

Emscripten: An LLVM-to-JavaScript Compiler

Alon Zakai
Mozilla
azakai@mozilla.com

Abstract

We present Emscripten, an LLVM-to-JavaScript compiler. Emscripten compiles LLVM assembly code into standard JavaScript, which opens up two avenues for running code written in languages other than JavaScript on the web: (1) Compile code written in a language directly into LLVM bitcode, and then compile that into JavaScript using Emscripten, or (2) Compile a language's entire runtime into LLVM and then JavaScript, as in the previous approach, and then use the compiled runtime to run code written in that language. For example, the former approach can work for C and C++, while the latter can work for Python; all three examples open up new opportunities for running code on the web.

Emscripten itself is written in JavaScript and is available under the MIT license (a permissive open source license), at <http://www.emscripten.org>. As a compiler from LLVM to JavaScript, the challenges in designing Emscripten are somewhat the reverse of the norm – one must go from a low-level assembly into a high-level language, and recreate parts of the original high-level structure of the code that were lost in the compilation to low-level LLVM. We detail the algorithms used in Emscripten to deal with those challenges.

©2011 Alon Zakai. License: Creative Commons Attribution-ShareAlike (CC BY-SA), <http://creativecommons.org/licenses/by-sa/3.0/>

1 Introduction

Since the mid 1990's, JavaScript has been present in most web browsers (sometimes with minor variations and under slightly different names, e.g., JScript in Internet Explorer), and today it is well-supported on essentially all web browsers, from desktop browsers like Internet Explorer, Firefox, Chrome and Safari, to mobile browsers on smartphones and tablets. Together with HTML and CSS, JavaScript forms the standards-based foundation of the web.

Running other programming languages on the web has been suggested many times, and browser plugins have allowed doing so, e.g., via the Java and Flash plugins. However, plugins must be manually installed and do not integrate in a perfect way with the outside HTML. Perhaps more problematic is that they cannot run at all on some platforms, for example, Java and Flash cannot run on iOS devices such as the iPhone and iPad. For those reasons, JavaScript remains the primary programming language of the web.

There are, however, reasonable motivations for running code from other programming languages on the web, for example, if one has a large amount of existing code already written in another language, or if one simply has a strong preference for another language and perhaps is more productive in it.

As a consequence, there have been work on tools to compile languages **into** JavaScript. Since JavaScript is present in essentially all web browsers, by compiling one's language of choice into JavaScript, one can still generate content that will run practically everywhere. Exam-

ples of this approach include the Google Web Toolkit [Google Web Toolkit], which compiles Java into JavaScript and Pyjamas [Pyjamas], which compiles Python into JavaScript; for a comprehensive list, see [Ashkenas]. Such tools usually only allow a subset of the original language to be compiled. For example, multi-threaded code (with shared memory) is not possible on the web, so compiling code of that sort is subject to not directly possible. There are also often limitations of the conversion process, for example, Pyjamas compiles Python to JavaScript in a nearly 1-to-1 manner, and as a consequence the underlying semantics are those of JavaScript, not Python so for example division of integers can yield unexpected results (it should yield an integer in Python 2.x, but in JavaScript and in Pyjamas a floating-point number can be generated).

In this paper we present another project along those lines: **Emscripten**, which compiles LLVM assembly code into JavaScript. LLVM (the Low Level Virtual Machine, [Low Level Virtual Machine]) is a compiler project primarily focused on C, C++ and Objective-C. It compiles those languages through a *frontend* (the main ones of which are Clang and LLVM-GCC) into the LLVM intermediary representation (which can be machine-readable bitcode, or human-readable assembly), and then passes it through a *backend* which generates actual machine code for a particular architecture. Emscripten plays the role of a backend which targets JavaScript.

By using Emscripten, potentially many languages can be run on the web, using one of the following methods:

- Compile **code** in a language recognized by one of the existing LLVM frontends into LLVM, and then compile that into JavaScript using Emscripten. Frontends for various languages exist, including many of the most popular programming languages such as C and C++, and also various new and emerging languages (e.g., Rust).
- Compile the **runtime** used to parse and execute code in a particular language into

LLVM, then compile that into JavaScript using Emscripten. It is then possible to run code in that runtime on the web. This is a useful approach if a language's runtime is written in a language for which an LLVM frontend exists, but the language itself has no such frontend. For example, there is currently no frontend for Python, however it is possible to compile CPython – the standard implementation of Python, written in C – into JavaScript, and run Python code on that (see Section ??).

From a technical standpoint, one challenge in designing and implementing Emscripten is that it compiles a low-level language – LLVM assembly – into a high-level one – JavaScript. This is somewhat the reverse of the usual situation one is in when building a compiler, and leads to some unique difficulties. For example, to get good performance in JavaScript one must use natural JavaScript code flow structures, like loops and ifs, but those structures do not exist in LLVM assembly (instead, what is present there is essentially ‘flat’ code with *goto* commands). Emscripten must therefore reconstruct a high-level representation from the low-level data it receives.

In theory that issue could have been avoided by compiling a higher-level language into JavaScript. For example, if compiling Java into JavaScript (as the Google Web Toolkit does), then one can benefit from the fact that Java's loops, ifs and so forth generally have a very direct parallel in JavaScript. (Of course the downside in that approach is it yields a compiler only for Java.) Compiling LLVM into JavaScript is less straightforward, but we will see later that it is possible to reconstruct a substantial part of the original high-level structure of the original code.

Another challenge in Emscripten is to achieve good performance. LLVM assembly is an abstraction of how modern CPUs are programmed for, and its basic operations are not all directly possible in JavaScript. For example, if in LLVM we are to add two unsigned 8-bit numbers x and y , with overflowing (e.g., 255 plus 1 should give 0), then there is no single operation in JavaScript which can do this - we cannot

just write $x + y$, as that would use the normal JavaScript semantics. It is possible to emulate a CPU in JavaScript, however doing so is very slow. Emscripten’s approach is to allow such emulation, but try to use it as little as possible, and to provide tools that help one find out which parts of the compiled code will need full emulation. We will later see that in practice it is often possible to not use full emulation, and thereby achieve far better performance than strict emulation would achieve.

We conclude this introduction with a list of this paper’s main contributions:

- We describe Emscripten itself, during which we detail its approach in compiling LLVM into JavaScript.
- We give details of Emscripten’s ‘Relooper’ algorithm, which generates high-level loop structures from low-level branching data.

In addition, the following are the main contributions of Emscripten itself, that to our knowledge were not previously possible:

- It allows compiling a very large subset of C and C++ code into JavaScript, which can then be run on the web.
- By compiling their runtimes, it allows running languages such as Python on the web.

XXX The remainder of this paper is structured as follows. In Section 2 we describe, from a high level, the approach taken to compiling LLVM assembly into JavaScript. In Section 3 we describe the workings of Emscripten on a lower, more concrete level. In Section ?? we describe Emscripten’s approach to achieving good compatibility with existing real-world code, while still achieving good performance. In Section 4 we give an overview of some uses of Emscripten. In Section 5 we summarize and give directions for future work on Emscripten and uses of it.

2 Compilation Approach

Let us begin by considering what the challenge is, when we want to compile LLVM assembly into

JavaScript. Assume we are given the following simple example of a C program, which we want to compile into JavaScript:

```
#include <stdio.h>
int main()
{
    int sum = 0;
    for (int i = 1; i < 100; i++)
        sum += i;
    printf("1+...+100=%d\n", sum);
    return 0;
}
```

This program calculates the sum of the integers from 1 to 100. When compiled by Clang, the generated LLVM assembly code includes the following:

```
@.str = private constant [14 x i8]
        c"1+...+100=%d\0A\00"

define i32 @main() {
    %1 = alloca i32, align 4
    %sum = alloca i32, align 4
    %i = alloca i32, align 4
    store i32 0, i32* %1
    store i32 0, i32* %sum, align 4
    store i32 1, i32* %i, align 4
    br label %2

; <label>:2
    %3 = load i32* %i, align 4
    %4 = icmp slt i32 %3, 100
    br i1 %4, label %5, label %12

; <label>:5
    %6 = load i32* %i, align 4
    %7 = load i32* %sum, align 4
    %8 = add nsw i32 %7, %6
    store i32 %8, i32* %sum, align 4
    br label %9

; <label>:9
    %10 = load i32* %i, align 4
    %11 = add nsw i32 %10, 1
    store i32 %11, i32* %i, align 4
    br label %2
```

```

; <label>:12
%13 = load i32* %sum, align 4
%14 = call i32 @i8*, ...)*
    @printf(i8* getelementptr inbounds
      ([14 x i8]* @.str, i32 0, i32 0),
      i32 %13)
ret i32 0
}

```

At first glance, this may look more difficult to translate into JavaScript than the original C++. However, compiling C++ in general would require writing code to handle preprocessing, classes, templates, and all the idiosyncrasies and complexities of C++. LLVM assembly, while more verbose in this example, is lower-level and simpler to work on. Compiling it also has the benefit we mentioned earlier, which is one of the main goals of Emscripten, that it allows many languages can be compiled into LLVM and not just C++.

A detailed overview of LLVM assembly is beyond our scope here (see <http://llvm.org/docs/LangRef.html>). Briefly, though, the example assembly above can be seen to define a function `main()`, then allocate some values on the stack (`alloca`), then load and store various values (load and store). We do not have the high-level code structure as we had in C++ (with a loop), instead we have labelled code fragments, called LLVM basic blocks, and code flow moves from one to another by branch (`br`) instructions. (Label 2 is the condition check in the loop; label 5 is the body, label 9 is the increment, and label 12 is the final part of the function, outside of the loop). Conditional branches can depend on calculations, for example the results of comparing two values (`icmp`). Other numerical operations include addition (`add`). Finally, `printf` is called (`call`). The challenge, then, is to convert this and things like it into JavaScript.

In general, Emscripten's main approach is to translate each line of LLVM assembly into JavaScript, 1-to-1, into 'normal' JavaScript as much as possible. So, for example, an `add` operation becomes a normal JavaScript addition, a function call becomes a JavaScript function call, etc. This 1 to 1 translation generates JavaScript

that resembles assembly code, for example, the LLVM assembly shown before for `main()` would be compiled into the following:

```

function _main() {
  var __stackBase__ = STACKTOP;
  STACKTOP += 12;
  var __label__ = -1;
  while(1) switch(__label__) {
    case -1:
      var $1 = __stackBase__;
      var $sum = __stackBase__+4;
      var $i = __stackBase__+8;
      HEAP[$1] = 0;
      HEAP[$sum] = 0;
      HEAP[$i] = 0;
      __label__ = 0; break;
    case 0:
      var $3 = HEAP[$i];
      var $4 = $3 < 100;
      if ($4) { __label__ = 1; break; }
      else { __label__ = 2; break; }
    case 1:
      var $6 = HEAP[$i];
      var $7 = HEAP[$sum];
      var $8 = $7 + $6;
      HEAP[$sum] = $8;
      __label__ = 3; break;
    case 3:
      var $10 = HEAP[$i];
      var $11 = $10 + 1;
      HEAP[$i] = $11;
      __label__ = 0; break;
    case 2:
      var $13 = HEAP[$sum];
      var $14 = _printf(__str, $13);
      STACKTOP = __stackBase__;
      return 0;
  }
}

```

Some things to take notice of:

- A switch-in-a-loop construction is used in order to let the flow of execution move between basic blocks of code in an arbitrary manner: We set `__label__` to the (numerical representation of the) label of the fragment we want to reach, and do a break, which

leads to the proper fragment being reached. Inside each fragment, every line of code corresponds to a line of LLVM assembly, generally in a very straightforward manner.

- Memory is implemented by *HEAP*, a JavaScript array. Reading from memory is a read from that array, and writing to memory is a write. *STACKTOP* is the current position of the stack. (Note that we allocate 4 memory locations for 32-bit integers on the stack, but only write to 1 of them. See Section 2.2 for why.)
- LLVM assembly functions become JavaScript functions, and function calls are normal JavaScript function calls. In general, we attempt to generate as ‘normal’ JavaScript as possible.
- We implemented the LLVM *add* operation using simple addition in JavaScript. As mentioned earlier, the semantics of that code are not entirely identical to those of the original LLVM assembly code (in this case, overflows will have very different effects). We will explain Emscripten’s approach to that problem in Section 2.2.1.

2.1 Performance

In this section we will deal with several topics regarding Emscripten’s approach to generating high-performance JavaScript code.

2.2 Load-Store Consistency (LSC)

We saw before that Emscripten’s memory usage allocates the usual number of bytes on the stack for variables (4 bytes for a 32-bit integer, etc.). However, we only wrote values into the first location, which appeared odd. We will now see the reason for that.

To get there, we must first step back, and note that Emscripten does not aim to achieve perfect compatibility with all possible LLVM assembly (and correspondingly, with all possible C or C++ code, etc.); instead, Emscripten targets a subset of LLVM assembly code, which is portable and

does not make crucial assumptions about the underlying CPU architecture on which the code is meant to run. That subset is meant to encompass the vast majority of real-world code that would be compiled into LLVM, while also being compilable into very performant JavaScript.

More specifically, Emscripten assumes that the LLVM assembly code it is compiling has **Load-Store Consistency** (LSC), which is the requirement that loads from and stores to a specific memory address will use the same type. Normal C and C++ code generally does so: If x is a variable containing a 32-bit floating point number, then both loads and stores of x will be of 32-bit floating point values, and not 16-bit unsigned integers or anything else.

To see why this is important for performance, consider the following C code fragment, which does *not* have LSC:

```
int x = 12345;
printf("first byte: %d\n", *((char*)&x));
```

Assuming an architecture with more than 8 bits, this code will read the first byte of x . (This might, for example, be used to detect the endianness of the CPU.) To compile this into JavaScript in a way that will run properly, we must do more than a single operation for either the read or the write, for example we could do this:

```
var x_value = 12345;
var x_addr = stackAlloc(4);
HEAP[x_addr] = (x_value >> 0) & 255;
HEAP[x_addr+1] = (x_value >> 8) & 255;
HEAP[x_addr+2] = (x_value >> 16) & 255;
HEAP[x_addr+3] = (x_value >> 24) & 255;
[...]
printf("first byte: %d\n", HEAP[x_addr]);
```

Here we allocate space for the value of x on the stack, and store that address in x_addr . The stack itself is part of the ‘memory space’, which is the array *HEAP*. In order for the read on the final line to give the proper value, we must go to the effort of doing 4 store operations, each of the value of a particular byte. In other words, *HEAP* is an array of bytes, and for each store into memory, we must deconstruct the value into bytes.¹

¹Note that we can use JavaScript typed arrays with

Alternatively, we can store the value in a single operation, and deconstruct into bytes as we load. This will be faster in some cases and slower in others, but is still more overhead than we would like, generally speaking – for if the code **does** have LSC, then we can translate that code fragment into the far more optimal

```
var x_value = 12345;
var x_addr = stackAlloc(4);
HEAP[x_addr] = x_value;
[...]
```

```
printf("first byte: %d\n", HEAP[x_addr]);
```

(Note that even this can be optimized even more – we can store x in a normal JavaScript variable. For now though we are just clarifying why it is useful to assume we are compiling code that has LSC.)

In practice the vast majority of C and C++ code does have LSC. Exceptions do exist, however, for example:

- Code that detects CPU features like endianness, the behavior of floats, etc. In general such code can be disabled before running it through Emscripten, as it is not actually needed.
- *memset* and related functions typically work on values of one kind, regardless of the underlying values. For example, *memset* may write 64-bit values on a 64-bit CPU since that is usually faster than writing individual bytes. This tends to not be a problem, as with *memset* the most common case is setting to 0, and with *memcpy*, the values end up copied properly anyhow (with a proper implementation of *memcpy* in Emscripten’s generated code).
- Even LSC-obeying C or C++ code may turn into LLVM assembly that does not, after being optimized. For example, when storing

a shared memory buffer, which would work as expected, assuming (1) we are running in a JavaScript engine which supports typed arrays, and (2) we are running on a CPU with the same architecture as we expect. This is therefore dangerous as the generated code may run differently on different JavaScript engines and different CPUs. Emscripten currently has optional experimental support for typed arrays.

two 32-bit integers constants into adjoining locations in a structure, the optimizer may generate a single 64-bit store of an appropriate constant. In other words, optimization can generate nonportable code, which runs faster on the current CPU, but nowhere else. Emscripten currently assumes that optimizations of this form are not being used.

In practice it may be hard to know if code has LSC or not, and requiring a time-consuming code audit is obviously impractical. Emscripten therefore has a compilation option, `SAFE_HEAP`, which generates code that checks that LSC holds, and warns if it doesn’t. It also warns about other memory-related issues like reading from memory before a value was written (somewhat similarly to tools like Valgrind). When such problems are detected, possible solutions are to ignore the issue (if it has no actual consequences), or altering the source code.

Note that it is somewhat wasteful to allocate 4 memory locations for a 32-bit integer, and use only one of them. It is possible to change that behavior with the `QUANTUM_SIZE` parameter to Emscripten, however, the difficulty is that LLVM assembly has hardcoded values that depend on the usual memory sizes being used. For example, *memcpy* can be called with the normal size each the value would have, and not a single memory address each as Emscripten would prefer. We are looking into modifications to LLVM itself to remedy that.

2.2.1 Emulating Code Semantics

The semantics of LLVM assembly and JavaScript are not identical: The former is very close to that of a modern CPU, while the latter is a high-level dynamic language. Both are of course Turing-complete, so it is possible to precisely emulate each in the other, but doing so with good performance is more challenging. For example, if we want to convert

```
add i8 %1, %2
```

to JavaScript, then to be completely accurate we must emulate the exact same behavior, in particular, we must handle overflows properly, which

would not be the case if we just implement this as `%1+%2` in JavaScript. For example, with inputs of 255 and 1, the correct output is 0, but simple addition in JavaScript will give us 256. We can emulate the proper behavior by adding additional code, one way (not necessarily the most optimal) would be to check for overflows after each addition, and correct them as necessary. This however makes each operation take significantly more CPU time than the original code.

Emscripten's approach to this problem is to allow the generation of both accurate code, that is identical in behavior to LLVM assembly, and inaccurate code which is faster. In practice, most addition operations in LLVM do not overflow, and can be translated into `%1+%2` in JavaScript which is fast. Emscripten provides tools that make it straightforward to find which code does require the slower, more accurate code, and to generate that code in those locations. In practice, this is done as follows:

- Compile the code using Emscripten with special options that generate runtime checking. `CHECK_OVERFLOW`s adds runtime checks for integer overflows, `CHECK_SIGNS` checks for signing issues (the behavior of signed and unsigned integers can be different, and JavaScript does not natively support that difference), and `CHECK_ROUNDINGS` checks for rounding issues (in C and C++, the convention is to round towards 0, while in JavaScript there is no simple operation that does the same).
- Run the compiled code on a representative sample of inputs, and notice which lines are warned about by the runtime checks.
- Recompile the code, telling Emscripten to add corrections (using `CORRECT_SIGNS`, `CORRECT_OVERFLOW`s or `CORRECT_ROUNDINGS`) only on the specific lines that actually need it.

This method is not guaranteed to work, as if we do not run on a truly representative sample of possible inputs, we may not compile with

all necessary corrections. It is of course possible to compile with all corrections applied to all the code, to make sure things will work properly (this is the default compilation setting), however, in practice the procedure above appears to work quite well, and can result in code that runs very significantly faster.

2.2.2 Code Optimizations

When comparing the example program from page 4, the generated code was fairly complicated and cumbersome, and unsurprisingly it performs quite poorly. There are two main reasons for that: First, that the code is simply unoptimized – there are many variables declared when fewer could suffice, for example, and second, that the code does not use ‘normal’ JavaScript, which JavaScript engines are optimized for – it stores all variables in an array (not normal JavaScript variables), and it controls the flow of execution using a switch-in-a-loop, not normal JavaScript loops and ifs.

Emscripten's approach to generating fast-performing code is as follows. Emscripten doesn't do any optimizations that can be done by other tools: LLVM can be used to perform optimizations before Emscripten, and the Closure Compiler [Closure Compiler] can perform optimizations on the generated JavaScript afterwards. Those tools will perform standard useful optimizations like removing unneeded variables, dead code, function inlining, etc. That leaves two major optimizations that are left for Emscripten to perform:

- **Variable nativization:** Convert variables that are on the stack – which is implemented using addresses in the *HEAP* array as mentioned earlier – into native JavaScript variables (that is to say, *var x*; and so forth). In general, a variable will be nativized unless it is used outside that function, e.g., if its address is taken and stored somewhere or passed to another function. When optimizing, Emscripten tries to nativize as many variables as possible.
- **Relooping:** Recreate high-level loop and

if structures from the low-level labels and branches that appear in LLVM assembly. We detail the algorithm Emscripten uses for this purpose in Section 3.2.

When run with Emscripten’s optimizations, the code on page 4 looks like this:

```
function _main() {
  var __label__;
  var $1;
  var $sum;
  var $i;
  $1 = 0;
  $sum = 0;
  $i = 0;
  $2$2: while(1) {
    var $3 = $i;
    var $4 = $3 < 100;
    if (!($4)) { __label__ = 2; break $2$2; }
    var $6 = $i;
    var $7 = $sum;
    var $8 = $7 + $6;
    $sum = $8;
    var $10 = $i;
    var $11 = $10 + 1;
    $i = $11;
    __label__ = 0; continue $2$2;
  }
  var $13 = $sum;
  var $14 = _printf(__str, $13);
  return 0;
}
```

If in addition the Closure Compiler is run on that output, we get

```
function K() {
  var a, b;
  b = a = 0;
  a:for(;;) {
    if(!(b < 100)) {
      break a
    }
    a += b;
    b += 1;
  }
  _printf(J, a);
  return 0;
}
```

which is fairly close to the original C++ (the differences, of having the loop’s condition inside the loop instead of inside the for() expression at the top of the original loop, are not important to performance). Thus, it is possible to recreate the original high-level structure of the code that was compiled into LLVM assembly, despite that structure not being explicitly available to Emscripten.

2.3 Benchmarks

We will now take a look at some performance benchmarks:

benchmark	JS	gcc	ratio
fannkuch (9)	0.577	0.054	10.68
fasta (100,000)	0.640	0.055	11.63
primes	0.490	0.160	3.06
raytrace (5, 64)	1.078	0.033	32.66

The first column is the name of the benchmark, and in parentheses any parameters used in running it. The source code to all the benchmarks can be found at <https://github.com/kripken/emscripten/tree/master/tests> (each in a separate file with its name, except for ‘primes’, which is embedded inside runner.py in the function test_primes). The second column is of running the compiled code (using all Emscripten and LLVM optimizations, as well as the closure compiler) in the SpiderMonkey JavaScript engine (specifically, the JaegerMonkey branch, running with *-m -n -a*). The third column is the results when compiling the original code with *gcc -O3*. The last column is the ratio, that is, how much slower the JavaScript code is when compared to gcc.

3 Emscripten’s Architecture

In the previous section we saw a general overview of Emscripten’s approach to compiling LLVM assembly into JavaScript. We will now get into more detail into how Emscripten itself is implemented.

Emscripten is written in JavaScript. A main reason for that decision was to simplify sharing code between the compiler and the runtime,

and to enable various dynamic compilation techniques. Two simple benefits of this approach are that (1) the compiler can create JavaScript objects that represent structures from the original assembly code, and convert them to a string using `JSON.stringify()` in a trivial manner, and (2) the compiler can simplify numerical operations by simply `eval()`ing the code (so “1+2” would become “3”, etc.). In both examples, the development of Emscripten was made simpler by having the exact same environment during compilation as the executing code will have.

Emscripten’s compilation has three main phases:

- The **intertyper**, which converts from LLVM assembly into Emscripten’s internal representation.
- The **analyzer**, which inspects the internal representation and generates various useful information for the final phase, including type and variable information, stack usage analysis, optional data for optimizations (variable nativization and relooping), etc.
- The **jsifier**, which does the final conversion of the internal representation plus additional analyzed data into JavaScript.

3.1 The Runtime Environment

Code generated from Emscripten is meant to run in a JavaScript engine, typically in a web browser. This has implications for the kind of runtime environment we can generate for it, for example, there is no direct access to the local filesystem.

Emscripten comes with a partial implementation of a C library, mostly written from scratch in JavaScript, which parts compiled from an existing C library (newlib [newlib]). Some aspects of the runtime environment, as implemented in that C library, are:

- Files to be read must be ‘preloaded’ in JavaScript. They can then be accessed using the usual C library methods (`fopen`, `fread`, etc.). Files that are written are cached, and

can be read by JavaScript later. While limiting, this approach can often be sufficient for many purposes.

- Emscripten allows writing pixel data to an HTML5 canvas element, using a subset of the SDL API. That is, one can write an application in C or C++ using SDL, and that same application can be compiled normally and run locally, or compiled using Emscripten and run on the web. See, for example, Emscripten’s raytracing demo at <http://syntensity.com/static/raytrace.html>.
- `sbrk` is implemented using the **HEAP** array which was mentioned previously. This allows a normal `malloc` implementation written in C to be compiled to JavaScript.

3.2 The Relooper: Recreating high-level loop structures

The Relooper is among the most complicated components in Emscripten. It receives a ‘soup of blocks’, which is a set of labeled fragments of code, each ending with a branch operation, and the goal is to generate normal high-level JavaScript code flow structures such as loops and ifs. As we saw before, generating such code structures is essential to generating good-performing code.

For example, the LLVM assembly on page 3 has the following block structure:

```

      /-----\
      |         |
      v         |
ENTRY --> 2 --> 5 --> 9
      |
      v
      12

```

In this simple example, it is fairly straightforward to see that a natural way to implement it using normal loop structures is

```

ENTRY
while (true) do
    2

```

```

    if (condition) break
5
9
12

```

In general though, this is not always easy or even practical – there may not be a straightforward high-level loop structure corresponding to the low-level one, if for example the original C code relied heavily on *goto* instructions. In practice, however, almost all real-world C and C++ code tends to be amenable to loop recreation.

Emscripten’s Relooper takes as input a ‘soup of LLVM blocks’ as described above, and generates a structured set of Emscripten code blocks, which are each a set of LLVM blocks with some logical structure. For simplicity we call LLVM blocks ‘labels’ and Emscripten blocks ‘blocks’ in the following.

There are three types of Emscripten blocks:

- **Simple block:** A block with one internal label, and a Next block, which the internal label branches to. The block is later translated simply into the code for that label, and the Next block appears right after it.
- **Loop:** A block that represents a basic loop, comprised of two internal sub-blocks:
 - **Inner:** A block that will appear inside the loop, i.e., when execution reaches the end of that block, flow will return to the beginning. Typically a loop will contain a conditional *break* defining where it is exited. When we exit, we reach the Next block, below.
 - **Next:** A block that will appear just outside the loop, in other words, that will be reached when the loop is exited.
- **Multiple:** A block that represents an divergence into several possible branches, that eventually rejoin. A Multiple block can implement an ‘if’, an ‘if-else’, a ‘switch’, etc. It is comprised of:
 - **Handled blocks:** A set of blocks to which execution can enter. When we reach the multiple block, we check

which of them should execute, and go there. When execution of that block is complete, or if none of the handled blocks was selected for execution, we proceed to the Next block, below.

- **Next:** A block that will appear just after the Handled blocks, in other words, that will be reached after code flow exits the Handled blocks.

Remember that we have a `__label__` variable that helps control the flow of execution: Whenever we enter a block with more than one entry, we set `__label__` before we branch into it, and we check its value when we enter that block. So, for example, when we create a Loop block, its Next block can have multiple entries – any label to which we branch out from the loop. By creating a Multiple block after the loop, we can enter the proper label when the loop is exited. (Having a `__label__` variable does add some overhead, but it greatly simplifies the problem that the Relooper needs to solve and allows us to only need three kinds of blocks as described above. Of course, it is possible to optimize away writes and reads to `__label__` in many or even most cases.)

We will use the term ‘entry’ to mean a label that can be reached immediately in a block. In other words, a block consists of labels l_1, \dots, l_n , and the entries are a subset of those labels, specifically the ones that execution can directly reach when we reach that block. With that definition, the Relooper algorithm can then be written as follows:

- **Receive a set of labels and which of them are entry points.** We wish to create a block comprised of all those labels.
- **Calculate, for each label, which other labels it *can* reach,** i.e., which labels we are able to reach if we start at the current label and follow one of the possible paths of execution.
- **If we have a single entry, and cannot return to it (by some other label later on branching to it) then create a Simple block,** with the entry as its internal la-

bel, and the Next block comprised of all the other labels. The entries for the Next block are the entries to which the internal label can branch.

- **If we can return to all of the entries, return a Loop block**, whose Inner block is comprised of all labels that can reach one of the entries, and whose Next block is comprised of all the others. The entry labels for the current block become entry labels for the Inner block (note that they must be in the Inner block by definition, as each one can reach itself). The Next block's entry labels are all the labels in the Next block that can be reached by the Inner block.
- **If we have more than one entry, try to create a Multiple block**: For each entry, find all the labels it reaches that cannot be reached by any other entry. If at least one entry has such labels, return a Multiple block, whose Handled blocks are blocks for those labels (and whose entries are those labels), and whose Next block is all the rest. Entries for the next block are entries that did not become part of the Handled blocks, and also labels that can be reached from the Handled blocks.
- **If we could not create a Multiple block, then create a Loop block as described above** (see proof below of why creating a Loop block is possible, i.e., why the labels contain a loop).

Note that we first create a Loop only if we must, then try to create a Multiple, then create a Loop if we have no other choice. We could have slightly simplified this in various ways, but the algorithm as presented above has given overall better results in practice, in terms of the 'niceness' of the shape of the generated code, both subjectively and at least in some simple benchmarks.

Additional details of the algorithm include 'fixing' branch instructions accordingly. For example, when we create a Loop block, then all branch instructions outside of the loop are converted

into *break* commands, and all branch instructions to the beginning of the loop are converted into *continue* commands. Those commands are then ignored when called recursively to generate the Inner block (that is, the *break* and *continue* commands are guaranteed, by the semantics of JavaScript, to get us to where we need to go – they do not need any further work for them to work properly).

Emscripten also does an additional pass after running the Relooper algorithm which has been described. The Relooper is guaranteed to produce valid output (see below). The second pass takes that valid output and optimizes it, by making minor changes such as removing *continue* commands that occur at the very end of loops (where they are not needed), etc. In other words, the first pass focuses on generating high-level code flow structures that are correct, while the second pass simplifies and optimizes that structure.

We now turn to an analysis of the Relooper algorithm. It is straightforward to see that the output of the algorithm, assuming it completes successfully – that is, that it finishes in finite time, and does not run into an error in the last part (where it is claimed that if we reach it we can return to at least one of the entry labels) – is correct in the sense of code execution being carried out as in the original data. We will now prove that the algorithm must in fact complete successfully.

First, note that if we successfully create a block, then we simplify the remaining problem, where the 'complexity' of the problem for our purposes here is the sum of labels plus the sum of branching operations:

- This is trivial for Simple blocks (since we now have a Next block which is strictly smaller).
- It is true for Loop blocks simply by removing branching operations (there must be a branching back to an entry, which becomes a *continue*).
- For Multiple blocks, if the Next block is non-empty then we have split into strictly

smaller blocks (in number of labels) than before. If the next block is empty, then since we built the Multiple block from a set of labels with more than one entry, then the Handled blocks are strictly smaller than the current one.

We have seen that whenever we successfully create a block, we simplify the remaining problem as defined above, which means that we must eventually halt successfully (since we strictly decrease a nonnegative integer). The remaining issue is whether we can reach a situation where we *cannot* successfully create a block, which is if we reach the final part of the relooper algorithm, but cannot create a Loop block there. For that to occur, we must not be able to return to any of the entries (or else we would create a Loop block). But if that is so, we can, at minimum, create a Multiple block with entries for all the current entries, since the entry labels themselves cannot be reached by the others as we have just assumed (when we ruled out creating a Loop block here), contradicting the assumption and concluding the proof.

(We have not, of course, proven that the shape of the blocks is optimal in any sense. However, even if it is possible to optimize them further, the Relooper already gives a very substantial speedup due to the move from the switch-in-a-loop construction to more natural JavaScript code flow structures.)

4 Example Uses

Emscripten has been run successfully on several real-world codebases, including the following:

- **Python.** It is possible to run variants of Python on the web in various ways, including pyjamas, IronPython on SilverLight and Jython in Java. However, all of these are slightly nonstandard in the the Python code they run, while the latter two also require plugins to be installed. With Emscripten, on the other hand, it is possible to compile CPython itself – the standard, reference implementation of Python – and run

that on the web, which allows running standard Python code. An online demo is available at <http://syntensity.com/static/python.html>.

- **Poppler and FreeType**
- **zlib**
- **Bullet**
- **Lua**

4.1 CPython

Other ways to run Python on web: pyjamas/pyjs, old pypy-js backend, ?? also IronPython on Silverlight and Jython in Java. Limitations etc.

CPython is the standard implementation of the Python programming language written in C. Compiling it in Emscripten is straightforward, except for needing to change some `#defines`, without which CPython creates platform-specific assembly code.

Compilation using llvm-gcc generates approximately 27.9MB of LLVM assembly. After running Emscripten and the Closure Compiler, the generated JavaScript code is approximately 2.76MB in size. For comparison, a native binary version of CPython is approximately 2.28MB in size, so the two differ by only 21%.

The **CORRECT_OVERFLOWS** option in Emscripten is necessary for proper operation of the generated code, as Python hash code relies on integer overflows to work normally.

The demo can be seen live at <http://www.syntensity.com/static/python.html>. Potential uses include ... etc.

4.2 Bullet Physics Engine

Mention other bullet-*js* manual port, via Java

4.3 Lua

Mention other lua-*js* compilers

5 Summary

We presented Emscripten, an LLVM-to-JavaScript compiler, which opens up numerous opportunities for running code written in languages other than JavaScript on the web, including some not previously possible. Emscripten can be used to, among other things, compile real-world C and C++ code and run that on the web. In addition, by compiling the runtimes of languages which are implemented in C and C++, we can run them on the web as well, for example Python and Lua.

One of the main tasks for future work in Emscripten is to broaden its standard library. Emscripten currently includes portions of libc and other standard C libraries, implemented in JavaScript. Portions of existing libc implementations written themselves in C can also be compiled into JavaScript using Emscripten, but in general the difficulty is in creating a suitable runtime environment on the web. For example, there is no filesystem accessible, nor normal system calls and so forth. Some of those features can be implemented in JavaScript, in particular new HTML features like the File API should help.

Another important task is to support multithreading. Emscripten currently assumes the code being compiled is single-threaded, since JavaScript does not have support for multithreading (Web Workers allow multiprocessing, but they do not have shared state, so implementing threads with them is not trivial). However, it would be possible to emulate multithreading in a single thread. One approach could be to not generate native JavaScript control flow structures, and instead to use a switch-in-a-loop for the entire program. Code can then be added in the loop to switch every so often between ‘threads’, while maintaining their state and so forth.

A third important task is to improve the speed of generated code. An optimal situation would be for code compiled by Emscripten to run at or near the speed of native code. In that respect it is worth comparing Emscripten to Portable Native Client (PNaCl), a project in development which aims to allow an LLVM-like format to be distributed and run securely on the web, with

speed comparable to native code.

Both Emscripten and PNaCl allow running compiled native code on the web, Emscripten by compiling that code into JavaScript, and PNaCl by compiling it into an LLVM-like format, which is then run in a special PNaCl runtime. The major differences are that Emscripten’s generated code can run on all web browsers, since it is standard JavaScript, while PNaCl’s generated code requires the PNaCl runtime to be installed; another major difference is that JavaScript engines do not yet run code at near-native speeds, while PNaCl does. In a broad summary, Emscripten’s approach allows the code to be run in more places, while PNaCl’s allows the code to run faster.

However, improvements in JavaScript engines may narrow the speed gap. In particular, for purposes of Emscripten we do not need to care about *all* JavaScript, but only the kind generated by Emscripten. Such code is **implicitly statically typed**, that is, types are not mixed, despite JavaScript in general allowing assigning, e.g., an integer to a variable and later a floating point value or even an object to that same variable. Implicitly statically typed code can be statically analyzed and converted into machine code that has no runtime type checks at all. While such static analysis can be time-consuming, there are practical ways for achieving similar results quickly, such as tracing and type inference, which would help on such code very significantly, and are already in use or being worked on in mainstream JavaScript engines (e.g., SpiderMonkey).

The limit of such an approach is to perform static analysis on an entire program compiled by Emscripten, generating highly-optimized machine code from that. As evidence of the potential in such an approach, the PyPy project can compile RPython – something very close to implicitly statically typed Python – into C, which can then be compiled and run at native speed. We may in the future see JavaScript engines perform such static compilation, when the code they run is implicitly statically typed, which would allow Emscripten’s generated code to run at native speeds as well. While not trivial, such an approach is possible, and if accomplished, would

mean that the combination of Emscripten and suitable JavaScript engines will let people write code in their languages of choice, and run them at native speeds on the web.

Finally, we conclude with another another avenue for optimization. Assume that we are compiling a C or C++ runtime of a language into JavaScript, and that that runtime uses JIT compilation to generate machine code. Typically code generators for JITs are written for the main CPU architectures, today x86, x86_64 and ARM. However, it would be possible for a JIT to generate JavaScript instead. Thus, the runtime would be compiled using Emscripten, and at runtime it would pass the JIT-generated JavaScript to *eval*. In this scenario, JavaScript is used as a low-level intermediate representation in the runtime, and the final conversion to machine code is left to the underlying JavaScript engine. This approach can potentially allow languages that greatly benefit from a JIT (such as Java, Lua, etc.) to be run on the web efficiently.

Thank you to the following people: Brian Crowder (multithreading), Robert O’Callahan (JIT into JS), add people that contributed

References

- [SOME REF et al.(2009)Someone, Another] A. B. Someone, and X. Y. Another. ...reference text...
- [Ashkenas] Jeremy Ashkenas’ list of languages that compile to JavaScript <https://github.com/jashkenas/coffee-script/wiki/List-of-languages-that-compile-to-JS>
- [Closure Compiler] Google, Inc.<http://code.google.com/closure/compiler/>
- [Google Web Toolkit] <http://code.google.com/webtoolkit/>
- [Low Level Virtual Machine] <http://llvm.org/>
- [newlib] Red Hat, Inc.<http://sourceware.org/newlib/>
- [Pyjamas] <http://pyjs.org/>