

# Using JavaScript and WebCL for Numerical Computations: A Comparative Study of Native and Web Technologies

Faiz Khan

Vincent Foley-Bourgon

Sujay Kathrotia

Erick Lavoie

Laurie Hendren

McGill University

{faiz.khan,vincent.foley-bourgon,sujay.kathrotia,erick.lavoie}@mail.mcgill.ca, hendren@cs.mcgill.ca

## Abstract

From its modest beginnings as a tool to validate forms, JavaScript is now an industrial-strength language used to power online applications such as spreadsheets, IDEs, image editors and even 3D games. Since all modern web browsers support JavaScript, it provides a medium that is both easy to distribute for developers and easy to access for users. This paper provides empirical data to answer the question: *Is JavaScript fast enough for numerical computations?* By measuring and comparing the runtime performance of benchmarks representative of a wide variety of scientific applications, we show that sequential JavaScript is within a factor of 2 of native code. Parallel code using WebCL shows speed improvements of up to 2.28 over JavaScript for the majority of the benchmarks.

**Categories and Subject Descriptors** D.3.4 [Programming Languages]: Processors - Run-time environments

**General Terms** Experimentation, Languages, Measurement

**Keywords** Benchmark; Numerical Computation; Web Browser; JavaScript; WebCL; OpenCL; C; Parallelism; Computational Dwarfs

## 1. Introduction

JavaScript was developed and released by Netscape Communications in 1995 to allow scripting in their browser[43]. By enabling embedded computations, simple tasks such as validating forms and changing the content of HTML elements could be performed in the browser without any interaction with the server. Since that time web applications have become increasingly more sophisticated and compute-intensive. To support this trend, JavaScript technology has improved to include sophisticated virtual machines and JIT compilers[18, 29], and associated technologies like WebCL[15], which provide GPU support.

Web-based JavaScript technologies provide numerous advantages to both application developers and end-users. Due to the increased standardization and support of JavaScript by all modern web browsers[7], application developers have a simple and portable mechanism for distributing their applications. Users can easily run

these applications because they are likely to have modern web browsers installed on their desktops, laptops and mobile devices.

Although one might normally associate web-based computing with e-commerce applications or social networking, we want to explore if JavaScript technology is ready for a wide variety of compute-intensive numerical applications. If so, then JavaScript could be a reasonable way to develop and distribute numerical applications, either through developing JavaScript applications directly, or by using JavaScript as a target in compilers for languages such as MATLAB or R. End-users would have easy access to such applications through the wide variety of computers and mobile devices at their disposal.

In addition to examining sequential performance, we also want to evaluate the performance of parallel code using WebCL, a JavaScript binding for OpenCL.

Our work has three main contributions: (1) the development of Ostrich, an extended benchmark suite which covers a wide spectrum of numerical applications; (2) an examination of the sequential performance of JavaScript on different browsers and platforms; and (3) an examination of parallel performance using WebCL. We elaborate on each of these contributions below.

**Ostrich, an extended benchmark suite:** Colella and a subsequent team at Berkeley have identified important common patterns of numerical computation[16, 24], and suites such as OpenDwarfs [27] and Rodinia[22] have been developed to evaluate native sequential and parallel languages. We have built upon this work to develop a benchmark suite called Ostrich, which can be used to evaluate these applications on sequential JavaScript and for most benchmarks we also provide a parallel version using WebCL.

**Sequential Performance Study:** We study three variants of JavaScript code: hand-written using regular arrays, hand-written using typed arrays, and compiler-generated asm.js. We compare them among themselves and give an assessment of their respective performance. We also compare hand-written JavaScript using typed arrays and asm.js with native C code and show that they are both competitive for numerical computation. Indeed, our results show that for our benchmark suite, on average, JavaScript's performance is within a factor of 2 of C. We also identify the elements of C and JavaScript that can skew the results.

**Parallel Performance Study:** We study the performance of parallel computation by executing our benchmarks with the WebCL and OpenCL technologies. We compare the speedups these technologies offer over their respective sequential version. We show that for a majority of benchmarks can benefit from parallel code in WebCL. We determine there is an incongruence between the two frameworks, however the problem is due to technical details in the current implementation of WebCL. The

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

DLS '14, October 20–24, 2014, Portland, OR, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3211-8/14/10...\$15.00.

<http://dx.doi.org/10.1145/2661088.2661090>

parallel tests are performed both on multi-core CPUs and on GPUs.

The rest of this paper is organized as follows: Section 2 gives a brief introduction about different technologies, including JavaScript, asm.js, and WebCL. Section 3 introduces Ostrich, our new extended benchmark suite. Section 4 describes the experimental methodology used for our performance studies, and the machine configurations used. Section 5 and Section 6 respectively discuss the performance results of sequential and parallel versions of the benchmarks. Section 7 cites the related works in the area of JavaScript performance measurement and Section 8 concludes the paper.

## 2. Background

In this section we provide some key background on the technologies we are using for our experiments. Our study targets the JavaScript language from a sequential perspective and a parallel perspective. On the sequential side we analyze regular JavaScript as well as an important subset of the language called asm.js. We employ a compiler that targets this subset called Emscripten. We use native C and C++ code as a baseline for our measurements. On the parallel side we compare a framework for JavaScript called WebCL to its native counterpart, OpenCL. In the remainder of this section we elaborate on these technologies.<sup>1</sup>

### 2.1 JavaScript

JavaScript[7] is a high-level, multi-paradigm language that was originally intended for browser-side scripting. Recently, the browser has become a target for non-traditional applications such as document and spreadsheet editing, graphics, games, and other multimedia. This has been supported by new technologies such as HTML5[6], CSS3[2], and WebGL[14] that interface through JavaScript. Browser engines have evolved and moved from interpreters to JIT compilers. The language has also introduced byte array buffers to provide an interface for dealing with raw binary data. Typed arrays for integer and floating point numbers have also been built on top of these buffers. Since JavaScript is the only language available on the browser, many other mainstream languages such as Java[5], Python[12] and Ruby[10] have compilers that target it. The dynamic features of JavaScript, such as dynamic typing and prototypical inheritance, require sophisticated dynamic optimization techniques to make the application run fast. Up until now, it was unclear whether they would perform well for numerical computation. Our goal is to determine the suitability of JavaScript and browser engines to support numerical applications.

### 2.2 Asm.js

Mozilla has developed a subset of JavaScript called asm.js[1] as a compilation target for native applications written in languages such as C and C++. It is composed of type annotations, specific program structure and low-level instruction like code. The primary advantage of programs written in this subset is that they can be directly compiled to assembly code thereby bringing near native performance to the browser. Currently, Firefox is the only browser that has an ahead-of-time compiler that compiles directly to assembly code. However, since asm.js is a subset of JavaScript, it can be executed in all browsers that support JavaScript. We are interested in determining the performance gap between asm.js code and native C and C++ code. We are also interested in the the gap between asm.js in different browsers as well as the gap between asm.js and regular JavaScript code.

### 2.3 Emscripten

Emscripten[45] is a compiler that translates LLVM bytecode to asm.js. Any language that is compiled to LLVM, such as C and C++ with clang, can then be used by Emscripten to generate asm.js code. Emscripten has been used successfully to compile large software such as the Unreal 4 game engine, the Bullet physics library and the Qt GUI framework[4]. The maturity of both the LLVM and the Emscripten compilers, and their successful application to complex and large code bases, suggest that the resulting asm.js output is of sufficient quality to be used to evaluate the performance of the asm.js subset on JavaScript VMs.

### 2.4 OpenCL

OpenCL[11] is a standard by the Khronos Group describing a framework for heterogeneous parallel computing. It abstracts GPUs, multi-core CPUs and Digital Signal Processors into a single programming model. The code is expressed through a C-like language for a single thread execution, called a kernel, with added constructs such as barriers for synchronization. The data is transferred from the host to the device with the required parameters and then executed on the device. Due to the heterogeneous nature of OpenCL, computations can execute in parallel on multiple devices. As a result the highly parallel nature of devices like GPUs can be harnessed through a more general purpose C-like expression. Since it is a general framework for heterogeneous devices we use it as a native code baseline for our parallel analysis.

### 2.5 WebCL

WebCL[15] is a proposed standard to bring OpenCL to the web. An implementation is available from Nokia as an extension for Firefox. WebCL exposes the OpenCL API to JavaScript with certain restrictions in order to maintain security. The same kernels written for OpenCL may be used with WebCL. It is the first standard that allows for heterogeneous parallel computing in a browser exposing CPUs, GPUs and DSPs. We use it as our target for parallel code in JavaScript because it is the web equivalent of OpenCL. However, certain required security checks, such as array-bound checking, are not implemented by the Nokia prototype. These are mandatory for safe execution in a web browser and would likely add an additional overhead compared to OpenCL.

## 3. Ostrich - an extended benchmark suite

Although many JavaScript benchmark suites have been contributed in the past both by researchers[40] and web browser vendors[8, 9, 38], they are not necessarily representative of the classes of numerical application that could be used by scientists and engineers on the web. Notably, numerical applications are often very compute-intensive, they should leverage the parallel capabilities of modern CPUs and GPUs, and achieve as much sequential performance as possible. At the moment, different proposals for exposing parallelism to web developers [15, 31] and for increasing the speed of sequential JavaScript [1] are being actively developed. There is now a need for introducing a new benchmark suite that is tailored to numerical computations. This will enable the wider web community to assess the current performance state of web technologies and entice web browser vendors to address the bottlenecks that prevent web technologies from being fully competitive with native ones.

Selecting benchmarks is an important, yet delicate task; the benchmarks should have certain desirable properties that we list below.

1. The benchmarks should be *representative* of the core computations found in actual numerical programs;

<sup>1</sup> Readers familiar with these JavaScript technologies may skip this section.

2. The benchmarks should have sufficient *breadth* in the classes of applications that they represent to ensure that the results are generally applicable;
3. The benchmarks should produce *correct* results;
4. The results of a given benchmark written in multiple languages should be *comparable*.

The Thirteen Dwarfs, as described in Section 3.1, is a classification in thirteen groups of numerical algorithms, and are representative of most numerical tasks. By selecting from as many of these dwarfs as possible, we ensure that the benchmarks are both representative (1) and cover a wide range of applications (2). By reusing the benchmarks from existing suites such as Rodinia[22] and OpenDwarfs[27], we profit from the previous work that went into ensuring they were correct (3). Finally, we ensure that results are comparable (4) across languages by using common types and programming techniques across all implementations.

In the remainder of this section we provide the background concerning the computational patterns that we used, existing benchmark suites, and then we provide details of our new extended Ostrich benchmark suite.

### 3.1 Thirteen Dwarfs

Colella identified seven algorithmic patterns that occur in numerical computing that are important for the coming decade[24]. These patterns are referred to as dwarfs. At a high level, a dwarf can be thought of as an abstraction that groups together related but different computational models. An example of a dwarf is Dense Linear Algebra that reflects dense matrix and vector computations and algorithms such as LU decomposition and matrix multiplication. Asanovic et al. expanded the concept to algorithms in Parallel Benchmark Suites, Machine Learning, Databases, Computer Graphics and Games. They identified six additional dwarfs that cover the majority of common algorithms in these domains. As a result, a benchmark suite that implements a majority of these dwarfs can be considered representative of a broad range of numerical computing applications. We apply this philosophy in our Ostrich benchmark suite.

### 3.2 Rodinia and OpenDwarfs

Rodinia[22] and OpenDwarfs[27] are two high performance benchmark computing suites for heterogeneous architectures. Both suites offer benchmarks that are classified by dwarf categories. Rodinia offers implementations of benchmarks in CUDA, OpenMP and OpenCL and OpenDwarfs has implementations in OpenCL. Rodinia is the most mature suite, but does not cover all possible dwarfs. We drew a set of benchmarks from Rodinia and the remainder from the younger OpenDwarfs suite. The map-reduce dwarf did not have an implementation in either suite, so we provide our own.

### 3.3 Ostrich

Our benchmark suite, Ostrich, contains 12 sequential benchmarks and 10 parallel benchmarks covering 12 dwarf categories. All but two benchmarks have five implementations: C, JavaScript, asm.js, OpenCL and WebCL.

The OpenCL implementations for the majority of benchmarks originate from the Rodinia and OpenDwarfs suites. The C and C++ code is extracted from OpenMP code when available, otherwise we provide our own implementation based on the OpenCL code. The C and OpenCL benchmarks were adapted to match the constraints of JavaScript running in the browser. When benchmarks read data from a file, we incorporated data generation into the benchmark since JavaScript running in the browser cannot access the file system. We tested various input sizes across many configurations to

find values that would work in all browsers. We avoided some non-deterministic behaviors in browsers when programs start to run against the memory limitations of the virtual machines. We supplied a custom random number generator, an adaptation of the Octane benchmark suite[9] PRNG, to ensure consistent results across languages. We ensured that all the benchmarks give the same result across all languages.

In addition, we ported the benchmarks to JavaScript and WebCL by hand. The JavaScript code is written as a direct translation of the C and C++ source code, using typed arrays wherever possible. We compile the C and C++ source code to asm.js code using Emscripten. We re-use the OpenCL kernels for WebCL and mirror the C and C++ wrapper code in JavaScript.

We created an infrastructure to automatically run the experiments in a repeatable manner.

The list below describes the different dwarfs and Table 1 describes the specific benchmarks we have selected, and which implementations were used from an existing suite (=), created by the authors (+) and modified by the authors (~).

**Dense Linear Algebra:** Dense matrix and vector operations typically found in BLAS implementations belong to this dwarf. Examples: matrix multiplication, LU decomposition.

**Sparse Linear Algebra:** Sparse Matrices have a large number of zeros. These matrices are stored in special formats in order to improve space and time efficiency. Operations on these special formats belong to the Sparse Linear Algebra dwarf. Examples: sparse matrix multiplication.

**Spectral Methods:** Data is often transformed from the time or space domain into to the spectral domain. Operations involving these transformations describe the Spectral Methods dwarf. Examples: FFT.

**N-body methods:** Computations that deal with the interaction between multiple discrete points fit into the N-body dwarf. Examples: lavamd.

**Structured grids:** Data is often organized in regular structured multi-dimensional grids. Computations that proceed as a sequence of upgrade steps describe this dwarf. Examples: image manipulation.

**Unstructured grids:** Data can also be organized in irregular grids or meshes. In this dwarf Operations involve a grid element being updated from neighbors within an irregular grid. Examples: neural networks, global illumination.

**Map Reduce:** Computations that run multiple independent trials and aggregate the result at the end describe this dwarf. Examples: Page Rank, distributed search.

**Combinatorial logic:** This dwarf covers computations that involve a large number of simple operations that exploit bit-level parallelism. Examples: hashing, cryptography.

**Graph traversal:** This dwarf typically involves the traversal of a number of objects and examination of their characteristics. Examples: depth-first search, breadth-first search.

**Dynamic programming:** Typical dynamic programming involves solving large problems by using the solutions of subproblems. Examples: DNA alignment, query optimization.

**Backtrack and Branch-and-bound:** Certain problems involve large search spaces. In many cases searching the entire space is infeasible. One technique is to calculate bounds on subregions in the space in order to identify areas that can be ignored. Applications involving this technique fit in to the Backtrack and Branch-and-bound category. Examples: Integer programming, Chess solver.

Benchmark	Dwarf	Provenance	Description	C	JS	OpenCL	WebCL
<i>back-prop</i>	Unstructured grid	Rodinia	a machine-learning algorithm that trains the weights of connecting nodes on a layered neural network	=	+	=	+
<i>bfs</i>	Graph traversal	Rodinia	a breadth-first search algorithm that assigns to each node of a randomly generated graph its distance from a source node	=	+	=	+
<i>crc</i>	Combinatorial logic	OpenDwarfs	an error-detecting code which is designed to detect errors caused by network transmission or any other accidental error	+	+	=	+
<i>fft</i>	Spectral methods	OpenDwarfs	the Fast Fourier Transform (FFT) function is applied to a random data set	+	+	=	
<i>hmm</i>	Graphical models	OpenDwarfs	a forward-backward algorithm to find the unknown parameters of a hidden Markov model	+	+	=	+
<i>lavamd</i>	N-body methods	Rodinia	an algorithm to calculate particle potential and relocation due to mutual forces between particles within a large 3D space	=	+	=	
<i>lud</i>	Dense linear algebra	Rodinia	a LU decomposition is performed on a randomly-generated matrix	~	+	=	+
<i>nqueens</i>	Branch and bound	OpenDwarfs	an algorithm to compute the number of ways to put down $n$ queens on an $n \times n$ chess board where no queens are attacking each other	+	+	=	+
<i>nw</i>	Dynamic programming	Rodinia	an algorithm to compute the optimal alignment of two protein sequences	~	+	=	+
<i>page-rank</i>	Map reduce	Authors	the algorithm famously used by Google Search to measure the popularity of a web site	+	+	+	+
<i>spmv</i>	Sparse linear algebra	OpenDwarfs	an algorithm to multiply a randomly-generated sparse matrix with a randomly generated vector	+	+	=	+
<i>srad</i>	Structured grid	Rodinia	a diffusion method for ultrasonic and radar imaging applications based on partial differential equations	=	+	=	+

**Table 1.** Benchmarks: provenance, description and implementation details. + indicates a new implementation, = means the algorithm was used without modification, ~ indicates that some changes were made to the algorithm.

**Graphical models:** Computations that run on top of a graph that has random variables as nodes and probabilities as edges fit into this dwarf. Examples: Hidden Markov models, Viterbi decode.

The source code, scripts, and instructions for repeating our experiments are publicly available in a GitHub repository at <http://github.com/Sable/Ostrich/wiki>.

## 4. Methodology

A web browser adds another layer of virtualization between a program and the machine hardware, in addition to the operating system. The browser safely executes code from untrusted sources alongside trusted programs on the same machine and provides a programming interface that is portable across many devices and hardware implementations. That extra layer can potentially lead to major performance degradations compared to running the same program natively. Much industrial effort in web browsers goes into ensuring the performance penalty is minimized to provide a maximum amount of hardware performance to the program.

In this paper, we quantify the overhead of web over native technologies for numerical programs using our Ostrich benchmark suite

and we study the performance impact of some JavaScript language features. In this section, we identify relevant research questions and group them into sequential and parallel sets. Answers to the sequential set provide web developers of numerical applications and compiler writers for numerical languages that target JavaScript with information on the performance currently available on the latest versions of major web browsers. Answers to the parallel set provide a glimpse of the upcoming performance of general-purpose parallel computing using WebCL, which is still at the prototype phase and currently only supported by the Firefox browser. They should help its current proponents to focus their efforts on minimizing the performance overhead of providing additional security guarantees compared to OpenCL.

### 4.1 Research questions

For the sequential experiments, we study the performance of hand-written JavaScript, hand-written JavaScript using typed arrays, and machine generated asm.js code and we report how they fare against C. Typed arrays provide a mechanism for accessing raw binary data much more efficiently than with JavaScript strings. In addition to using typed arrays, asm.js code introduces specialized Java-

Script operations that helps the JIT compiler to better perform type-inference and it performs manual memory management. In a sense, it is a C-like subset that provides a bytecode-equivalent target for compilers. The following sequential questions therefore help both quantify the overhead of sequential code and quantify the performance impact of some key JavaScript features for numerical computation:

- *SQ1*: Is the performance of JavaScript competitive with C?
- *SQ2*: Do typed arrays improve the performance of JavaScript code?
- *SQ3*: Does the asm.js subset offer performance improvements over hand-written JavaScript?

WebCL is the first general-purpose and low-level parallel programming language that is proposed for the web platform and has the potential to provide performance near what can be achieved with the native OpenCL technology. For the parallel experiments, we are interested in questions regarding how WebCL compares to sequential JavaScript and if WebCL's gains compare to the speedups obtained in OpenCL:

- *PQ1*: Does WebCL provide performance improvement versus sequential JavaScript?
- *PQ2*: Does WebCL provide performance improvements for JavaScript which are congruent with the performance improvements of OpenCL versus C?

In the following sections, we describe the configurations of machines and browsers we used to answer the research questions and provide details on measurements methods for repeatability.

## 4.2 Experimental set-up

To represent the computer that a typical web user owns, we use two consumer-grade machines, a desktop and a laptop, to perform the sequential benchmarks; the relevant hardware and software specs are detailed in Table 2. For the parallel benchmarks, we use two desktop machines with different GPUs that our research group uses as compute servers; their specs are detailed in Table 3.

	Desktop	MacBook
CPU	Intel Core i7, 3.10GHz × 8	Intel Core 2 Duo, 2.66GHz × 2
Cache	12 MiB	4 MiB
Memory	8 GiB	4 GiB
OS	Windows 7 Enterprise	Mac OS X 10.8.5
GCC	4.6.3	llvm-gcc 4.2.1
Emscripten	1.12.0	1.12.0

**Table 2.** Specifications of the sequential machines.

The benchmarks for the sequential experiments come in four distinct configurations listed below.

- Native code compiled from C with *gcc -O3*;
- asm.js code compiled from C with *emscripten -O2* and executed in Chrome, Firefox, Internet Explorer, and Safari;
- Hand-written JavaScript using regular arrays and executed in Chrome, Firefox, Internet Explorer, and Safari;
- Hand-written JavaScript using typed arrays and executed in Chrome, Firefox, Internet Explorer, and Safari.

For the parallel part, the benchmarks come in two configurations:

- Native code with an OpenCL kernel;

	Tiger	Lion
CPU	Intel Core i7 930, 2.80GHz × 8	Intel Core i7-3820, 3.60GHz × 8
Cache	8 MiB	10 MiB
GPU	NVIDIA Tesla C2050	AMD Radeon HD 7970
Memory	6 GiB	16 GiB
OS	Ubuntu 12.04 LTS	Ubuntu 12.04 LTS
GCC	4.6.3	gcc 4.6.3

**Table 3.** Specifications of the parallel machines.

- JavaScript code with an OpenCL kernel made available through WebCL.

All our tests are performed against version 1.0.0 of the Ostrich benchmark suite.

## 4.3 Measurements

The sequential, native code version compiled from C serves as the performance baseline for the sequential and parallel experiments. In the sequential case, we are interested in how JavaScript fares against typical numerical code which is most often written in languages that compile to native code. In the parallel case, we are interested in the speed-up that a parallel implementation can provide against a sequential one.

For the sequential experiments, we test hand-written JavaScript and asm.js code generated by compiling C code with Emscripten. The code is run on Firefox 34, Chrome 38 and Internet Explorer 11 on Windows 7 and in Firefox 34, Chrome 38, and Safari 6.1.5 on MacOS X. All executions use only a single core on both machines.

At the time of writing, the Firefox and Chrome versions are the Nightly and Canary builds respectively. We chose these versions for two reasons: 1) they possess the most up-to-date JIT compilers, 2) on Windows, they are available as 64-bit executables while the current stable version (Firefox 31 and Chrome 36) are only available as 32-bit executables. On MacOS X, only Safari is available as a 64-bit executable.

For the parallel experiments, we compare WebCL code running in Firefox on Linux against native OpenCL code. Since Firefox is the only browser that supported WebCL at the time of our study, it is the only one we test.

Every benchmark configuration is executed 10 times on each machine and we compute the arithmetic mean of the execution times. We calculate the execution time ratio against the C version and report the geometric mean of the ratios. The execution time covers only the core algorithm of the benchmark, not any initialization or finalization processing. The execution times are measured with *gettimeofday* for C, *Date.now* for asm.js and Safari, and with the high-resolution *performance.now* for JavaScript in Firefox, Chrome, and Internet Explorer.

## 5. Sequential Results

In this section, we compare the performance of different JavaScript variants between themselves and against C to answer the *SQ1*, *SQ2*, and *SQ3*.

### 5.1 JavaScript and asm.js vs C

Let us start by tackling *SQ1* by comparing the performance of JavaScript and asm.js against C. This question is of particular interest because in order to extract the maximum performance out of the hardware, a large number of numerical programs are written in C, C++, and other low-level languages that compile to efficient native code.

The three contestants in this bout are hand-written JavaScript using typed arrays, asm.js generated by compiling C code with

Emscripten, and C code. We start by looking at the results in Figure 1 obtained on our desktop machine using three popular browsers for Windows: Chrome, Firefox, and Internet Explorer.

In Figure 1(a) for Chrome, we observe that 16 out of 24 results have a slowdown below 2x, and no JavaScript program is more than 3 times slower than C. The most interesting benchmark is definitely *lavamd* which is four times faster than C with both asm.js and JavaScript! Comparing the execution profiles of the benchmark reveals that the culprit for JavaScript’s surprising performance is a call to the exponentiation function in a deeply nested loop body. An article on Google Developers explains that a faster `exp()` function was implemented in V8, and that it is faster than the one in the standard system libraries [13]. On average, the slowdowns for Chrome with respect to C for JavaScript and asm.js are 1.62 and 1.35 respectively.

In Figure 1(b) for Firefox, 20 results show a slowdown below 2x and only one asm.js benchmark shows a slowdown higher than 2x. In 6 benchmarks, at least one of the JavaScript versions is faster or nearly equal to C. On average, the slowdowns for Firefox with respect to C for JavaScript and asm.js are 1.65 and 1.17 respectively.

In Figure 1(c) for Internet Explorer, the results are noticeably worse than in the other browsers. For instance, *crc* in JavaScript is 35 times slower than C. The profiling tool of Internet Explorer confirms that more than 98% of the execution time is spent in the *crc32\_8bytes* function, however the tool does not offer more fine-grained information. Interestingly, the asm.js version of the same benchmark is only 3 times slower than C. On average, the slowdowns for Internet Explorer with respect to C for JavaScript and asm.js are 3.11 and 2.81 respectively.

With these results in mind, let us now see if the less powerful MacBook laptop gives similar ratios in Chrome and Firefox. We also look at the performance of Safari, the browser included with OS X.

In Figure 2(a) for Chrome, the ratios are consistent with the ones obtained on the desktop machine. Like in the desktop results, no benchmark shows slowdown over 3x. The mean slowdowns for JavaScript and asm.js are 1.51 and 1.37 respectively.

For Firefox in Figure 2(b), the ratios are also fairly consistent with the desktop numbers. 22 out of 24 benchmarks have a slowdown below 2x and six of the asm.js benchmarks perform at least as fast as the native version. The mean slowdowns for JavaScript and asm.js are both below 2x, 1.55 and 1.11 respectively.

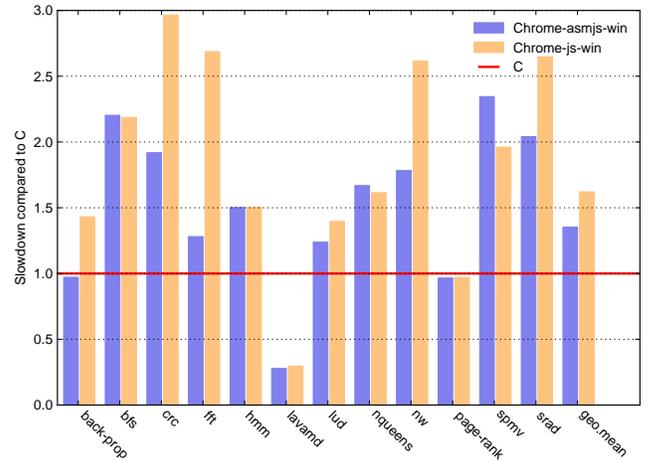
Finally, Safari in Figure 2(c) shows average slowdowns of 2.69 and 2.15 for JavaScript and asm.js respectively. Two benchmarks, *crc* and *fft*, have a severe penalty in JavaScript, 17x and 10x. This high overhead of *crc* and *fft* suggests a performance optimization that is being missed by the Safari JIT compiler.

The results in this section show that, in a modern browser, the performance of JavaScript for most of the numerical benchmarks in our suite is competitive with native code.

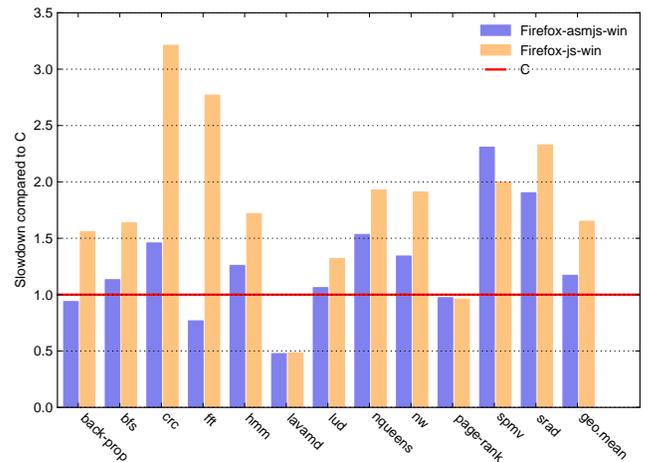
## 5.2 Arrays vs Typed Arrays

Our next experiment addresses *SQ2*. Recently, browsers have started supporting typed arrays which can contain only values of a given type, e.g. a `Float32Array` contains only single-precision floats. In contrast, JavaScript’s regular arrays can contain values of any type, and the values need not be of the same type. For our tests in the previous section we used typed arrays, and in this section we investigate if using these specialized arrays is an important factor for the good overall performance of our benchmarks.

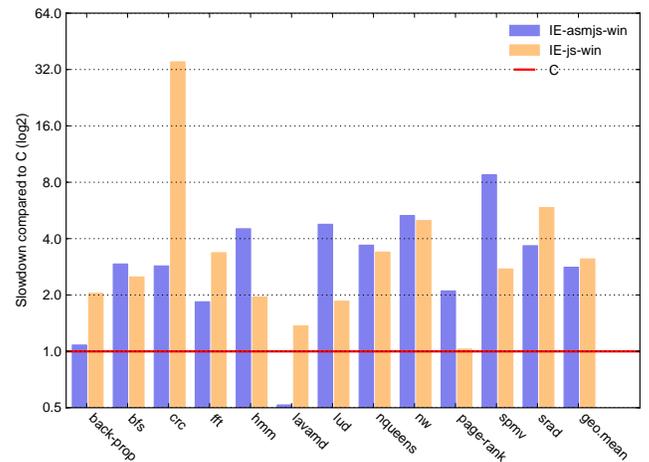
For this experiment, we created JavaScript versions of the benchmarks where we replaced specialized arrays with regular ar-



(a) Chrome’s JavaScript and asm.js ratios compared to C

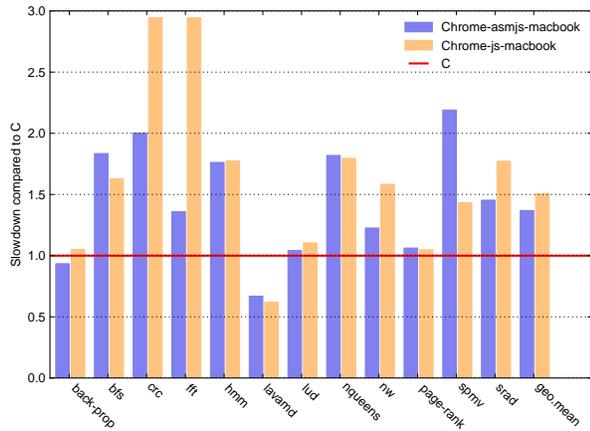


(b) Firefox’s JavaScript and asm.js ratios compared to C

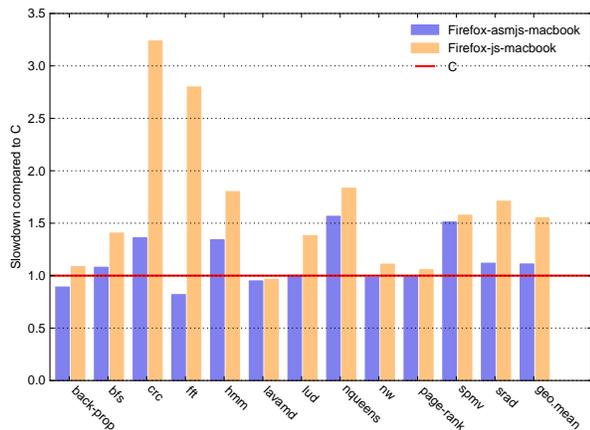


(c) Internet Explorer’s JavaScript and asm.js ratios compared to C

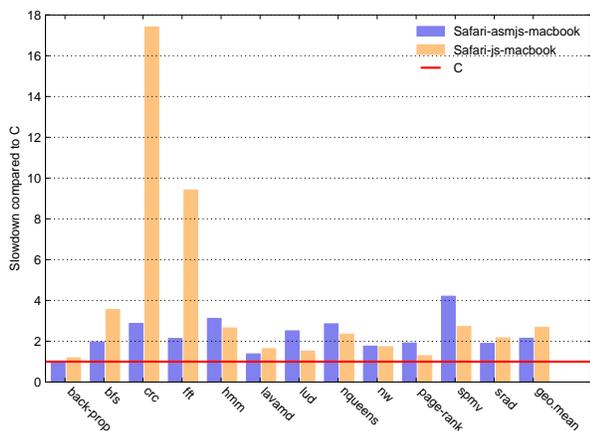
**Figure 1.** Slowdown of JavaScript and asm.js to C on Desktop (lower is better)



(a) Chrome's JavaScript and asm.js ratios compared to C



(b) Firefox's JavaScript and asm.js ratios compared to C



(c) Safari's JavaScript and asm.js ratios compared to C

**Figure 2.** Slowdown of JavaScript and asm.js to C on MacBook (lower is better)

rays.<sup>2</sup> In Figure 3, we report the speedups that we observe by using typed arrays. There is a noticeable improvement for using typed arrays in many of the benchmarks, with a geometric mean speedup hovering around 2x for all browsers on both machines. In some cases, the speedup exceeds 4x and in two cases, the speedup is over 16x! In one benchmark, *fft* in all browsers, the typed arrays version is slower. This is because multiple arrays are allocated inside a recursive function and for typed arrays, all the elements are set to zero, making their initialization slower than a regular array. Note that although the slowdown of Firefox is less severe than Chrome's on Windows, the execution time reveals that Chrome is faster than Firefox by a factor of 3.1.

We also notice that, except for a couple of cases (e.g. *nw* and *back-prop*), the speedups obtained in Chrome, Firefox, and Internet Explorer are similar. The results are consistent on the laptop both between Chrome and Firefox, and between Safari and the other browsers. One noticeable exception is the *fft* benchmark which is significantly slower in Safari when using typed arrays. We note that when using regular arrays on Safari, the input size of *back-prop* and *nw* had to be reduced to prevent swapping. The memory allocation pattern for Safari is non-linear and the amount of memory used by the program in these two cases would go from 1.7GB to 2.7GB for a slight change of input size, exhausting the amount of memory available in main memory on the laptop. We therefore ensured that all benchmarks ran using only main memory.

These results show that a good first step in improving the performance of a hand-written numerical JavaScript program is to replace the declarations of arrays that are manipulated in tight loops to typed arrays if possible. These results give a clear answer to *SQ2*: typed arrays improve the performance of JavaScript programs.

### 5.3 JavaScript vs asm.js

In this section, we will study whether the compiler-friendly asm.js can offer better performance than hand-written JavaScript using typed arrays.

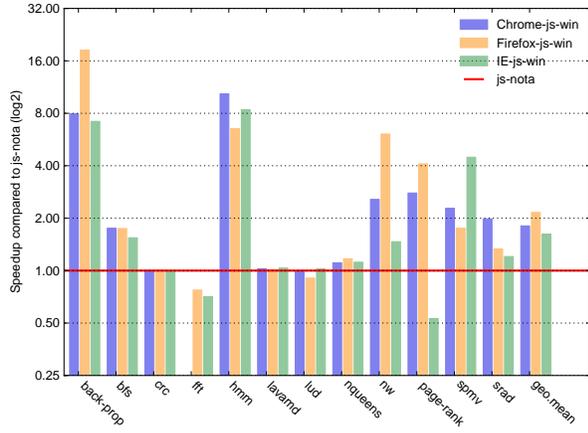
On the desktop, the chart in Figure 4(a) shows that 11 out of 12 asm.js benchmarks in Firefox have equal or better performance than hand-written JavaScript with typed arrays. In Chrome, 11 of the 12 asm.js benchmarks run equally fast or faster than the hand-written JavaScript version. Internet Explorer shows great variability going from 3x slower (*spmv*) to 12x faster (*crc*). The mean speedups are respectable, 1.20 for Chrome and 1.41 for Firefox and 1.11 for Internet Explorer.

On the MacBook, in Figure 4(b), the results are similar. The majority of the asm.js benchmarks in Firefox are faster. The highest improvement is 3.5x for *fft*. Chrome shows speedups similar to the desktop for the majority of the benchmarks. Safari shows impressive speedups of 4.5x and 6x for *crc* and *fft* and slowdowns for *hmm*, *lud*, *nqueens*, *page-rank* and *spmv*. The mean speedups for Chrome, Firefox, and Safari are 1.1, 1.39, and 1.25 respectively.

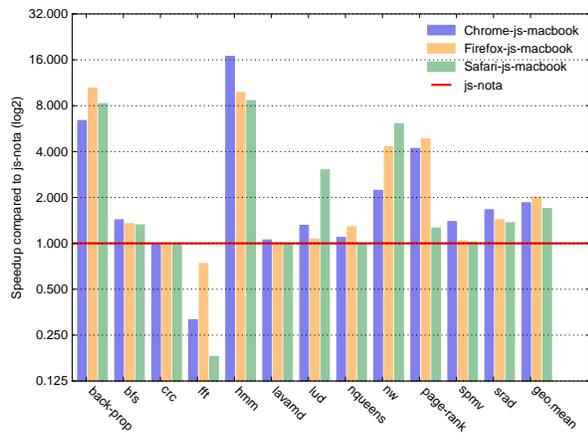
Since asm.js also uses typed arrays, these results are a strong indication that the one of the most important factor in improving the performance of numerical JavaScript programs is to use typed arrays when appropriate. They also provide an answer to *SQ3*: asm.js offers performance improvements over regular JavaScript.

Let us note that although we compare their performance, JavaScript and asm.js address two different approaches to writing web applications. Asm.js is not a human-friendly format, and should be generated by a compiler; it is most useful when the developers of an application written in a language like C want to port it to the web without having to do a complete rewrite. JavaScript and typed arrays are going to be most interesting to developers who want to

<sup>2</sup>Some typed arrays were kept if they were necessary to obtain correct results, e.g. when doing unsigned 32-bit integer manipulations.



(a) Desktop



(b) MacBook

**Figure 3.** Speedup obtained by using typed arrays (js) compared to the same code without typed arrays (js-nota) (higher is better)

directly write their application for the web and want to interact with other JavaScript libraries and services.

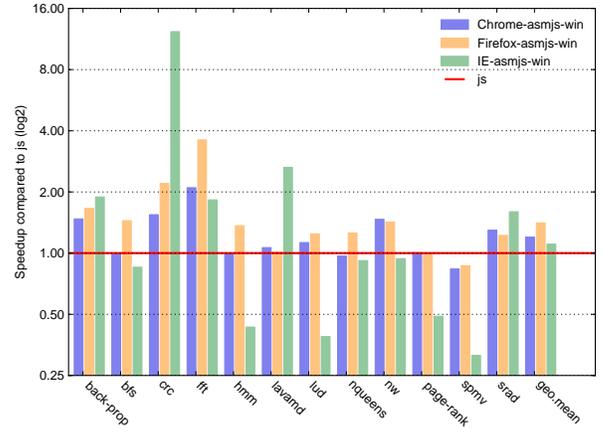
#### 5.4 Other Observations

Although the results presented in the previous sections show that JavaScript’s execution speed is competitive with C, there are other considerations that may need to be taken into account when using it to write numerical code.

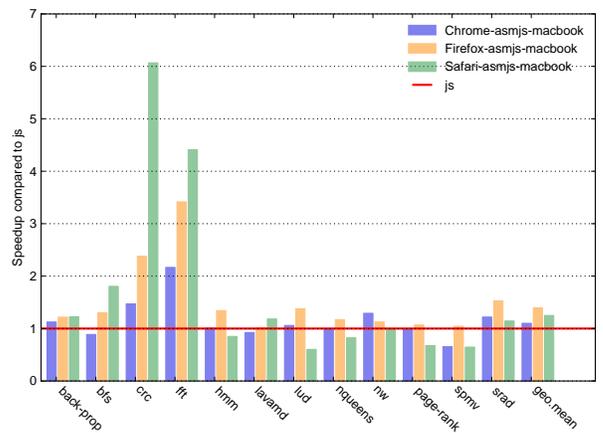
The first point to be aware of is that not all instance sizes that can be handled by C can be handled by JavaScript. *Crc* and *bfs* are examples of benchmarks where we had to reduce the size of the input to accommodate JavaScript. In *bfs*, we initially had a graph of 4 million nodes, and it could be traversed by C, in both Chrome and Firefox with *asm.js*, and in Firefox with JavaScript and typed arrays. However, the execution never finished in Chrome, the process just stopped. We had to reduce the number of nodes to 3 million in order to obtain results in Chrome.

The choice of browser can also greatly affect how fast a program executes: the JavaScript implementation of *crc* executes about 10 times faster in Chrome and Firefox than in Internet Explorer for example.

Another point to consider when creating a numerical application for the web is how the browser reacts to long computations. In



(a) Desktop



(b) MacBook

**Figure 4.** Speedup of *asm.js* against JavaScript using typed arrays (higher is better)

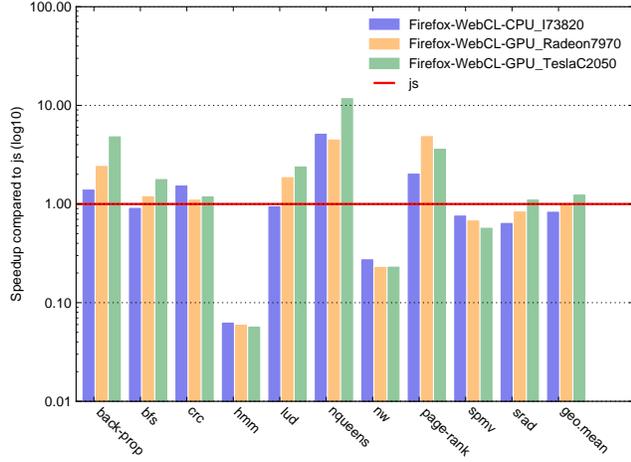
Chrome, with its one process per tab design[39], a long running computation in one tab does not affect the ability of the user to perform tasks in other tabs. In Firefox, a long running JavaScript computation will freeze the entire browser until it has finished.

#### 5.5 Summary

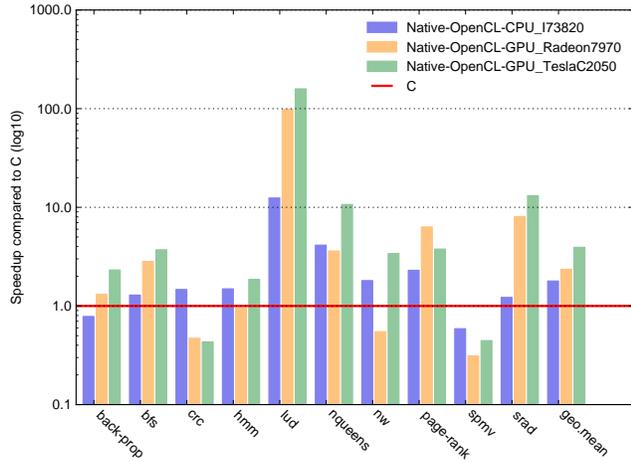
In this section, we have shown that the developer writing a numerical JavaScript application directly and the developer compiling his C application to *asm.js* will both obtain code that can be easily distributed to their users and is often reasonably competitive with the performance of native code (*SQ1*). We have shown that using typed arrays where applicable is an easy and effective way of improving the performance of a JavaScript application, especially if there are a lot of array accesses in the hot loops (*SQ2*). We have shown that *asm.js* delivers on its promise of being an efficient target for compilers; depending on the browser, it offers a speedup between 11% and 40% over hand-written JavaScript (*SQ3*).

### 6. Parallel Results

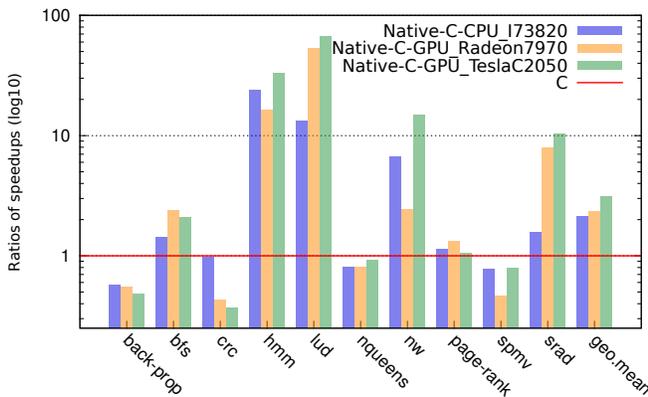
In this section we study the performance of benchmarks in WebCL. We compare how they perform against sequential JavaScript using typed arrays to answer *PQ1*. Then we answer *PQ2* by comparing the congruency of WebCL gains to OpenCL gains.



(a) Speedup of WebCL over JavaScript (higher is better)



(b) Speedup of OpenCL over C (higher is better)



(c) WebCL-JS speedup over OpenCL-C speedup (higher is worse)

Figure 5. Parallel performance results for WebCL and OpenCL

## 6.1 WebCL Parallel Performance Against Sequential JavaScript

We start with answering *PQ1*. In Figure 5(a), we present the speedup of WebCL over JavaScript. The TeslaC2050 experiments were done on the Tiger machine and the i7-3820 along with the Radeon-7970 were done on the Lion machine mentioned in Section 4.2. The entire set of benchmarks shows means of 0.81, 1.01 and 1.21 for i7-3820, Radeon-7970 and Tesla-C2050, respectively. Only the Tesla platform seems to offer a performance gain on average. It seems that WebCL does not have an advantage but we note that not all algorithms benefit equally from parallel implementations. With this in mind, we consider the effect of the slowest benchmarks and the performance gains over the most parallel benchmarks.

The two slowest benchmarks are *hmm* and *nw*. We explore the reasons for their slow performance in Section 6.2. *Hmm* is slower by a factor of 16.67 and *nw* is slower by a factor of 4.17 over all platforms. Removing these two benchmarks from our analysis yields mean speedups of 1.34, 1.74, 2.28 for i7-3820, Radeon-7970 and Tesla-C2050, respectively. With the majority of benchmarks offering speedups, there is an observable advantage in having a parallel platform such as WebCL.

The benchmarks that are most improved by using WebCL are *nqueens*, *page-rank* and *back-prop*. The *nqueens* kernel is highly optimized thus yields significant gains. *Page-rank* is a map-reduce problem that is embarrassingly parallel, therefore it performs better on a parallel platform. Similarly, *Back-prop* also falls on a parallel side as weights between the neural network layers can be calculated independently and is reflected in the speedup. The mean speedup of these three benchmarks are 2.42, 3.72 and 5.85 for i7-3820, Radeon-7970 and Tesla-C2050, respectively. It is clear that certain benchmarks are better suited for WebCL over others.

It is not surprising that there is a spectrum of performance gains from parallel code since some code is more parallel in nature than other code. With respect to *PQ1* we see that there are significant gains that can be made with WebCL, albeit they are dependant on the application.

## 6.2 WebCL Specific Performance

If we observe Figure 5(b), we see the speedups offered by OpenCL over C. The average OpenCL speedups are 1.79, 1.99 and 3.93 for i7-3820, Radeon-7970 and Tesla-C2050, respectively. If we remove the two slowest benchmarks, *hmm* and *nw*, the speedups are 1.83, 2.56 and 4.40. The best performers are *back-prop*, *lud*, *nqueens*, and *page-rank*; the average speedup on i7-3820, Radeon-7970 and Tesla-C2050, for these four benchmarks, are 3.10, 7.36 and 11.03 respectively. OpenCL is certainly giving us a higher return over C than WebCL is over JavaScript, we will quantify this observation further.

Figure 5(c) presents the ratio of the two speedups: the speedup of OpenCL over C against the speedup of WebCL over JavaScript. This offers us a perspective on how the gains realized in OpenCL carry over to WebCL. A higher bar means that OpenCL provided a greater performance improvement than WebCL. The WebCL gains over JavaScript were 2.50 times lower than the gains OpenCL realized over C. We find this to be interesting because although WebCL is a binding to OpenCL there seems to be excessive overhead somewhere in between.

Figure 5(c) shows that half of the benchmarks suffer significant execution penalties in WebCL. These benchmarks are *bfs*, *hmm*, *lud*, *nw* and *srad*. Compared to the OpenCL implementations, these benchmarks have mean slowdowns of 1.93, 23.42, 36.16, 6.20 and 5.04, respectively. To find the source of these slowdowns, we profiled the JavaScript code that wraps the kernels. Our investigation

Function	back-prop	bfs	crc	hmm	lud	nqueens	nw	page-rank	spmv	srad	average
setArg	4.1%	5.6%	20.1%	45.7%	34.2%	4.7%	41.9%	0.6%	24.2%	15.0%	19.6%
._validateKernel	1.1%	16.7%	10.9%	21.0%	25.6%	1.2%	22.1%	22.2%	7.1%	41.1%	16.9%
device.getInfo	-	-	1.9%	-	5.4%	0.5%	3.1%	3.6%	0.7%	6.5%	2.2%
queue.getInfo	0.4%	-	1.5%	-	2.4%	0.2%	2.8%	4.2%	1.0%	6.1%	1.5%
readBuffer	15.4%	11.2%	9.7%	7.6%	-	1.7%	-	10.7%	32.6%	12.5%	10.1%
writeBuffer	73.7%	33.3%	9.4%	7.6%	-	7.0%	-	10.4%	28.8%	-	17.0%
Total	94.7%	66.7%	53.4%	82.0%	67.7%	15.4%	69.9%	51.6%	94.3%	81.2%	67.7%

**Table 4.** Percentage of WebCL execution spent on setup functions (results that are close to zero are marked as “-”)

revealed that six functions were taking up most of the execution time in the code for the worst performing benchmarks.

Table 4 shows the name and total run time percentage for the most time consuming functions on all of the benchmarks<sup>3</sup>. The first function, `setArg`, points the kernel to the address in memory for a specific parameter to the kernel. The `._validateKernel` function is an internal function that ensures that the kernel is a valid object, compiled and associated with the right context. The `device` and `queue.getInfo` functions are information query functions for a specific device and the command queue working on a device. The buffer read and write functions are responsible for reading from and writing to device memory from the host.

Host and device memory transfers are a known bottleneck in parallel applications, and we will examine them separately, and focus solely on the `setArg`, `._validateKernel`, and `getInfo` functions. The average chunk of execution taken up by `setArg` is 19.6% for all benchmarks and the worst cases, `hmm` and `nw`, have 45.7% and 41.9% of their execution taken up by it. Similarly, `._validateKernel` had a mean chunk of 16.9% and the worst performer, `srad` had over 40% of its execution taken up by it. The `getInfo` functions in comparison on average took less than 3% of the execution but over 5% for `lud` and `srad`. Profiling the OpenCL equivalents of these functions showed negligible time (i.e. they were too quick for the profiler to give non-zero results) on all benchmarks. This suggests these functions are bottlenecked in the extension or through the browser. Since these slow functions are mainly used for setting up the kernels and could be improved through caching or other optimizations.

Now we address the memory transfer operations, `readBuffer` and `writeBuffer`. Looking at Table 4 we notice that these operations take up a significant portion of the time on all benchmarks, including the ones that showed gains or smaller losses by WebCL over JavaScript compared to their native counter parts. The mean chunks we were 10.1% and 17.0% for `readBuffer` and `writeBuffer`, respectively. The worst performer `back-prop` has 70% of its execution taken by `back-prop`. Profiling the equivalent functions on OpenCL showed negligible time for almost the entire set of benchmarks. Out of the detectable benchmarks, for `readBuffer` we saw that `hmm` and `srad` were on average 67.8 and 12.4 times slower, respectively, in WebCL than in OpenCL for all platforms. Similarly, the `writeBuffer` operation was 93.45 and 51.5 times slower for `hmm` and `back-prop`; it was negligible for the rest. It is clear that WebCL is slower in these functions than OpenCL.

The mean of the totals shows that these functions took over 65% of running time on average. The WebCL functions are simple bindings to OpenCL therefore we realize there are significant bottlenecks in the extension or the browser that are causing these slowdowns. We were not able to determine these values since the Firefox profiler does not allow for traversal into functions past the JavaScript API for an extension.

With regards to *PQ2* we state that currently the gains from WebCL are not congruent with the gains from OpenCL. The WebCL benchmarks showed an overall loss of 2.50. Half of the benchmarks had significant slowdowns on the WebCL side and most of this was due to setup code. This, however, may be addressable through optimizations or a direct implementation in the browser.

### 6.3 Other Performance Observations

Figure 5(c) shows that three benchmarks have better gains in WebCL than OpenCL: `back-prop`, `crc` and `spmv`. Looking back at Figure 5(a) and Figure 5(b) we see that `back-prop` actually executes faster in WebCL than in OpenCL.

Profiling reveals that the `back-prop` implementation slowdown occurs in the sequential portion of the code. The original OpenCL implementation of the benchmark uses a two-dimensional representation of matrices, and has to convert it to a one-dimensional array before passing it to the kernel. In WebCL, we use a typed array, which are always one-dimensional, to represent the matrix, and the conversion was not necessary. The actual kernel execution time was similar.

In `crc` the WebCL execution time is faster than JavaScript, yet OpenCL is slower than C. This illustrates an interesting fact, the cost of overhead of parallelizing was high for the already fast C computations but it was less for the slower JavaScript computations. In such cases, it is beneficial to use WebCL even when it isn't beneficial to use OpenCL.

A similar event occurs in `spmv`. In Figure 5(a) and Figure 5(b), we see that `spmv` had worse performance for the parallel code for both WebCL and OpenCL than their sequential counterparts. Although Figure 5(c) seems to indicate that the WebCL gains are better, in reality WebCL is incurring less of a loss than OpenCL. This happens because the gap between JavaScript and C is wider than between WebCL and OpenCL.

### 6.4 WebCL Limits and Caveats

In this section, we report some observations that we obtained in our experiences with WebCL.

The WebCL specification does not allow for structs, but the same is not true for OpenCL. This imposes a limit on the ease of expression of complex data types in kernels for WebCL. Our `lavamd` benchmark makes liberal use of nested structs and would require a large number of arrays and careful indexing to convert to WebCL. This method would change the nature of the memory access patterns in our data structures and would skew our performance results.

As mentioned in the previous sections, a large portion of the slowdown occurs in the setup code in the WebCL extensions. The effect of this setup code diminishes for very large data sets. Massive data set sizes are a common occurrence in numerical computations. Although WebCL itself is capable of handling large data sets, the same is not true for the browser. The browser memory limitations forced us to reduce the sizes of inputs that we tested. As an

<sup>3</sup>The function names in Table 4 are truncated for readability.

example, even though OpenCL can handle 200 million non-zero elements on *spmv*, the browser often crashed on our machines with the same instance size for WebCL.

OpenCL allows for data to be copied to device memory with an offset. However, no offset can be specified for the host memory. In C, this is fixed with a simple increment of the pointer that is being copied. In JavaScript, a new sub-array view object on the typed array or a new typed array must be created in order to achieve the same result.

## 6.5 Summary

In this section, we have shown that WebCL can offer performance gains over JavaScript when the problem is a good parallel candidate (*PQ1*). We have also shown that currently, the performance gains of WebCL over JavaScript are 2.5x less than the gains of OpenCL over C (*PQ2*).

Overall, we find that WebCL does offer advantages as an extension to JavaScript. It shows that a parallel JavaScript framework can aid in numerical computations, particularly those that are amenable to parallelization. However, the inefficiencies of the current implementation do not make the gains over sequential JavaScript as important as OpenCL does with C.

## 7. Related Work

There has been a surge of interest towards numerical computing in the programming language community recently, with the coming of age of newer languages, such as Julia[19] and the study and improvement of older alternatives, such as R[37] and MATLAB[21]. The increasing and widely available parallelism in hardware, plus the establishment of the Web as a major development and deployment platform, means the ground is fertile for work that helps bring the available performance and convenience of the platform to users of numerical languages. Our work in this paper focused on the comparative study of native and web technologies performance for numerical computing using a benchmark suite inspired by the dwarf categories[16, 27]. In this section we situate our work by discussing contributions from the larger context of understanding and improving the performance characteristics of the Web platform for numerical computing.

A number of benchmark suites have been proposed by vendors over the years to measure the performance of browsers, such as Apple's JetStream[38] (previously SunSpider), Google's Octane[9] (previously V8 benchmarks), and Mozilla's Kraken[8] suites. These benchmark suites are all compute-intensive and their 2010 versions were shown not to be representative of real interactive web applications[42]. This has sparked initiatives to build infrastructure to automatically extract benchmarks from real web applications[41], which has been used to produce a suite of reference applications called JSBench[40]. To the best of our knowledge, however, there is no comprehensive web benchmark suite for the upcoming numerical applications typical of engineering and scientific activities, such as the Rodinia suite[22], which is tailored for multicore CPU and GPU computation. Our collection of benchmark implementations for both sequential (JavaScript) and parallel (WebCL) computation aim to answer that need by taking the Rodinia and OpenDwarfs suites and adding new web implementations.

The advent of multiple numerical computing libraries built entirely in JavaScript, such as *math.js*[25], *sylvester*[23], and *numeric.js*[34] shows the need for numerical computing abstractions in web applications. In addition, there have been numerous initiatives to leverage available parallelism in today's hardware, by extending JavaScript implementations without breaking the sequential semantics. Sluice provides JavaScript with a library that can dynamically recompile programs to target multicores for high perfor-

mance stream processing[28]. ParaScript provides automatic runtime parallelization that is fully transparent to the programmer by optimistically identifying code regions for parallelization, generating on-the-fly parallel code, and speculatively executing it[35]. River Trail[31] extends the array datatype with functional operations that automatically leverage multicores and GPUs for execution. We believe our benchmarks could be adapted in a straightforward manner to compare the relative merits of these and similar approaches, and we hope that they will be useful in that context.

JavaScript is increasingly used as a compilation target and runtime system. Examples include Racket with the Whalesong compiler[44], Dart with *dart2js*[3], and SML with *SML2JS*[26]. Krahn and al. also introduced JavaScript-specific optimization techniques for dealing with modularization constructs with ContextJS [32]. When it is used as a target, compilers need to be aware of the underlying implementation to extract maximum performance as shown by Auler et al.[17]. In addition to performance-oriented contributions, the recent formalization of the semantics of the major parts of JavaScript standard[20] is paving the way for high-assurance translations of other languages to JavaScript with machine-checked implementation. One of the aims of our work is to shed light on the suitability of JavaScript and associated web technologies to serve as a competitive compilation target for numerical languages. Superconductor[36] has shown the suitability of WebCL, WebGL and web workers for large scale data visualization. In complement, our work provides quantitative measurements for the relative merits of both sequential (JavaScript) and parallel (WebCL) computation in the context of numerical computation on the web.

Web browsers are complicated pieces of software and multiple factors make performance evaluation non-trivial. A study of the performance variations in the Mozilla Firefox web browser[33] has shown that external factors to the browser, in particular memory randomization done by the operating system, are the dominant factors that explain noise in performance measurements. Grosskurth and Godfrey[30] introduced a reference architecture for web browsers in 2005, which identify and explain the major sub-system in web browsers such as Firefox. Chrome, a more recent architecture, was presented and studied later in a paper by Reis and Gribble[39]. Our work has shown that the current performance of web technologies is competitive with native technologies on many benchmarks typical of numerical computation. Pinpointing the bottlenecks with regard to the exact services in web browsers will be addressed in future work.

## 8. Conclusions and future work

In this paper, we have presented our benchmark suite, Ostrich, which draws benchmarks from two reputable sources, Rodinia and OpenDwarfs, and adds JavaScript and WebCL implementations. Ostrich provides implementations for numerical algorithms in C, JavaScript, *asm.js*, OpenCL and WebCL for twelve of the thirteen dwarfs.

Using the benchmarks in this suite, we have shown that the performance of both hand-written JavaScript (using typed arrays) and *asm.js* is competitive (within a factor of 2) with C, in a large number of benchmarks (*SQ1*). We have shown that to achieve this performance, the use of typed arrays is important. Compared to regular JavaScript arrays, typed arrays provide nearly double the performance in all browsers (*SQ2*). We have also shown that the *asm.js* subset can give additional performance improvements over just using typed arrays. Our experiments showed that the *asm.js* versions run 15% to 30% faster than hand-written JavaScript (using typed arrays), depending on the browser (*SQ3*). These results should help web developers of numerical applications and compiler writers for

numerical languages that target JavaScript in choosing languages features to use and target.

We have also shown that there are significant gains from parallel code for a majority of benchmarks. The fastest three benchmarks were around 5.85 times faster than JavaScript on GPUs (*PQ1*). We found that the WebCL code speedup over JavaScript had significant losses when compared to the OpenCL speedup over C. We attribute this to technical bottlenecks in the current implementation. Therefore, we find that WebCL gains are incongruent with OpenCL gains but the issue is not inherent to the WebCL specification (*PQ2*). These results provide a glimpse of the upcoming performance of general-purpose parallel computing using WebCL, which is still at the prototype phase and currently only supported by the Firefox browser. They should help its current proponents to focus their efforts on minimizing the performance overhead of providing additional security guarantees compared to OpenCL. Our benchmark suite will help to evaluate the performance evolution of the technology.

In addition to extending the suite with more benchmarks, one interesting possibility of extension for our work is to port the benchmarks to other numerical languages, such as MATLAB and R, to evaluate the performance of the generated JavaScript code produced by JavaScript compilers for these languages. Another interesting case is to analyze the sequential code to understand if different algorithmic variations could offer enhanced performance.

## Acknowledgments

We would like to thank Ismail Badawi for his help with our experimental infrastructure and Rahul Garg for providing us with his expertise on OpenCL and GPUs. This work was supported, in part, by NSERC and FQRNT.

## References

- [1] asm.js. URL <http://asmjs.org/spec/latest/>.
- [2] Cascading style sheets (CSS) snapshot 2010. URL <http://www.w3.org/TR/css-2010/>.
- [3] dart2js. URL <https://www.dartlang.org/tools/dart2js>.
- [4] Emscripten. URL <http://emscripten.org>.
- [5] Google web toolkit. URL <http://www.gwtproject.org/>.
- [6] A vocabulary and associated APIs for HTML and XHTML. URL <http://www.w3.org/TR/html5/>.
- [7] Standard ECMA-262 ECMAScript language specification edition 5.1 (june 2011). URL <http://www.ecma-international.org/publications/standards/Ecma-262.htm>.
- [8] Kraken Benchmark Suite. URL <http://krakenbenchmark.mozilla.org/>.
- [9] Octane Benchmark Suite. URL <https://developers.google.com/octane/>.
- [10] Opal: Ruby to JavaScript compiler. URL <http://opalrb.org/>.
- [11] OpenCL. URL <https://www.khronos.org/opencv/>.
- [12] Pyjs. URL <http://pyjs.org/>.
- [13] Profile your web application with v8's internal profiler. URL [https://developers.google.com/v8/profiler/\\_example](https://developers.google.com/v8/profiler/_example).
- [14] WebGL specifications. URL <http://www.khronos.org/registry/webgl/specs/latest/>.
- [15] WebCL specification. May 2014. URL <http://www.khronos.org/registry/webcl/specs/latest/1.0/>.
- [16] K. Asanovic, R. Bodik, B. C. Catanzaro, J. J. Gebis, P. Husbands, et al. The landscape of parallel computing research: A view from berkeley. Technical Report UCB/EECS-2006-183, EECS Department, University of California, Berkeley.
- [17] R. Auler, E. Borin, P. de Halleux, M. Moskal, and N. Tillmann. Addressing JavaScript JIT Engines Performance Quirks: A Crowd-sourced Adaptive Compiler. In *CC*.(2014).
- [18] L. Bak. Google chrome's need for speed. URL [http://blog.chromium.org/2008/09/google-chromes-need-for-speed\\_02.html](http://blog.chromium.org/2008/09/google-chromes-need-for-speed_02.html).
- [19] J. Bezanson, S. Karpinski, V. B. Shah, and A. Edelman. Julia: A Fast Dynamic Language for Technical Computing. *CoRR*, abs/1209.5145, 2012.
- [20] M. Bodin, A. Chargueraud, D. Filaretti, P. Gardner, et al. A Trusted Mechanised JavaScript Specification. In *POPL*.(2014).
- [21] A. Casey, J. Li, J. Doherty, M. C. Boisvert, T. Aslam, et al. McLab: An Extensible Compiler Toolkit for MATLAB and Related Languages. In *C3S2E*.(2010).
- [22] S. Che, M. Boyer, J. Meng, D. Tarjan, et al. Rodinia: A benchmark suite for heterogeneous computing. In *IISWC*.(2009).
- [23] J. Coglan. Sylvester: Vector and Matrix math for JavaScript.
- [24] P. Colella. Defining software requirements for scientific computing, 2004.
- [25] J. de Jong. math.js. URL <http://mathjs.org/>.
- [26] M. Elsman. SMLtoJs: Hosting a standard ML compiler in a web browser. In *PLASTIC*.(2011).
- [27] W. C. Feng, H. Lin, T. Scogland, and J. Zhang. OpenCL and the 13 Dwarfs: A Work in Progress. In *ICPE*.(2012).
- [28] J. Fifield and D. Grunwald. A Methodology for Fine-Grained Parallelism in JavaScript Applications. In *Languages and Compilers for Parallel Computing*, volume 7146 of *LNCS*, pages 16–30. 2013.
- [29] A. Gal, B. Eich, M. Shaver, D. Anderson, et al. Trace-based Just-in-time Type Specialization for Dynamic Languages. In *PLDI*.(2009).
- [30] A. Grosskurth and M. W. Godfrey. A Reference Architecture for Web Browsers. In *ICSM*.(2005).
- [31] S. Herhut, R. L. Hudson, T. Shpeisman, and J. Sreeram. River Trail: A Path to Parallelism in JavaScript. In *OOPSLA*.(2013).
- [32] R. Krahn, J. Lincke, and R. Hirschfeld. Efficient Layer Activation in Context JS. In *C5*.(2012).
- [33] J. Larres, A. Potanin, and Y. Hirose. A Study of Performance Variations in the Mozilla Firefox Web Browser. In *ACSC*.(2013).
- [34] S. Loisel. Numeric Javascript. URL <http://numericjs.com/>.
- [35] M. Mehrara, P.-C. Hsu, M. Samadi, and S. Mahlke. Dynamic parallelization of JavaScript applications using an ultra-lightweight speculation mechanism. In *HPCA*.(2011).
- [36] L. A. Meyerovich, M. E. Torok, E. Atkinson, and R. Bodik. Superconductor: A Language for Big Data Visualization, 2013.
- [37] F. Morandat, B. Hill, L. Osvald, and J. Vitek. Evaluating the Design of the R Language: Objects and Functions for Data Analysis. In *ECOOP*.(2012).
- [38] F. Pizlo. JetStream Benchmark Suite. URL <https://www.webkit.org/blog/3418/introducing-the-jetstream-benchmark-suite/>.
- [39] C. Reis and S. D. Gribble. Isolating Web Programs in Modern Browser Architectures. In *EuroSys*.(2009).
- [40] G. Richards. JSBench Benchmark Suite. URL <http://jsbench.cs.purdue.edu/>.
- [41] G. Richards, A. Gal, B. Eich, and J. Vitek. Automated Construction of JavaScript Benchmarks. In *OOPSLA*.(2011), .
- [42] G. Richards, S. Lebresne, B. Burg, and J. Vitek. An analysis of the dynamic behavior of javascript programs. In *PLDI*.(2010), .
- [43] C. Severance. JavaScript: Designing a language in 10 days. *Computer*, 45(2):7–8, Feb. 2012. ISSN 0018-9162. .
- [44] D. Yoo and S. Krishnamurthi. Whalesong: Running racket in the browser. In *DLS*.(2013).
- [45] A. Zakai. Emscripten: An LLVM-to-JavaScript compiler. In *OOPSLA*.(2011).