

Emscripten: An LLVM-to-JavaScript Compiler

Alon Zakai

Mozilla

azakai@mozilla.com

Abstract

JavaScript is the standard language of the web, supported on essentially all web browsers. Despite efforts to allow other languages to be run as well, none have come close to being universally available on all browsers, which severely limits their usefulness on the web. However, there are reasonable reasons why allowing other languages would be beneficial, including reusing existing code and allowing developers to use their languages of choice.

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 other languages on the web: (1) Compiling code written in another language directly into LLVM, if such a compiler is available, and then compiling that into JavaScript using Emscripten, or (2) Compiling a language's entire runtime (typically written in C or C++) into JavaScript using Emscripten, and using that to run code written in that language. Thus, Emscripten opens up many opportunities for running additional languages on the web, some of which were previously impossible.

Emscripten itself is written in JavaScript (to enable various dynamic compilation techniques), and is available under the MIT license (a permissive open source license), at <http://www.emscripten.org>. As an LLVM-to-JavaScript compiler, 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. We detail the algorithms used in Emscripten to deal with those challenges.

[Copyright notice will appear here once 'preprint' option is removed.]