

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	Heuristics	7
2.3.2	Algorithms	7
2.3.3	Summary	8
2.4	Web runtimes	8
2.4.1	Introduction to Webassembly	8
3	Methodology	10
3.1	Implementation of a performance benchmark	10
3.1.1	State of the art: Simplify.js	10
3.1.2	The webassembly solution	12
3.1.3	The implementation of a web framework	16
3.1.4	External libraries	16
3.1.5	The framework	16
3.1.6	The user interface	18
3.1.7	The test data	20
3.2	Algorithm comparison	20
3.2.1	State of the art: psimpl	20
3.2.2	Compiling to WebAssembly	21
3.2.3	The implementation	22
3.2.4	The user interface	24
4	Results	27
4.1	Benchmark results	27
4.2	Comparing the results of different algorithms	27

5 Discussion	28
6 Conclusion	29
6.1 Enhancements	29
6.2 Future Work	29

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 the feature richness of locally installed applications. Even though these

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>

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

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

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

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


```
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

2.3.1 Heuristics

Positional errors

Length errors

Area Errors

2.3.2 Algorithms

Compression algorithms.

n-th point algorithm

Random-point routine

Radial distance algorithm

Perpendicular distance algorithm

Reumann-Witkam simplification

Opheim simplification

Lang simplification

Douglas-Peucker simplification

with reduction parameter

Jenks simplification

Visvalingam-Whyatt simplification

Zhao-Saalfeld simplification

2.3.3 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

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>

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 algorithms for this task. It will be further introduced in chapter 3.2.1.

Existing compilers

emscripten

assemblyscript

rust

Technical hurdles

Managing memory

passing arrays

3 Methodology

The benefits that WebAssembly promises shall be tested in two separate web pages. One for the performance measurements and one to test the integration of existing libraries.

Performance As it is the most applied algorithm the Douglas-Peucker algorithm will be used for measuring performance. A JavaScript implementation is quickly found. SimplifyJS. It is used by Turf, a geospatial analysis library. To produce comparable results the implementation will be based on this package. A separate library, called SimplifyWASM, written in C will be created that mimics the JavaScript original.

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.

3.1 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.

3.1.1 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.

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

¹⁰<https://leafletjs.com/>

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.

```
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] : [
6       coords.x, coords.y];
7   });
}
```

Listing 4: Turf.js usage of simplify.js

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 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

3.1.2 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.

```
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

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 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.

More
about
the com-
piler call


```

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

```

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

Listing 3.1.2 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 utility 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,
5     wasmUrl)
6   return await emscriptenModule

```

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.

¹¹lib/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>

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

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.

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.

programming:

Check
for
coords
length <
2

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

```

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

3.1.3 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.

3.1.4 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 3.1.2 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 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.

3.1.5 The framework

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 3.1.5 shows the corresponding

¹⁴<https://webpack.js.org/>

¹⁵<https://babeljs.io/>

¹⁶<https://benchmarkjs.com/>

¹⁷<https://reactjs.org/>

¹⁸<https://www.chartjs.org/>

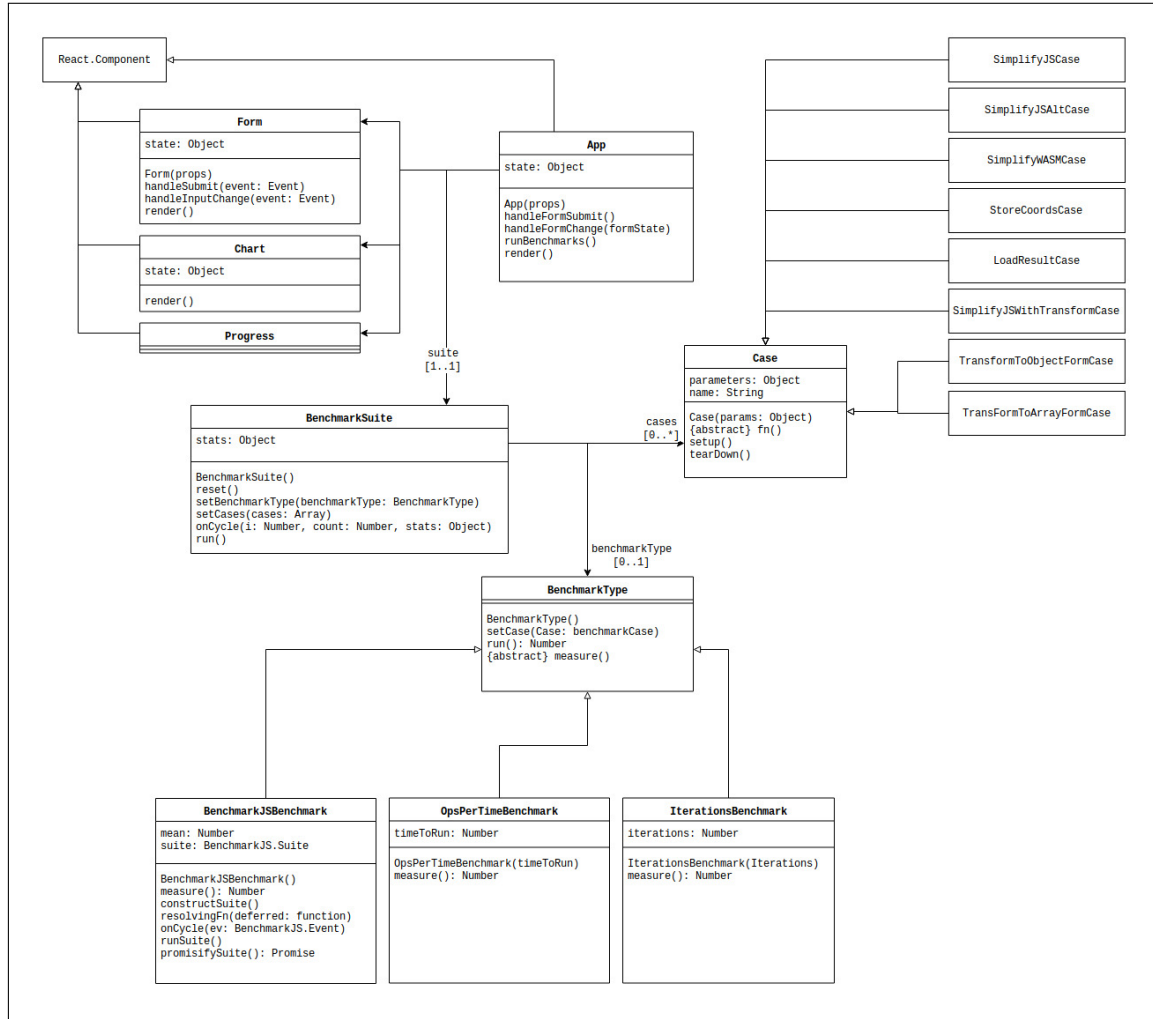


Figure 1: UML diagram of the benchmarking application

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.

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 Sim-

Cases

ply.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`).

- **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 Turf.js.

Cases

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

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.

more
about
Bench-
mark.js

Add digram: state machine for suite

Explain state machine

3.1.6 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.

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

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




Figure 2: The user interface for benchmarking application. (not final)

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 3.1.7. 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 3.1.5. 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 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.

3.1.7 The test data

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

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.

3.2.1 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 3.1.1). 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

3.2.2 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`²¹ and `val.h`²². `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.

```

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

²¹https://emscripten.org/docs/api_reference/bind.h.html#bind-h

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

Compiler call (-bind)

The library code on JavaScript side is similar to the one in chapter 3.1.2. This time a function is exported per routine.

More about javascript glue code with listing callSimplification.

3.2.3 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. MobX²³ 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 Redux²⁴ does. The final state diagram can be seen in listing 3. 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 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

²³<https://mobx.js.org/>

²⁴<https://redux.js.org/>

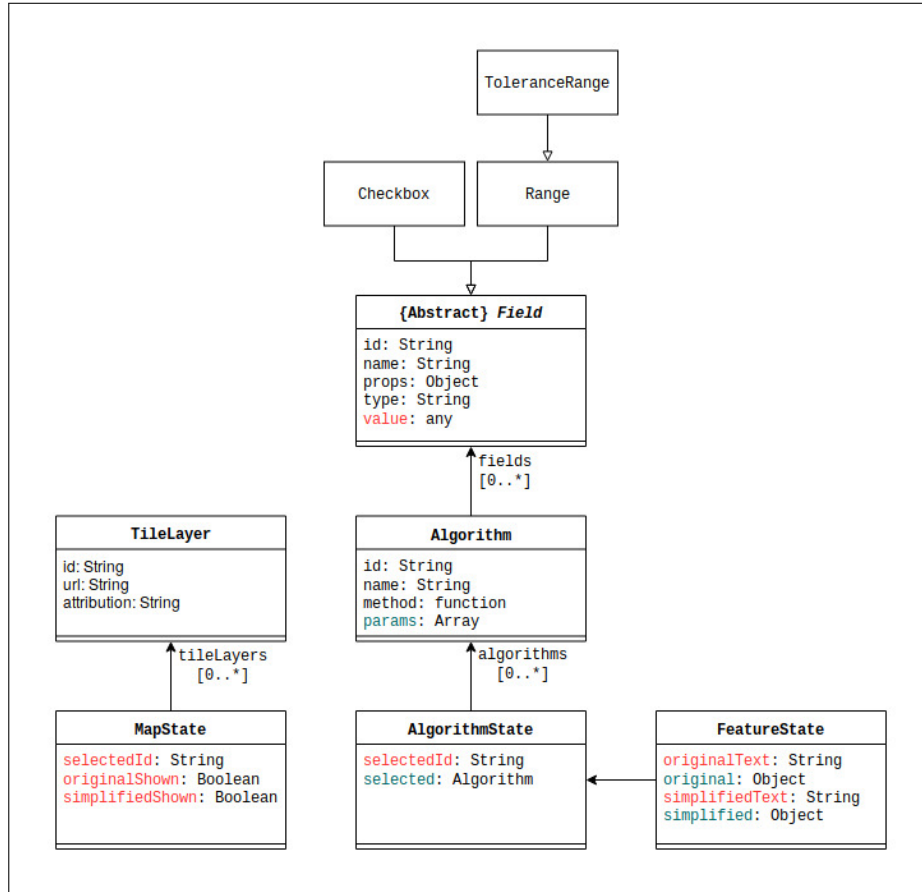


Figure 3: The state model of the application

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.

3.2.4 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 4 the resulting web page is shown. The labeled regions correspond to the components. Their behavior will

be explained in the following.

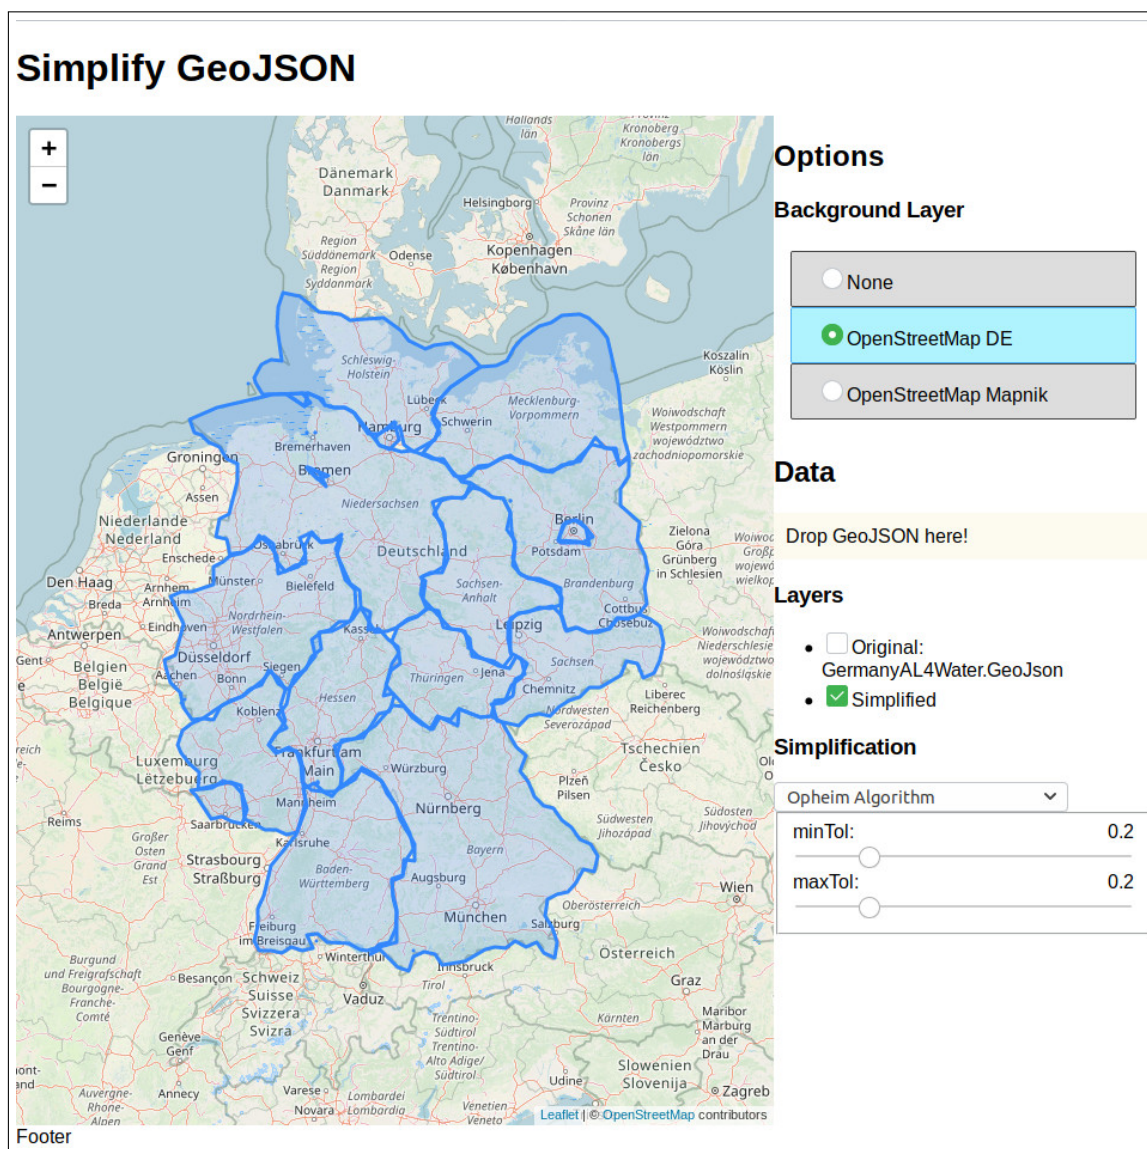


Figure 4: The user interface for the algorithm comparison. (not final)

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 an 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) foundation²⁵. By experience the layer "OpenStreetMap DE" provides better loading times in Germany. "OpenStreetMap Mapnik" is considered the standard OSM tile layer²⁶.

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 used²⁷. It is a custom element based on the rather recent Custom Elements specification²⁸. 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.

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.

²⁵https://wiki.osmfoundation.org/wiki/Main_Page

²⁶https://wiki.openstreetmap.org/wiki/Featured_tile_layers

²⁷<https://github.com/GoogleChromeLabs/file-drop#readme>

²⁸<https://w3c.github.io/webcomponents/spec/custom/>

4 Results

4.1 Benchmark results

4.2 Comparing the results of different algorithms

5 Discussion

6 Conclusion

6.1 Enhancements

Enhancement: Line Smoothing as preprocessing step

6.2 Future Work

List of Figures

1	UML diagram of the benchmarking application	17
2	The user interface for benchmarking application. (not final)	19
3	The state model of the application	23
4	The user interface for the algorithm comparison. (not final)	25

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	11
5	Snippet of the difference between the original Simplify.js and alternative	11
6	The compiler call	12
7	My Caption	13
8	The storeCoords function	14
9	Entrypoint in the C-file	15
10	Loading coordinates back from module memory	15
11	One entrypoint to the C++ code	21
12	Emscripten bindings	22
13	The file-drop element in use	26