality mode.




Figure 15: Simplify.wasm vs. Simplify.js benchmark result of MacBook Pro device with Firefox browser on dataset "Bavaria outline" without high quality mode.




Figure 16: Simplify.wasm vs. Simplify.js benchmark result of MacBook Pro device with Safari browser on dataset "Bavaria outline" with high quality mode.




Figure 17: Simplify.wasm vs. Simplify.js benchmark result of MacBook Pro device with Firefox browser on dataset "Bavaria outline" with high quality mode.

Device:	Desktop		Mobile	
Browser:	Firefox	Chrome	Edge	Safari
Dataset:	Simplify.js example		Bavaria outline	
High Quality:	On		Off	
Charts:	Simplify.js vs. Simplify.wasm		Simplify.wasm runtime analysis	Turf.js runtime analysis

Table 4: Problem dimensions of Case 4

4.4 Case 4 - Ubuntu - turf.js analysis




Figure 18: Simplify.wasm vs. Simplify.js benchmark result of Ubuntu device with Firefox browser on dataset "Bavaria outline" without high quality mode.

Lenovo
Miix 510
4 charts
Firefox be-
cause orig
simplify is
faster




Figure 19: Turf.js simplify benchmark result of Ubuntu device with Firefox browser on dataset "Bavaria outline" without high quality mode.




Figure 20: Simplify.wasm vs. Simplify.js benchmark result of Ubuntu device with Firefox browser on dataset "Bavaria outline" with high quality mode.




Figure 21: Turf.js simplify benchmark result of Ubuntu device with Firefox browser on dataset "Bavaria outline" with high quality mode.

4.5 Case 5 - iPhone - mobile testing

iPhone??
6 charts

Device:	Desktop		Mobile	
Browser:	Firefox	Chrome	Edge	Safari
Dataset:	Simplify.js example		Bavaria outline	
High Quality:	On		Off	
Charts:	Simplify.js vs. Simplify.wasm		Simplify.wasm runtime analysis	Turf.js runtime analysis

Table 5: Problem dimensions of Case 5

5 Discussion

In this section the results are interpreted. This section is structured in different questions to answer. First it will be analyzed what the browser differences are. One section will deal with the performance of the pure JavaScript implementations while the next will inspect how Simplify.wasm performs. Then further insights to the performance of the WebAssembly implementation will be given. It will be investigated how long it takes to set up the WebAssembly call and how much time is spent to actually execute the simplification routines. Next the case of Turf.js will be addressed and if its format conversions are reasonable under specific circumstances. Finally the performance of mobile devices will be evaluated.

5.1 Browser differences for the JavaScript implementations

The first thing to see from the results of chapter 4.1 and 4.3 is that there is actually a considerable performance difference in the two versions of Simplify.js. So here we take a closer look at the JavaScript performance of the browsers. Interestingly clear winner between the similar algorithms cannot be determined as the performance is inconsistent across browsers. While the original version is faster in Firefox and Safari, the altered version is superior in Chrome and Edge. This is regardless of whether the high quality mode is switched on or not. The difference is however more significant when the preprocessing step is disabled.

In figure 9 and 11 one can see how similar Chrome and Edge perform with high quality mode enabled. When disabled however the algorithms perform similar in Edge (figure 11) while in Chrome the alternative version still improves upon the original.

In Firefox the result is very different. Without the high quality mode the original version performs about 2.5 times better than the alternative. Figure 7 shows this. When disabling the preprocessing the performance gain is even higher. the original performs constantly 3x faster as seen in figure 10.

The same results can be reproduced under Firefox on macOS with the "Bavarian outline" dataset (figures 15 and 17). Interestingly under safari the algorithms perform similarly with a small preference to the original version. This applies to either case tested (figures 14 and 16).

With so much variance it is hard to determine the best performing browser regarding the JavaScript implementation. Under the right circumstances Chrome can produce the fastest results with the alternative implementation. Safari is consistently very fast. Even while it falls short to Firefox's results with the original algorithm when high quality is

turned on. The greatest discrepancy was produced by Firefox with high quality requested. There the alternate version produced the slowest results while the results with Simplify.js can compete with Chrome's results with the Simplify.js alternative. Edge lies between these two browsers with not too bad but also not the fastest results.

5.2 Browser differences for Simplify.wasm

So diverse the results from last chapter were, so monotonous they will be here. The performance of the Simplify.wasm function is consistent across all browsers tested. This is a major benefit brought by WebAssembly often described as predictable performance.

The variance is very low when the preprocessing is turned off through the high quality mode. The browsers produce about the same runtimes under the same conditions. When high quality is off the Chrome browser got its nose ahead with a mean runtime of 0.66ms. Edge follows with 1.02ms and Firefox takes an average 1.10ms. The results of chapter 4.3 show that Safari is a bit faster at the high quality mode than Firefox but slower without.

5.3 Insights into Simplify.wasm

So far when the performance of Simplify.wasm was addressed it meant the time spent for the whole process of preparing memory to running the algorithm in wasm context to loading back the result to JavaScript. This makes sense when comparing it to the JavaScript library with the motive to replace it one for one. It does however not produce meaningful comparisons of WebAssembly performance in contrast to the native JavaScript runtime.

This is because until Simplify.wasm actually gets to the WebAssembly computation it will take some time to prepare the data and load it back afterwards. To see how long that is the additional chart type "Simplify.wasm runtime insights" was created. In figures 12 and 13 the results can be seen. Each data point contains a stacked column to represent the proportion of time spent for each task. The blue section represents the time spent to initialize the memory, the red one the execution of the compiled WebAssembly code. At last the green part will show the time spent for getting the coordinates back in the right format.

First the parts where JavaScript is run will be examined. There is as good as no variance in the memory initialization. This is obviously due to the fact that this step is not dependent on any other parameter than the polyline length. Initial versions of the library produced in this thesis were not as efficient in flattening the coordinate array as the final version. By replacing the built-in `Array.prototype.flat`-method with a

simple for loop a good amount optimization was achieved on the JavaScript side of the *Simplify.wasm* process. The flat method is a rather new feature of ECMAScript and its performance might be enhanced in future browser versions. This example shows however that when writing JavaScript code one can quickly deviate from the "fast path" even when dealing with simple problems.

On the other side of process lies the function `loadResult`. It is dependent on the size of the resulting polyline. Since this is often very low in the examples used the green bar can be rarely seen. Merely at low tolerance values like in figure 12 the influence is visible. The maximum fraction there is at tolerance value 0.05 where the operation takes 4.26% of the total execution time.

Now when comparing the two graphs one can clearly see that the influence of the JavaScript portions is much greater when the high quality mode is turned of. The time taken for preparing the memory in both cases is about 0.67ms. The execution time of the algorithms is so low in the first case, that it comes down to making up only 24,47% when taking the median values. In case where high quality is enabled the results do not look as drastic. The median value of execution time is 4.31ms and with that much greater than preparation time. If JavaScript is at advantage in the first case and the high execution time justifies the switch of runtimes in the latter will be examined in the next chapter.

5.4 Comparison *Simplify.wasm* vs *Simplify.js*

when is what faster

5.5 Analysis of *Turf.js* implementation

When is *turf.js* faster

5.6 Mobile device analysis

6 Conclusion

6.1 Enhancements

Enhancement: Line Smoothing as preprocessing step

6.2 Future Work

7 Compiling an existing C++ library for use on the web

maybe remove whole chapter :’(

In this chapter I will explain how an existing C++ library was utilized compare different simplification algorithms in a web browser. The library is named *psimpl* and was written in 2011 from Elmar de Koning. It implements various Algorithms used for polyline simplification. This library will be compiled to WebAssembly using the Emscripten compiler. Furthermore a Web-Application will be created for interactively exploring the Algorithms. The main case of application is simplifying polygons, but also polylines will be supported. The data format used to read in the data will be GeoJSON. To maintain topological correctness a intermediate conversion to TopoJSON will be applied if requested.

Integrating an existing C++ library An existing implementation of several simplification algorithms has been found in the C++ ecosystem. *psimpl* implements 8 algorithms distributed as a single header file. It also provides a function for measuring positional errors making it ideal for use in a quality analysis tool for those algorithms.

7.1 State of the art: *psimpl*

psimpl is a generic C++ library for various polyline simplification algorithms. It consists of a single header file *psimpl.h*. The algorithms implemented are *Nth point*, *distance between points*, *perpendicular distance*, *Reumann-Witkam*, *Opheim*, *Lang*, *Douglas-Peucker* and *Douglas-Peucker variation*. It has to be noted, that the *Douglas-Peucker* implementation uses the *distance between points* routine, also named the *radial distance* routine, as preprocessing step just like *Simplify.js* (Section 3.1). All these algorithms have a similar templated interface. The goal now is to prepare the library for a compiler.

Describe the error statistics function of *psimpl*

7.2 Compiling to WebAssembly

As in the previous chapter the compiler created by the Emscripten project will be used. This time the code is not directly meant to be consumed by a web application. It is a generic library. There are no entry points defined that Emscripten can export in WebAssembly. So the entry points will be defined in a new package named *psimpl-js*. It

will contain a C++ file that uses the library, the compiled code and the JavaScript files needed for consumption in a JavaScript project. *psimpl* makes heavy use of C++ template functions which cannot be handled by JavaScript. So there will be entry points written for each exported algorithm. These entry points are the point of intersection between JavaScript and the library. Listing 11 shows one example. They all follow the same procedure. First the pointer given by JavaScript is interpreted as a double-pointer in line 2. This is the beginning of the coordinates array. *psimpl* expects the first and last point of an iterator so the pointer to the last point is calculated (line 3). The appropriate function template from *psimpl* is instantiated and called with the other given parameters (line 5). The result is stored in an intermediate vector.

```

1  val douglas_peucker(uintptr_t ptr, int length, double tol) {
2      double* begin = reinterpret_cast<double*>(ptr);
3      double* end = begin + length;
4      std::vector<double> resultCoords;
5      psimpl::simplify_douglas_peucker<2>(begin, end, tol, std::
      back_inserter(resultCoords));
6      return val(typed_memory_view(resultCoords.size(), &resultCoords[0]))
      ;
7  }

```

Listing 11: One entrypoint to the C++ code

Since this is C++ the capabilities of Emscripten's Embind can be utilized. Embind is realized in the libraries `bind.h`²² and `val.h`²³. `val.h` is used for transliterating JavaScript to C++. In this case it is used for the type conversion of C++ Vectors to JavaScript's Typed Arrays as seen at the end of listing 11. On the other hand `bind.h` is used for binding C++ functions, classes, or enumerations to from JavaScript callable names. Aside from providing a better developer experience this also prevents name mangling in cases where functions are overloaded. Instead of listing the exported functions in the compiler command or annotating it with `EMSCRIPTEN_KEEPALIVE` the developer gives a pointer to the object to bind. Listing 12 shows each entry point bound to a readable name and at last the registered vector datatype. The parameter `my_module` is merely for marking a group of related bindings to avoid name conflicts in bigger projects.

Compiler call (`-bind`)

The library code on JavaScript side is similar to the one in chapter 3.2. This time a function is exported per routine.

More about javascript glue code with listing `callSimplification`.

²²https://emscripten.org/docs/api_reference/bind.h.html#bind-h

²³https://emscripten.org/docs/api_reference/val.h.html#val-h

```

1 EMSCRIPTEN_BINDINGS(my_module) {
2     function("nth_point", &nth_point);
3     function("radial_distance", &radial_distance);
4     function("perpendicular_distance", &perpendicular_distance);
5     function("reumann_witkam", &reumann_witkam);
6     function("opheim", &opheim);
7     function("lang", &lang);
8     function("douglas_peucker", &douglas_peucker);
9     function("douglas_peucker_n", &douglas_peucker_n);
10    register_vector<double>("vector<double>");
11 }

```

Listing 12: Emscripten bindings

7.3 The implementation

The implementation is just as in the last chapter a web page and thus JavaScript is used for the interaction. The source code is bundled with Webpack. React is the UI Component library and babel is used to transform JSX to JavaScript. MobX²⁴ is introduced as a state management library. It applies functional reactive programming by giving the utility to declare observable variables and triggering the update of derived state and other observers intelligently. To do that MobX observes the usage of observable variables so that only dependent observers react on updates. In contrast to other state libraries MobX does not require the state to be serializable. Many existing data structures can be observed like objects, arrays and class instances. It also does not constrain the state to a single centralized store like Redux²⁵ does. The final state diagram can be seen in listing 22. It represents the application state in an object model. Since this has drawbacks in showing the information flow the observable variables are marked in red, and computed ones in blue.

On the bottom the three main state objects can be seen. They are implemented as singletons as they represent global application state. Each of them will now be explained.

MapState holds state relevant for the map display. An array of `TileLayers` defines all possible background layers to choose from. The selected one is stored in `selectedTileLayerId`. The other two variables toggle the display of the vector layers to show.

AlgorithmState stores all the information about the simplification algorithms to choose from. The class `Algorithm` acts as a generalization interface. Each algorithm defines which fields are used to interact with its parameters. These fields hold their current

²⁴<https://mobx.js.org/>

²⁵<https://redux.js.org/>

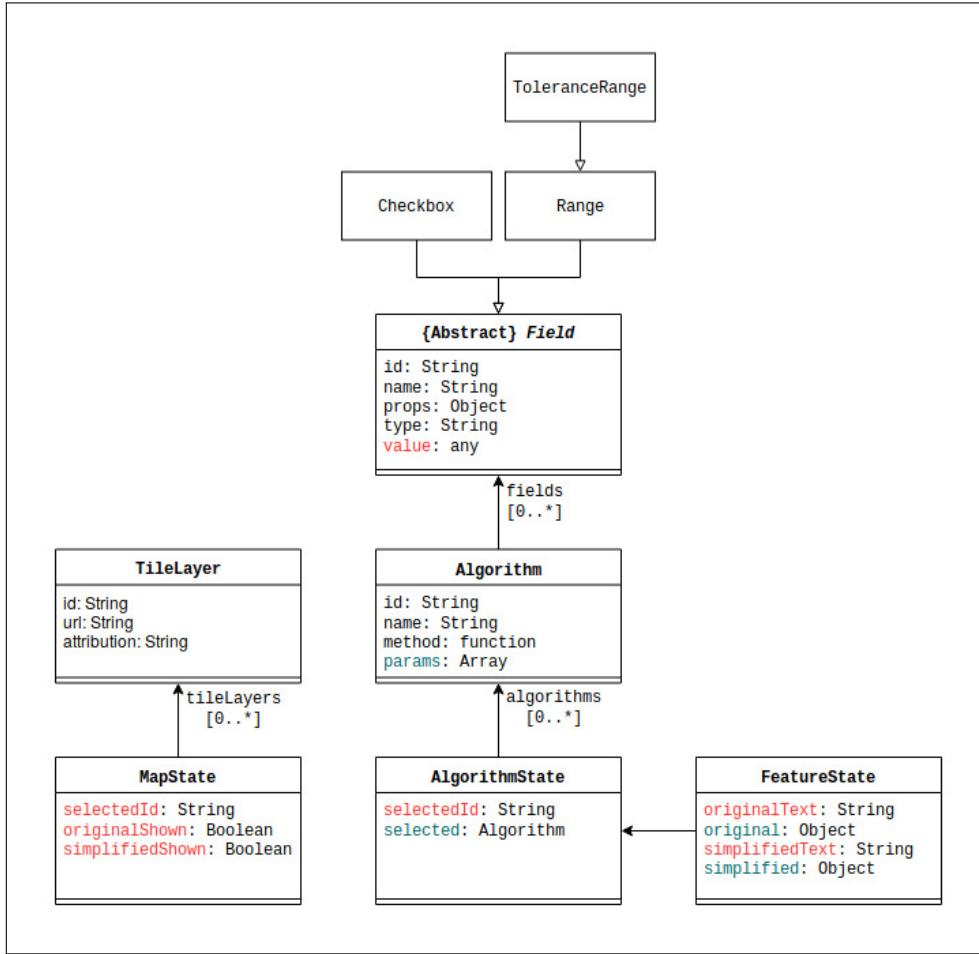


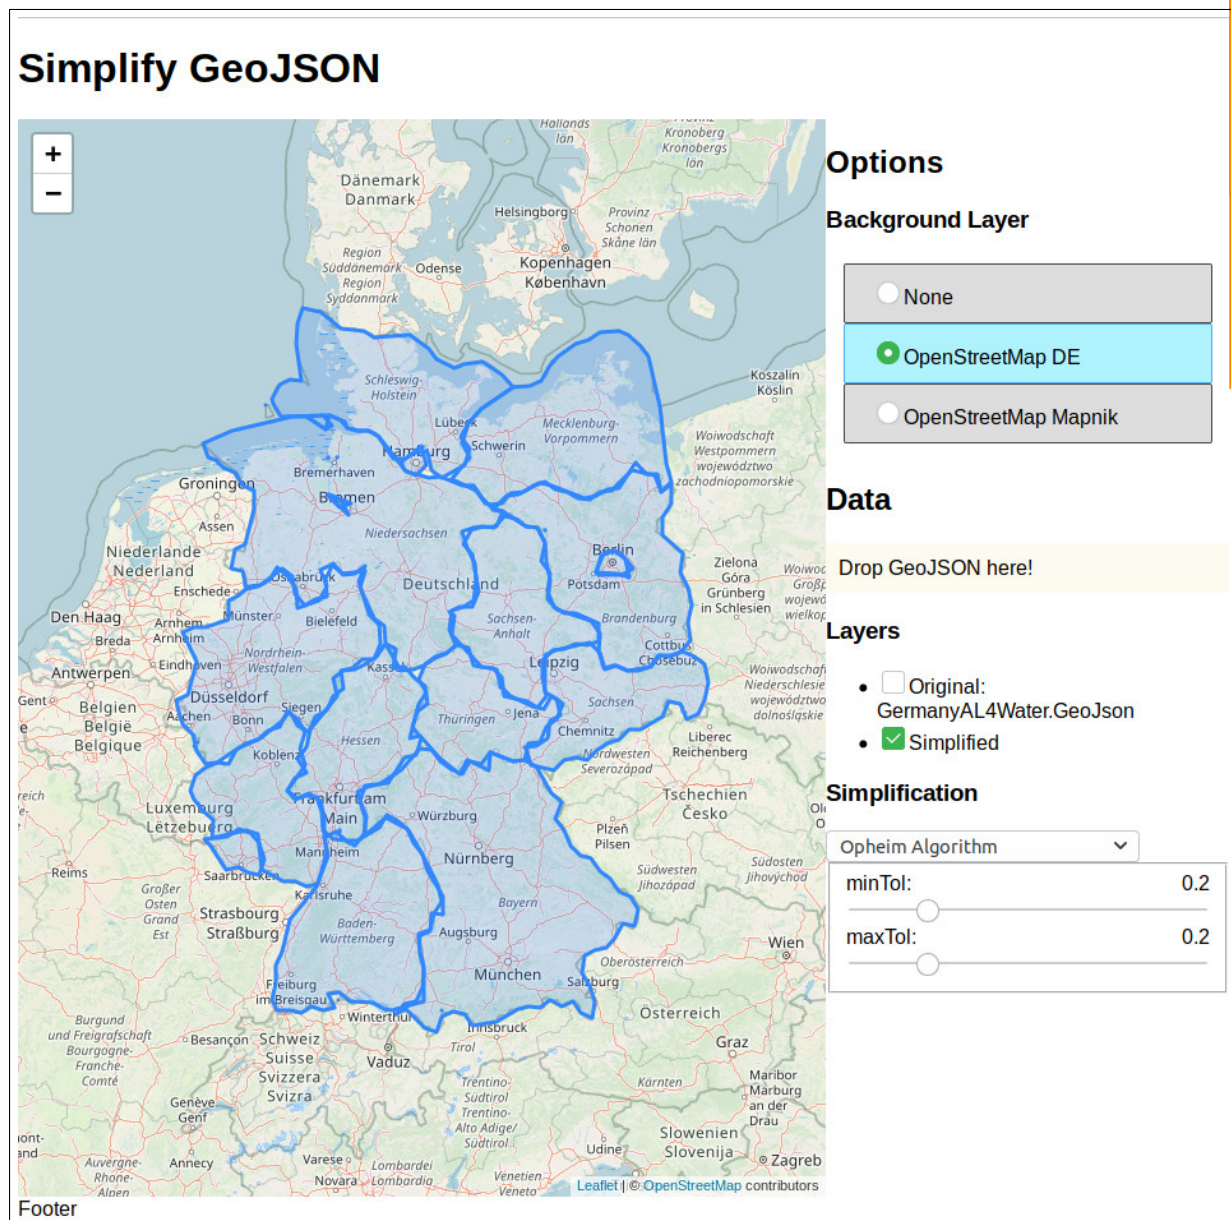
Figure 22: The state model of the application

value, so the algorithm can compute its parameters array at any time. The fields also define additional restrictions in their `props` attribute like the number range from which to choose from. An integer field for example, like the `n` value in the *Nth point* algorithm, would instantiate a range field with a step value of one. The **ToleranceRange** however, which is modeled as its own subclass due to its frequent usage, allows for smaller steps to represent decimal numbers.

FeatureState encapsulates the state of the vector features. Each layer is represented in text form and object format of the GeoJSON standard. The text form is needed as a serializable form for detecting whether the map display needs to update on an action. As the original features come from file or the server, the text representation is the source of truth and the object format derives from it. The simplified features are asynchronously calculated. This process is outsourced to a debounced reaction that updates the state upon finish.

7.4 The user interface

After explaining the state model the User Interface (UI) shall be explained. The interface is implemented in components which are modeled in a shallow hierarchy. They represent and update the application state. In listing 23 the resulting web page is shown. The labeled regions correspond to the components. Their behavior will be explained in the following.



Insert
final pic-
ture.

Red
boxes
around
regions

Make ui
fit de-
scription

Figure 23: The user interface for the algorithm comparison. (not final)

Leaflet Map The big region on the left marks the Leaflet map. Its main use is the visualization of Features. The layers to show are one background tile layer, the original and the simplified features. Original marks the user specified input features for simplification. These are marked in blue with a thin border. The simplified features are laid on top in a red styling. Aside from the default control for zooming on the top left the map contains a info widget showing the length of the currently specified tolerance on the top right.

Background Layers Control The first component in the Options panel is a simple radio button group for choosing the background layer of the map or none at all. They are provided by the OpenStreetMap (OSM) foundation²⁶. By experience the layer "OpenStreetMap DE" provides better loading times in Germany. "OpenStreetMap Mapnik" is considered the standard OSM tile layer²⁷.

Data Selection Here the input layer can be specified. Either by choosing one of the prepared data sets or by selecting a locally stored GeoJSON file. The prepared data will be loaded from the server upon selection by an Ajax call. Ajax stands for asynchronous JavaScript and XML and describes the method of dispatching an HTTP request from the web application without reloading the page. This way not all of the data has to be loaded on initial page load. On the other hand the user can select a file with an HTML input or via drag & drop. For the latter the external package "file-drop-element" is used²⁸. It is a custom element based on the rather recent Custom Elements specification²⁹. It allows the creation of new HTML elements. In this case it is an element called "file-drop" that encapsulates the drag & drop logic and provides a simple interface using attributes and events. Listing 13 shows the use of the element. The mime type is restricted by the `accept` attribute to GeoJSON files.

```
1 <file-drop accept="application/geo+json">Drop area</file-drop>
```

Listing 13: The file-drop element in use

Layer Control This element serves the purpose of toggling the display of the vector layers. The original and the simplified features can be independently displayed or be hidden. If features have been loaded, the filename will be shown here.

²⁶https://wiki.osmfoundation.org/wiki/Main_Page

²⁷https://wiki.openstreetmap.org/wiki/Featured_tile_layers

²⁸<https://github.com/GoogleChromeLabs/file-drop#readme>

²⁹<https://w3c.github.io/webcomponents/spec/custom/>

Simplification Control The last element in this section is the control for the simplification parameters. At first the user can choose if a conversion to TopoJSON should be performed before simplification. Then the algorithm itself can be selected. The parameters change to fit the requirements of the algorithm. The update of one of the parameters trigger live changes in the application state so the user can get direct feedback how the changes affect the geometries.

List of Figures

1	UML diagram of the benchmarking application	18
2	The state machine for the benchmark suite	19
3	The user interface for benchmarking application.	21
4	The Simplify.js test data visualized	22
5	The Bavaria test data visualized	23
6	Simplify.wasm vs. Simplify.js benchmark result of Windows device with Chrome browser on dataset "Simplify.js example" without high quality mode.	25
7	Simplify.wasm vs. Simplify.js benchmark result of Windows device with Firefox browser on dataset "Simplify.js example" without high quality mode.	25
8	Simplify.wasm vs. Simplify.js benchmark result of Windows device with Edge browser on dataset "Simplify.js example" without high quality mode.	26
9	Simplify.wasm vs. Simplify.js benchmark result of Windows device with Chrome browser on dataset "Simplify.js example" with high quality mode.	26
10	Simplify.wasm vs. Simplify.js benchmark result of Windows device with Firefox browser on dataset "Simplify.js example" with high quality mode.	27
11	Simplify.wasm vs. Simplify.js benchmark result of Windows device with Edge browser on dataset "Simplify.js example" with high quality mode.	27
12	Simplify.wasm runtime analysis benchmark result of Windows device with Edge browser on dataset "Simplify.js example" without high quality mode.	28
13	Simplify.wasm runtime analysis benchmark result of Windows device with Edge browser on dataset "Simplify.js example" with high quality mode.	29
14	Simplify.wasm vs. Simplify.js benchmark result of MacBook Pro device with Safari browser on dataset "Bavaria outline" without high quality mode.	30
15	Simplify.wasm vs. Simplify.js benchmark result of MacBook Pro device with Firefox browser on dataset "Bavaria outline" without high quality mode.	30
16	Simplify.wasm vs. Simplify.js benchmark result of MacBook Pro device with Safari browser on dataset "Bavaria outline" with high quality mode.	31
17	Simplify.wasm vs. Simplify.js benchmark result of MacBook Pro device with Firefox browser on dataset "Bavaria outline" with high quality mode.	31
18	Simplify.wasm vs. Simplify.js benchmark result of Ubuntu device with Firefox browser on dataset "Bavaria outline" without high quality mode.	32
19	Turf.js simplify benchmark result of Ubuntu device with Firefox browser on dataset "Bavaria outline" without high quality mode.	33

20	Simplify.wasm vs. Simplify.js benchmark result of Ubuntu device with Firefox browser on dataset "Bavaria outline" with high quality mode. . . .	33
21	Turf.js simplify benchmark result of Ubuntu device with Firefox browser on dataset "Bavaria outline" with high quality mode.	34
22	The state model of the application	42
23	The user interface for the algorithm comparison. (not final)	43

List of Tables

1	Problem dimensions of Case 1	24
2	Problem dimensions of Case 2	28
3	Problem dimensions of Case 3	29
4	Problem dimensions of Case 4	32
5	Problem dimensions of Case 5	34

Listings

1	An example for a GeoJSON object	5
2	Polyline coordinates in nested-array form	6
3	Polyline in array-of-objects form	6
4	Turf.js usage of simplify.js	12
5	Snippet of the difference between the original Simplify.js and alternative . .	12
6	The compiler call	13
7	My Caption	14
8	The storeCoords function	14
9	Entrypoint in the C-file	15
10	Loading coordinates back from module memory	16
11	One entrypoint to the C++ code	40
12	Emscripten bindings	41
13	The file-drop element in use	44