Putting in All the Stops

Execution Control for JavaScript

Samuel Baxter
University of Massachusetts Amherst

Rachit Nigam
University of Massachusetts Amherst

Arjun Guha
University of Massachusetts Amherst

Joe Gibbs Politz
University of California San Diego

Shriram Krishnamurthi
Brown University

Abstract

Scores of compilers produce JavaScript, enabling programmers to use many languages on the Web, reuse existing code, and even use Web IDEs. Unfortunately, most compilers expose the browser's compromised execution model, so long programs freeze the browser tab, infinite loops crash IDEs, and so on. The few compilers that avoid these problems suffer poor performance and are difficult to engineer.

This paper presents Breakout, a source-to-source compiler that extends JavaScript with debugging abstractions and blocking operations, and easily integrates with existing compilers. We apply Breakout to ten programming languages and develop a Web IDE that supports stopping, single-stepping, breakpointing, and long-running computations. For nine languages, Breakout requires no or trivial compiler changes. For eight, our IDE is the first that provides these features. Two of our subject languages have compilers with similar features. Breakout's performance is competitive with these compilers and it makes them dramatically simpler.

Breakout's abstractions rely on first-class continuations, which it provides by compiling JavaScript to JavaScript. We also identify sub-languages of JavaScript that compilers implicitly use, and exploit these to improve performance. Finally, Breakout needs to repeatedly interrupt and resume program execution. We use a sampling-based technique to estimate program speed that outperforms other systems.

1 Programming On the Web

Over the last decade, Web browsers have evolved into a powerful and universally-supported software platform. Traditional desktop programs, such as word processors, spreadsheets, and photo editors now compete with Web-based alternatives. More recently, the Web has started to affect programming technology too. Scores of programming languages compile to run on the Web [24] and there are several Web IDEs in widespread use [5, 8, 14, 38, 46, 47, 56, 67, 75, 79]. This growing audience for Web IDEs, and correspondingly languages that run in the browser, includes professionals and students. Whereas sales of traditional computers are shrinking [2], ChromeBooks now account for 58% of new devices sold to U.S. schools [64].

Unfortunately, JavaScript and the Web platform lack abstractions necessary to build IDEs and serve as a complete compilation target for high-level languages. As a result, we observe compilers for high-level languages that compromise on semantics, and popular Web IDEs that compromise on basic debugging and usability features. Moreover, as we explain in §7, new technologies such as WebAssembly and Web Workers do *not* address most of the problems that we address. Indeed, our work may be viewed as presenting additional challenges to the creators of those technologies.

Limitations in Web IDEs The key problem facing a Web IDE is that JavaScript has a single-threaded execution environment.[1] An IDE that wants to provide a "stop" button to halt runaway computation faces a nontrivial problem, because the callback for that button gets queued for execution behind the running code—which is not going to terminate. Related to this, to make pages responsive, the browser threatens to interrupt computations that run longer than a few seconds. This means long-running computations need to be broken up into short events. In turn, because computations cannot run for very long, JavaScript engines in browsers tend to provide very shallow stacks, which is problematic for functional programs that rely on recursion.

These limitations are not at all hypothetical. Codecademy, which claims over 25 million users [37], has a Web IDE for JavaScript and for Python. In response to numerous message board requests, CodeAcademy has a help article that explicitly addresses infinite loops that freeze the browser [77]: they suggest refreshing the page, which loses browser state and recent changes. Other systems, such as CodeSchool, Khan Academy, and Python Tutor [19], address this problem by killing all programs that run for more than a few seconds. These problems also afflict research-driven programming languages, such as Elm [11] and Lively Kernel [46], which crash when given an infinite loop. Appendix A discusses all these systems in more detail.

[1] An IDE that runs code on a server has several weaknesses too: the logistics of scaling and paying for large computing infrastructure; the security concerns of running untrusted code; users having to trust their computation to a server; running off-line; using the browser's rich DOM environment; and more. This paper focuses on IDEs that run user code in the browser.

Under double-blind review. Please do not distribute.

Preliminary Solutions There are a handful of robust programming language implementations for the Web: GopherJS (Go) [18], Pyret [56], Skulpt (Python) [66], Doppio (JVM) [74], GambitJS (Scheme) [69], and Whalesong (Racket) [79]. They use sophisticated compilers and runtime systems to support some subset of long-running computations, shared-memory concurrency, blocking I/O, proper tail calls, debugging, and other features that are difficult to implement in the browser. However, these systems have several shortcomings.

First, these systems are difficult to build and maintain because of the complexity of, effectively, implementing expressive control operators for which JavaScript has no native support. For example, GopherJS has had several issues in its implementation of goroutines [10, 12, 13, 16, 42, 70]; Skulpt [66] has had bugs in its debugger [15, 59]; and Pyret has had issues in its Web IDE [23, 28, 29, 31–33, 49–54], and several of these issues remain unresolved. As another example of implementation difficulty, the team that built Pyret previously developed a Web IDE for Scheme [80], but could not reuse the compiler from the Scheme system, because the execution management techniques and the language's semantics were too tightly coupled.

Second, these systems force all programs to pay for all features. For example, Pyret forces all programs to pay the cost of instrumentation for the browser, even if they are running in production or at the command-line; GopherJS forces all programs to pay the cost of goroutines, even if they don't use concurrency; and Doppio forces all programs to pay the cost of threads, even if they are single-threaded. Third, these compilers have a single back-end—one that is presumably already complex enough—for all browsers, and hence do not maximize performance on any particular browser. Finally, it is very difficult for a compiler author to try a new approach, since small conceptual changes can require the entire compiler and runtime system to change. Furthermore, although these systems use implementation techniques that are interchangeable in principle, in practice they cannot share code to benefit from each others' performance improvements and bug fixes. What is called for is a clear separation of concerns.

Contributions

Our goal is to enhance JavaScript to make it a suitable target for Web IDEs and, more broadly, to run a variety of languages atop JavaScript without compromising their semantics and programming styles. Our system, Breakout, is a *compiler from JavaScript to JavaScript*. Given a naïve compiler from language *L* to JavaScript—call it *L*JS—we can compose it with Breakout. Breakout prepares the code for Web execution while leaving *L*JS mostly or entirely unchanged.

Breakout relies on four key ideas. The first is to reify continuations with a family of novel implementation strategies (§3). The second is to identify reasonable sub-languages of JavaScript—as targeted by compilers—to reduce overhead and hence improve performance (§4). The third is to dynamically determine the rate at which to yield control to the browser, improving performance without hurting responsiveness (§5). Finally, we study how these different techniques vary in performance across browsers, enabling browser-specific performance improvements (§6).

```
1 type Opts = {
2   cont: 'checked' | 'exceptional' | 'eager', // continuation representation
3   ctor: 'direct' | 'wrapped',                // constructor representation
4   timer: 'exact' | 'countdown' | 'approx',   // time estimator
5   yieldInterval: number,                     // yield interval
6   stacks: 'normal' | 'deep',                 //  deep stacks
7   implicits: true | false | '+',             // implicit conversions
8   args: true | false | 'mixed' | 'varargs',  // arity-mismatch behavior
9   getters: true | false,                     // support getters
10  eval: true | false                         // apply breakout to eval'd code
11 }
12
13 type AsyncRun = {
14   run: (onDone: () => void) => void,
15   pause: (onPause: () => void) => void,
16   resume: () => void
17 }
18
19 function breakout(source: string, opts: Opts): AsyncRun;
```

Figure 1. A portion of the `breakout` API.

Continuations and execution control features enable new capabilities for languages that compile to JavaScript. We show several: (1) Breakout supports long running computations by periodically yielding control to the browser; (2) Breakout provides abstractions that help compilers simulate blocking operations atop nonblocking APIs; (3) the abstractions of Breakout enable stepping debuggers via source maps; (4) Breakout allows simulating a much deeper stack than most browsers provide; and (5) Breakout can simulate tail calls on browsers that don't implement them natively.

We evaluate the effectiveness of Breakout in four ways:

1. We evaluate Breakout on ten compilers that produce JavaScript, nine of which require no changes, and quantify the cost of Breakout using 147 benchmarks (§6.1).
2. We use Breakout to build an IDE for nine languages. For eight of them, ours is the first Web IDE that supports long-running computation and graceful termination. For seven languages, our IDE also supports breakpoints and single-stepping (§5.2).
3. We show that our Breakout-based Python IDE is faster and more reliable than a widely-used alternative (§6.2).
4. We present a case study of Pyret, integrating Breakout into its compiler and IDE. We show that Breakout makes the compiler significantly simpler, helps fix outstanding bugs, has competitive performance, and presents new opportunities for optimization (§6.2).

2 Overview

Breakout, which is written in TypeScript, provides a function called `breakout` (line 19 in figure 1) that takes two arguments: (1) JavaScript code to run and (2) a set of options that affect

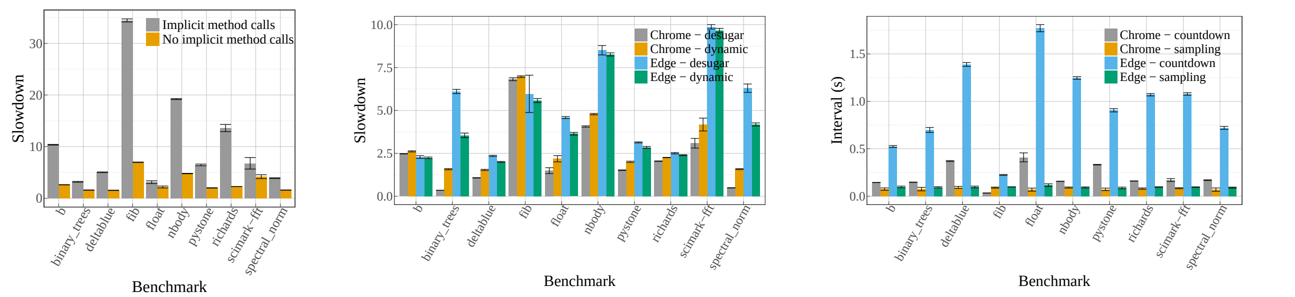

(a) Nonterminating arithmetic.

(b) Constructor encoding.

(c) Average time between yields.

Figure 2. Performance of Breakout relative to unmodified PyJS on a suite of 10 Python benchmarks run 10 times each. Each graph shows how an option setting in Breakout affects running time or latency.

the Breakout compiler and runtime system, which we elucidate in the rest of this section. The function produces an object with three methods:

1. The `run` method returns immediately and starts to evaluate the program. Breakout instruments every function and loop in the program to periodically save the continuation, yield control to the browser, and schedule resumption. In the absence of intervening events, the resumption event runs immediately (by applying the continuation). These periodic yields ensure that the browser tab remains responsive. When execution finally concludes, Breakout applies the callback passed to `run` to notify the caller.
2. The `pause` method interrupts the running program by setting a runtime flag that Breakout checks before applying a continuation in the resumption event. When the flag is set, Breakout calls the callback passed to `onPause` instead. Therefore, in an IDE, the event handler for a "stop" button can simply call `pause`—as if it is operating in a cooperative environment—and rely on Breakout to handle the rest.
3. Finally, `resume` resumes a paused program.

Breakout has additional methods to work with breakpoints, which we elide here. The rest of this section highlights a few features of Breakout using, as an example, the PyJS Python-to-JavaScript compiler's output as Breakout's input. This section presents only 10 benchmarks running on the Chrome and Edge Web browsers. §6 describes our experimental setup and an extensive evaluation.

Sub-languages of JavaScript There are some surprising ways to write nonterminating programs in JavaScript. For instance, an apparent arithmetic expression like `x + 1` or field access like `o.f` can fail to terminate if the program has an infinite loop in a `toString` method or a getter respectively. Breakout is conservative by default and can stop these nonterminating programs, but at a significant performance cost. However, most compilers generate programs in a "well-behaved" sub-language of JavaScript. As figure 2a shows, when we compile PyJS with conservative settings, the slowdown is much higher than when we specify that arithmetic expressions do not cause infinite loops, which is a safe assumption to make *of the code generated by PyJS*. §4 presents more sub-languages of this kind.

Browser-specific optimizations Different browsers optimize JavaScript differently, which gives Breakout the opportunity to implement browser-specific optimizations. For example, any function may serve as a constructor and there are two ways to capture a constructor's continuation frame: we can (1) *desugar* constructor calls to ordinary function calls (using `Object.create`) or (2) *dynamically* check if a function is a constructor. Figure 2b plots the slowdown incurred by both approaches on Chrome and Edge. Desugaring is better than the dynamic approach on Chrome ($p = 0.001$), but the other way around in Edge ($p = 0.042$). To our knowledge, none of the works we cite in this paper implement browser-specific optimizations. §3 discusses compilation issues in more detail.

Estimating elapsed time Breakout needs to periodically interrupt execution and yield control to the browser's event loop, which can be done in several ways. Skulpt [66] continuously checks the system time, but this is needlessly expensive. Pyret [56] counts basic blocks and interrupts after executing a fixed number of them. However, this approach results in high variability in page responsiveness between benchmarks and browsers. Figure 2c shows the average time between interrupts when counting up to 1 million using this mechanism. On Chrome, a few benchmarks yield every 30 milliseconds, which is too frequent and needlessly slows programs down. In contrast, on Edge, some benchmarks only yield every 1.5 seconds, which is slow enough for the browser to display a *"script is not responding"* warning. In Breakout, we sample the system time and dynamically estimate the rate at which statements are executed. Using this mechanism, Breakout takes a desired interrupt interval, δ, as a parameter and ensures that the program interrupts every δ milliseconds with high probability. Figure 2c shows that the average time between interrupts with $\delta = 100$ is very close to δ on both browsers.

In sum, we've identified the JavaScript sub-language that PyJS targets, applied browser-specific optimizations, and used an estimator of system time. This combination makes Breakout's Python faster and more reliable than a widely-used alternative (§6.2), while also being simpler to maintain. Breakout has several more features and options that we use to support a variety of programming languages that the rest of this paper describes in detail.

3 Continuations for JavaScript

This section presents Breakout's first-class continuations for a fragment of JavaScript that excludes some of its most egregious features (which we defer to §4). We have experimented with several ways to implement continuations and found that the best approach varies by browser. To exploit this, we present a parameterized family of approaches.

Language Extension We extend JavaScript with a unary operator, which we write as $\mathcal{C}$.[2] When applied to a function, `C(function(k) { body })`, the operator reifies its continuation to a value, binds it to *k*, and then evaluates *body* in an empty continuation. Within *body*, applying *k* aborts the current continuation and restores the saved continuation. For example, the following expression produces `0`:

```
10 + C(function (k) { return 0; })
```

Above, $\mathcal{C}$ evaluates the body in an empty continuation, thus produces `0`. In contrast, the next expression produces `11`:

```
10 + C(function (k) { return k(1) + 2; })
```

Above, $\mathcal{C}$ evaluates the body in an empty continuation and the body applies `k(1)`. This application discards the current continuation (which adds `2`) and restores the saved continuation (which adds `10`), thus produces `11`.

Breakout breaks up long running computations using a combination of $\mathcal{C}$ and browsers' timers. For example, the following function saves its continuation and uses `setTimeout` to schedule an event that restores the saved continuation:

```
function suspend() {
  C(function(k) { window.setTimeout(k, 0); }); }
```

Therefore, `suspend()` gives the browser a chance to process other queued events before resuming the computation. For example, the following infinite loop does not lock up the browser since each iteration occurs in a separate event:

```
while(true) { suspend(); }
```

Note that it is unnecessary to suspend computation on each iteration and we defer this issues to §5. The rest of this section shows how we compile $\mathcal{C}$ to ordinary JavaScript.

[2] $\mathcal{C}$ is equivalent to Sitaram and Felleisen's *control* operator [65].

```
var mode = 'normal';
var stack = [];

function P(f, g, x) {
  var t0, label, nextFrame;
  if (mode === 'restore') {
    const frame = stack.pop();
    [t0] = frame.locals;
    label = frame.label;
    nextFrame = stack[stack.length - 1];
  }
  function locals() { return [t0]; }
  function restore() { return P(f,g,x); }
  if (mode === 'normal' || (mode === 'restore' && label === 0)) {
    t1 = mode === 'normal' ? g(x) : nextFrame.restore();
    if (mode === 'capture') {
      stack.push({ label: 0, locals: locals(), restore: restore });
      return;
    }
  }
  return f(t0);
}
```

Figure 3. An instrumented function.

Strawman Solutions There are two obvious ways to implement $\mathcal{C}$ in modern JavaScript. One approach is to transform the program to continuation-passing style (with trampolining on browsers that don't implement proper tail calls). Alternatively, we could use generators to implement one-shot continuations [4]. We have investigated but abandoned both approaches for two reasons. First, on several benchmarks, we find that CPS and generators are 3x and 2x slower than the approach we present below. Second, both approaches change the type of instrumented functions, which makes it hard to support key language features such as constructors and prototype inheritance: CPS requires functions to take extra arguments, and generators turn all functions into generator objects.

3.1 Our Approach

We compile $\mathcal{C}$ to ordinary JavaScript in three steps: (1) A-normalize [17] programs, thus translating to a JavaScript subset where all applications are either in tail position or name their result; (2) box assignable variables that are captured by nested functions (discussed in §3.1.1); (3) instrument every function to operate in three modes: in *normal mode*, the program executes normally; in *capture mode* the program unwinds and reifies its stack; and in *restore mode* the program restores a saved stack.

We use the following function `P` as a running example:

```
function P(f, g, x) { return f(g(x)); }
```

Although `P` does not use $\mathcal{C}$ itself, any of the functions that it applies may use $\mathcal{C}$. Figure 3 shows an instrumented version of `P` and two global variables that determine the current execution mode (`mode`) and hold a reified stack (`stack`). The reified stack is only relevant in capture and restore mode. Note that in normal mode, the instrumented function is equivalent to the original.

Suppose `g(x)` applies $\mathcal{C}$, which switches execution to capture mode. To address this case, `P` checks to see if the program is in capture mode after `g(x)` returns (line 16). If so, `P` reifies its stack frame (line 17) and returns immediately. A reified stack frame contains (1) a copy of the local variables, (2) a label that identifies the current position within the function, and (3) a thunk called `restore` that re-enters the function.

Now, consider how `P` operates in restore mode. The function begins with a block of code that restores the saved local variables (lines 6—11). Next, on line 14, the function checks the saved label to see if `g(x)` should be applied. (the only label in this function is `0`.) Finally, instead of applying `g(x)` again, we apply the `restore()` function that `g(x)` had pushed onto the stack during capture mode. When `g(x)` returns normally, the last line of the function calculates `f(t0)`.

In general, to implement the three execution modes, we transform a function `f` as follows. (1) We define a nested thunk called `locals` that produces an array containing the values of `f`'s local variables. (2) We define a nested thunk called `restore` that applies `f` to its original arguments. (3) We give a unique label to every non-tail call in `f`. (4) After every non-tail call, we check if the program is in capture mode. If so, we push on object onto `stack` that contains the (a) label of the function call, (b) an array of local variable values (produced by `locals`), and (c) a reference to the `restore` function. This object is the continuation frame for `f`. We then immediately return and allow `f`'s caller to do the same. (5) At the top of `f`, we add a block of code to check if the program is in restore mode, which indicates that a continuation is being applied. If so, we use the reified stack frame to restore local variables and the label saved on the stack. (6) We instrument `f` such that in restore mode, the function effectively jumps to the call site that captured the continuation. Here, we invoke the `restore` method of the next stack frame. When the continuation is fully restored, execution switches back to normal mode. Finally, the top-level of the program needs to be wrapped by a thunk and executed within a driver loop that manages execution in two ways. (1) The expression `C(f)` switches execution to capture mode and reifies the stack. Therefore, the driver loop has to apply `f` to the reified stack. (2) When a continuation is applied, the program throws a special exception to unwind the current stack. The driver loop has to start restoring the saved stack by invoking the `restore` method of the bottommost frame.

3.1.1 Assignable Variables

A problem that arises with this approach involves assignable variables that are captured by nested functions. To restore a function `f`, we have to reapply `f`, which allocates new local variables, and thus we restore local variables' values (e.g., line 8 in figure 3). However, suppose `f` contains a nested function `g` that closes over a variable `x` that is local to `f`. If `x` is an assignable variable, we must ensure that after restoration `g` closes over the new copy of `x` too. We resolve this problem by boxing assignable variables that are captured by nested functions. This is the solution that *scheme2js* uses [36].

3.1.2 Proper tail calls

Our approach preserves proper tail calls. Notice that the application of `f` is in tail position and is not instrumented (line 21 of figure 3). Consider what happens if `f` uses `C`: `f` would first reify its own stack frame and then return immediately to `P`'s caller (instead of returning to `P`). In restore mode, `P`'s caller will jump to the label that called `P` and then call `nextFrame.restore()`. Since `P` did not save its own stack frame, this method call would jump into `f`, thus preserving proper tail calls. On browsers that do not support proper tail calls, Breakout uses trampolines.

3.1.3 Exception handlers and finalizers

Suppose a program captures a continuation within a **`catch`** clause. In restore mode, the only way to re-enter the **`catch`** clause is for the associated **`try`** block to throw an exception. Throwing an arbitrary exception will lose the original exception value, so we need to throw the same exception that was caught before. Therefore, we instrument each **`catch`** block to create a new local variable that stores the caught exception value and instrument each **`try`** block to throw the saved exception in restore mode.

Now, suppose a program captures a continuation within a **`finally`** clause. In restore mode, the only way to re-enter the **`finally`** block is to **`return`** within its associated **`try`** block. However, the returned value is not available within the **`finally`** block. Therefore, we instrument **`try`** blocks with finalizers to save their returned value in a local variable. In restore mode, the **`try`** block returns the saved value to transfer control to the **`finally`** block, which preserves the returned value.

3.2 Variations of Our Approach

The above approach is parameterizable in two major ways: (1) ways to capture stack frames and (2) ways to support continuations within constructors.

Capturing Stack Frames Breakout captures stack frames in three different ways. The previous section presented what we call *checked-return continuations*: every function application is instrumented to check whether the program is in capture mode. An alternative that involves fewer checks is an *exceptional continuations* approach [36]. Here, `C` throws a special exception and every function application is guarded by an exception handler that catches the special exception, reifies the stack frame, and re-throws it. Although this approach involves fewer conditionals than the checked return approach, exception handlers come at a cost. Both checked-return and exceptional continuations reify the stack lazily. An alternative approach, which we call *eager continuations*, maintains a shadow stack at all times [69]. This makes capturing continuations trivial and very fast. However, maintaining the shadow stack slows down normal execution. Breakout support all three approaches and all compose with the several other configuration options that Breakout provides.

Compiler	Impl	Args	Getters	Eval	(#) Benchmarks
PyJS	✗	**M**	✗	✗	(16) [9, 60]
ScalaJS	+	✗	✗	✗	(18) [9]
scheme2js	✗	**V**	✗	✗	(13) [30]
ClojureScript	+	**M**	✗	✗	(8) [9]
dart2js	+	✗	**T**	**T**	(15) [9, 71]
Emscripten	✗	**V**	✗	✗	(13) [9, 25]
BuckleScript	✗	✗	✗	✗	(15) [9, 45]
JSweet	+	**M**	✗	✗	(9) [9, 58]
JavaScript	✓	✓	✓	✓	(19) [9, 27]
Pyret	✗	✗	✗	**T**	(21) [56]

Figure 4. Compilers, their sub-language of JavaScript, and benchmark sources. A ✓ or ✗ indicates that a JavaScript feature is used in full or completely unused. The other symbols indicate restricted variants of the feature (discussed in §4).

Constructors JavaScript allows almost any function to be used as a constructor and a single function may assume both roles. Breakout allows constructors to capture their continuation using two different approaches. The simple approach is to desugar `new`-expressions to ordinary function calls (using `Object.create`), which effectively eliminates constructors. (Constructors for builtin types, such as `new Date()`, cannot be eliminated in this way.) Unfortunately, this desugaring can perform poorly on some browsers.

It is more challenging to preserve `new`-expressions. Consider the following constructor, which create two fields and calls a function `f`:

```
function N(x, f) {this.x = x; this.y = f(); return 0;}
```

We need to instrument `N` to address the case where `f` captures its continuation. The problem is that `new N(x, f)` allocates a new object (bound to `this`) every time it is called. Therefore, in restore mode, we have to apply `N` to the originally allocated object so that we do not lose the write to `this.x`. We address this problem in the `restore` function, by calling `N` as a function (not a constructor) and passing the original value of `this`:

```
N.call(this, x, f)
```

This presents another problem—when `N` is invoked as a constructor (as it originally was) it returns `this`, but when `N` is applied as an ordinary function (during restoration), it returns `0`. Therefore, we also need to track if the function was originally invoked as a constructor, so that we can return the right value in `restore` mode.

4 Sub-Languages of JavaScript

Breakout's continuations work with arbitrary JavaScript (ECMAScript 5) code, but this can come at a significant cost. Fortunately, compilers do not generate arbitrary code and every compiler we've studied only generates code in a restricted *sub-language of JavaScript*. This section identifies sub-languages that compilers (implicitly) use. In fact, in several cases, we find that multiple, independently developers compilers use the same sub-language of JavaScript. Breakout makes these sub-languages explicit: the `breakout` function (figure 1) consumes a program p and a specification of p's sub-language and then exploits properties of the sub-language to produce simpler and faster code.

We classify each sub-language as a composition of four orthogonal JavaScript language features. Each feature is either completely unused (✗), used to its fullest extend (✓), or used in a restricted manner. Figure 4 summarizes the sub-languages that our ten compilers inhabit. Note that the only language that requires all JavaScript features is JavaScript itself! A compiler author can always use Breakout with all features enabled. But, targeting a sub-language of JavaScript will improve performance dramatically.

4.1 Implicit Operations

In JavaScript, an apparent arithmetic expression like `x - 1` may run forever. This occurs when `x` is bound to an object that defines a `valueOf` or `toString` method. JavaScript defines operations to invoke these methods when applied to objects, and the program may define these methods such that they do not terminate. For completeness, Breakout supports all implicit operations, but they are expensive (figure 2a). Fortunately, the only language that requires all implicits is JavaScript itself (the ✓ in the **Impl** column of figure 4).

No Implicits Several compilers do not need JavaScript to make any implicit calls (✗ in the **Impl** column). For these compilers, Breakout can safely assume that all arithmetic expressions terminate.

Concatenation Only In JavaScript, the `+` operator is overloaded to perform addition and string concatenation, and some compilers rely on `+` to invoke `toString` to concatenate strings (`+` in the **Impl** column). For example, the JSweet Java compiler relies on this behavior, since Java overloads `+` and implicitly invokes `toString` in a similar way. For these compilers, Breakout desugars the `+` operator to expose implicit calls to `toString` and assumes that other operators do not invoke implicit methods.

4.2 Arity Mismatch Behavior

JavaScript has no arity-mismatch errors: any function may receive more arguments (or fewer) that it declares in its formal argument list. When a function receives fewer arguments, the elided arguments are set to the default value `undefined`. All the arguments, including extra ones, are available as properties of a special `arguments` object, which is an implicit, array-like object that every function receives. Some compilers do not leverage this behavior at all (✗ in the **Args** column). Breakout has full support for `arguments`, but we also identify two restricted variants that compilers use in practice.

Variable Arity Functions Many compilers use `arguments` to simulate variable-arity functions (**V** in the **Args** column).

To restore the continuation frame of a variable-arity function, Breakout applies it to its `arguments` object instead of explicitly applying it to its formal arguments:

```
f.apply(this, arguments)
```

However, this simple approach breaks down when `arguments` is used in other ways.

Optional Arguments A different problem arises when `arguments` simulates optional arguments. Suppose a function `f` has formals `x` and `y` and that both are optional. If so, `f` can use `arguments.length` to check if it received fewer arguments that it declared and then initialize the missing arguments to default values. However, this does not affect `arguments.length`. Therefore, if we restore `f` by applying it to its `arguments` object (i.e., like a variable-arity function), then the default values will be lost. So, we need to restore `f` by explicitly applying it to its formal arguments, i.e., `f.call(this, x, y)`. This is how we restore ordinary functions (§3.1), thus we don't need a sub-language when the program uses optional arguments.

Mixing Optional and Variable Arity We've seen that variable-arity functions and functions with optional arguments need to be restored in two different ways. However, we also need to tackle functions that mix both features (**M** in the **Args** column). To restore these functions, we effectively apply both approaches simultaneously: we pass formal parameters explicitly and the `arguments` object as a field on the reified continuation frame.

However, even this does not cover the full range of possible behaviors. For example, JavaScript allows formal arguments to be aliased with fields in the `arguments` array. Therefore, if a function uses a formal argument name and the `arguments` object to access the same location, the approach above will break. Breakout supports this behavior, but it comes at a cost and is only necessary when the source language is JavaScript.

4.3 Getters, Setters, and Eval

In principle, any read or write to a field may invoke a getter or a setter, and any getter or setter may have an infinite loop. Fortunately, this issue does not arise for most source languages (✗ in the **Getters** column). For example, Scala and Python support getters and setters, but ScalaJS and PyJS do not use JavaScript's getters and setters to implement them. On the other hand, Dart does make use of getters and setters but only calls trivial internal functions that terminate (**T** in the **Getters** column). Therefore, we can safely omit instrumentation for getters for Dart. If a compiler generates JavaScript with getters and setters, Breakout can instrument all field reads and writes to capture continuations. However, Breakout also supports a simple annotation (written as a JavaScript comment) that indicates that an expression may trigger a getter or a setter. Therefore, a compiler can be modified to produce this annotation where necessary, which

```
var d = 0, i = 0, ticks = 0, lastTime = 0, v = 0;

function resetTime() { d = 0; }

function estimateElapsed() {
  d++;
  if (i-- === 0) {
    const now = Date.now();
    v = ticks / (now - lastTime); lastTime = now; ticks = t * velocity;
    i = ticks;
  }
  return d / v;
}

function defer(k) { setTimeout(k, 0); } // enqueues k() in the event queue

var mustPause = false, saved;
function maySuspend() {
  if (estimateElapsed() >= δ) {
    resetTime();
    C(function (k) {
      return defer(function() {
        if (mustPause) { saved = k; onPause(); } else { k(); }
      }); }); } }
```

Figure 5. Breakout inserts `maySuspend()` into programs to support long computations.

avoids the cost of instrumenting every field access in the program, which can be prohibitively expensive.

Breakout supports `eval` by rewriting occurrences of `eval` to invoke the Breakout compiler, which is straightforward since it is written in TypeScript. However, the compiler and its dependencies are nearly 5MB of JavaScript and takes much longer than the browser's native `eval` function. Fortunately, most languages do not require JavaScript's `eval` (✗ in the **Eval** column). (In fact, compilers tend not to support `eval` even when the source language features it.) Dart and Pyret are two exceptions: they use `eval` as a form of compression: they dynamically generate lots of trivial functions (e.g., value constructors) that very obviously terminate (**T** in the **Eval** column). In these cases, it makes sense to leave `eval` uninstrumented. Finally, we note that the best way to support `eval` in the source language is to lightly modify the source language compiler to pass the `breakout` function an AST instead of a string. (Breakout uses a standard representation of JavaScript ASTs.) This avoids needlessly parsing and regenerating JavaScript at runtime.

5 Execution Control

We can now say more about the options that `breakout` (figure 1) takes along with the program p to compile: (1) the implementation strategy for continuations (§3.1), (2) the sub-language that p inhabits (§4), (3) whether breakpoints and single-stepping are desired, and (4) whether p requires a deep stack. Breakout transforms p into an equivalent program that (1) runs without freezing the browser tab, (2) can be gracefully terminated at any time, (3) that can simulate blocking operations, and (4) optionally supports deep stacks and debugger operations.

5.1 Long Computations and Graceful Termination

To support long-running computations, Breakout instruments p such that every function and loop calls the `maySuspend` function (figure 5), which may interrupt the computation (by capturing its continuation with $\mathcal{C}$) and schedule it for resumption (using `defer`[3]). The parameter δ determines how frequently these interrupts occur. These interruptions gives the browser an opportunity to process other events, which may include, for example, a user's click on a "Pause" button. To support pausing, `maySuspend` checks if the `mustPause` flag is set and calls the `onPause` callback (from the `pause` method in figure 1) instead of restoring the saved continuation.

The critical piece of this function is the mechanism it uses to calculate elapsed time. The `estimateElapsed` function produces an estimate of time that has elapsed since the program started or the last call to `resetTime`. To do so, it counts the number of times it is called (d) and maintains an estimate of the rate at which it is called (v). The parameter t determines how frequently the function checks the system time and thus the accuracy of the estimate. If the true function-call rate is v', then the estimated time will be off by a factor of $\frac{v}{v'}$, until we resample the system time.

This approach is significantly less expensive that Skulpt's approach—which is to report the exact time elapsed—and more accurate than Pyret's approach—which is to assume a fixed execution rate for all programs. Since this approach relies on sampling, it does lose precision. The table in appendix B applies all three techniques to a subset of our benchmarks and reports the mean interrupt interval and its standard deviation for each case. In one particularly bad example, the mean interrupt interval (μ) is 108.3ms. with standard deviation (σ) 88.91ms. However, by Chebyshev's inequality—$\Pr(|X - \mu| \geq k\sigma) \leq 1/k^2$—even in this case, 95% of interrupt intervals will be less than 553ms, which is still low enough for browsers to be responsive.

5.2 Blocking, Deep Stacks, and Debugging

Breakout provides programmatic access to the $\mathcal{C}$ operator so programs compiled to use Breakout can directly suspend their normal execution with an API call, rather than only during a time-scheduled suspension. This can be used to pause the program while it completes a nonblocking operation, like a network request or an input event. The runtime system of a programming language can use this feature to simulate a blocking operation. Breakout's continuations allow us to suspend an arbitrary JavaScript context while a nonblocking operation completes.

In addition, certain browsers (e.g., Firefox and mobile browsers) have very shallow stacks, which is a problem for recursive programs. With the `'deep'` option for `stacks` (figure 1), Breakout can simulate an arbitrary stack depth (up to heap size). This mode tracks a stack depth counter that is updated in every function call. On reaching a predefined limit, the stack is captured with $\mathcal{C}$, and computation resumes with an empty stack (that closes over the captured one). This counter needs just one variable, and so has negligible impact on performance for programs that don't trigger the stack depth counter: i.e., programs that don't need it hardly pay for it. This feature is important to languages like Pyret (§6.3), a functional language designed for introductory curricula that skip low-level machine details.

Finally, Breakout can be configured to enable breakpoints and stepping. It does this by instrumenting the program to invoke `maySuspend` before every statement. For breakpoints, `maySuspend` checks if the currently-paused statement has a breakpoint set. For stepping, we treat the program as if breakpoints are set on every statement. Breakout exploits JavaScript *source maps* to allow breakpoints to be set using locations in the source program. Source maps are an established technology that is supported by seven of the compilers that we use in our evaluation (§6). For these languages, Breakout is the first Web IDE that supports breakpoints and single-stepping, and with no language-specific effort. (Figure 6 shows our Web IDE providing full execution control, with languages selected from the dropdown.)

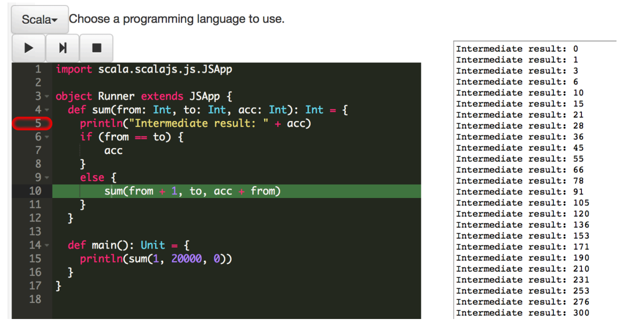

Figure 6. Debugging Scala with Breakout, single-stepped to the highlighted line. The red gutter marker is a breakpoint.

Browser	Platform
Google Chrome 60	Windows 10, Core i5-4690K, 16GB RAM
Mozilla Firefox 56	Windows 10, Core i5-4690K, 16GB RAM
Microsoft Edge 40	Windows 10, Core i5-4690K, 16GB RAM
Apple Safari 11	MacOS 10.13, Core i5-3470S, 8GB RAM
ChromeBook	ChromeOS, Celeron N3060, 4GB RAM

Figure 7. The platforms that we use to benchmark Breakout.

6 Evaluation

We now evaluate Breakout on a suite of ten languages. Breakout has a cost, but its cost must be understood in context. Without Breakout, almost all benchmarks either freeze the browser tab or overflow the JavaScript stack. Furthermore, Breakout supports the first Web IDE with execution control

[3]We present a simple `defer` function in figure 5; our implementation uses a faster but more complex version using `postMessage`.

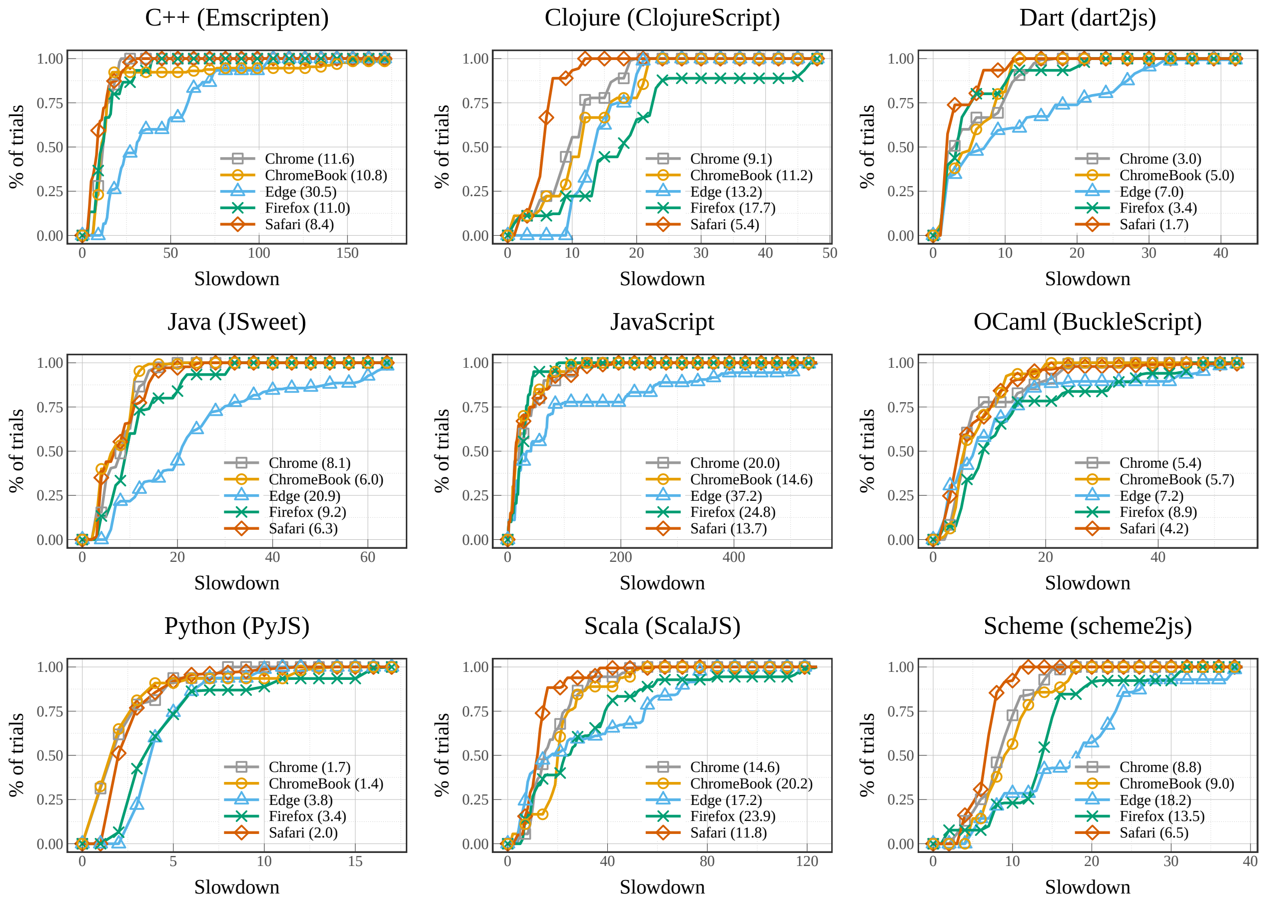

Figure 8. CDFs of Breakout's slowdown on nine languages. The median slowdown is in the legend.

Platform	Continuations	Constructors
Edge	checked-return	dynamic
Safari	exceptional	desugar
Firefox	exceptional	desugar
Chrome	exceptional	desugar

Figure 9. The best implementation strategy for continuations and constructors for each browser ($p < 0.01$).

for eight of these languages. The two exceptions are Python and Pyret, whose Web IDEs we can compare against directly.

Platform selection We run all programs on the four major browsers (Chrome, Firefox, Edge, and Safari) and on a $200 ChromeBook. (See Figure 7 for more detailed specifications.)

6.1 Breakout on Nine Languages

For our first experiment, we report the cost of Breakout using the compilers and benchmarks in figure 4. (We exclude Pyret here, devoting §6.3 to it.) For eight of these, we make *no changes to the compiler* and simply apply Breakout to the compiler output. The only exception is PyJS, which produces JavaScript embedded in a webpage: we modify the compiler to produce a standalone JavaScript file. When possible, we use the Computer Language Benchmarks Game [9] (formerly, the Great Computer Language Shootout), which is a collection of toy problems with solutions in several languages, as our benchmarks. We supplement the Shootout benchmarks with language-specific benchmarks in some cases. We only exclude benchmarks that use language or platform features that the compiler does not support (e.g., files, foreign function interface, etc.). Therefore, the excluded benchmarks cannot run in the browser even without Breakout, so this is not a Breakout limitation. We run each benchmark ten times, both with and without Breakout. Finally, we ensure all benchmarks run for at least two seconds without Breakout by editing the number of iterations. This ensures that that the benchmark yields control several times and thus Breakout is properly exercised by every benchmark. We have 147 benchmarks across all ten languages.

For each compiler, we configure Breakout to exploit the sub-language it generates (as discussed in figure 4). Breakout also provides three strategies for implementing continuations and two strategies for supporting constructors (§3.1). We use microbenchmarks to choose the best settings for each browser, which are shown in figure 9. Finally, we configure Breakout to yield control every 100 ms, which ensures that the browser is very responsive. The slowdown that we report is the ratio of running time with and without Breakout.

We summarize this experiment in figure 8, which shows empirical CDFs of the slowdown for each language. In these graphs, the x-axis shows the slowdown and y-axis shows the fraction of trials with slowdown less than x. We also report the median slowdown for each platform in each graph's legend (in parentheses). For brevity, we do not report the slowdowns for individual benchmarks in the main body of the paper. However, appendix D has bar graphs that report the mean slowdown and confidence intervals for the mean for each benchmark and platform.

Discussion Figure 8 shows that (1) there is no best platform for Breakout, (2) the cost of Breakout depends on the source language compiler, and (3) the sub-language of JavaScript has a significant impact on performance.

We find that the slowdown tends to be much lower on Chrome and Safari than Edge and Firefox. However, we spent months developing Breakout using Chrome and Safari; thus, we speculate that the slowdowns on Firefox and Edge can be made much lower. We are pleasantly surprised that the slowdown on our ChromeBook is comparable to the slowdown on Chrome, despite the fact that the ChromeBook has far less cache and RAM than the desktop.

We also find that the cost of Breakout varies significantly by source language compiler. For example, the median slowdown on PyJS ranges from 1.7x—3.8x across all platforms. In contrast, Breakout performs more poorly on ScalaJS, with slowdowns ranging from 11.8x—23.9x. We attribute these differences to how these languages are compiled to JavaScript. ScalaJS directly translates the Scala standard library implementation to JavaScript, instead of mapping Scala data structures to JavaScript's builtin data structures. (ScalaJS exposes JavaScript data structures to Scala programs, but our benchmarks do not exploit them.) PyJS is more lightweight and maps Python data structures to JavaScript's builtin data structures. We could improve performance on ScalaJS by applying Breakout more selectively to the standard library. (Note that selective transformation does not work if the calling convention changes, as it does in CPS.)

Finally, the cost of Breakout crucially depends on identifying a sub-language of JavaScript that the compiler produces. The most extreme example is when the source language is JavaScript itself, so we cannot make any restrictive assumptions. Implicit operations and getters are the main culprit: the other nine languages either don't need them or use restricted variants of these features (figure 4). We advise compiler writers who want to leverage Breakout to avoid these features. Fortunately, existing compilers already do.

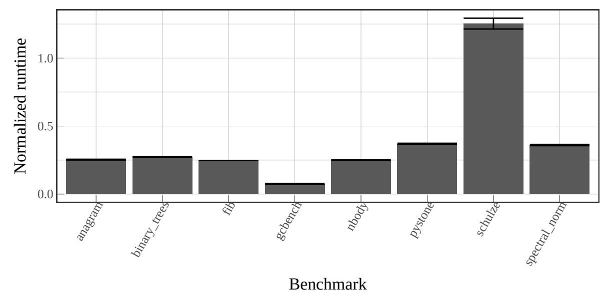

Figure 10. Breakout vs. Skulpt (lower ⇒ Breakout is better).

Native baselines The slowdown that we report is the only meaningful slowdown for platforms that cannot run native code (e.g., ChromeBooks) and languages that only compile to JavaScript (e.g., Pyret). For other cases, we report the slowdown incurred by compiling to JavaScript instead of running natively in appendix C. These slowdowns are not cause by Breakout, but by running code in the browser.

6.2 Case Study: Python (Skulpt)

In §2 and §6.1, we applied Breakout to PyJS to get execution control for Python in the browser. We now compare our approach to Skulpt, which is another Python to JavaScript compiler that has its own execution control. Skulpt is also widely used by online Python courses at Coursera [62, 63], by Rice University's introductory computer science courses [67], by online textbooks [20, 40, 41], and by several schools [61, 72, 73]. Skulpt can be configured to either timeout execution after a few seconds or to yield control at regular intervals, similar to Breakout. Therefore, it is natural to compare Breakout to Skulpt. Unfortunately, Skulpt's implementation of continuations is a new feature that is quite unreliable and often fails to resume execution after yielding. Therefore, we perform the following experiment that puts Breakout at a disadvantage: we configure Skulpt to neither yield nor timeout and compare its performance to Breakout configured to yield every 100 ms. Figure 10 shows the normalized runtime of Breakout: a running time of 1 is the same as the running time of Skulpt; lower means Breakout is faster. Breakout is substantially faster or competitive with Skulpt on all benchmarks, despite its handicap in this experiment.

Though both PyJS and Skulpt are reasonably mature implementations of Python, they have their differences: they each seem to pass only portions of the CPython test suite, and each fail on some benchmarks. Indeed, Skulpt passes only 8 of our 16 benchmarks. Nevertheless, we believe that this experiment shows that Breakout is already fast enough

```
function eachLoop(fun, start, stop) {
  var i = start;
  function restart(_) {
    var res = thisRuntime.nothing;
    if (--thisRuntime.GAS <= 0) { res = thisRuntime.makeCont(); }
    while(!thisRuntime.isContinuation(res)) {
      if (--thisRuntime.RUNGAS <= 0) { res = thisRuntime.makeCont(); }
      else {
        if(i >= stop) {
          ++thisRuntime.GAS;
          return thisRuntime.nothing;
        } else {
          res = fun.app(i);
          i = i + 1;
    } } }
    res.stack[thisRuntime.EXN_STACKHEIGHT++] =
      thisRuntime.makeActivationRecord("eachLoop", restart, true, [], []);
    return res;
  }
  return restart();
}
```

(a) Original, hand-instrumented implementation.

```
function eachLoop(fun, start, stop) {
  for (var i = start; i < stop; i++) { fun.app(i); }
  return thisRuntime.nothing;
}
```

(b) With Breakout, no instrumentation is necessary.

Figure 11. A higher-order function from Pyret's standard library that applies a function to a sequence of numbers.

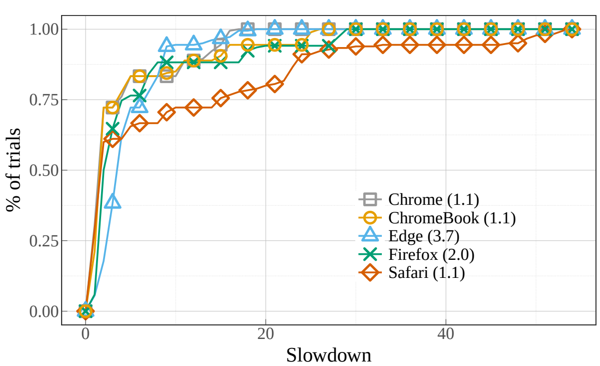

Figure 12. Pyret with Breakout. Running time normalized to Pyret without Breakout.

to support Python Web IDEs, while increasing reliability and thus illustrating its value.

6.3 Case Study: Pyret

Pyret [56] is a mostly-functional programming language that runs entirely in the browser; it supports proper tail calls, blocking I/O, a REPL, and interactive animations; it allows users to gracefully terminate running programs; and it takes care to not freeze the browser tab. Despite five years of engineering and thousands of users, Pyret still suffers from issues that produce wrong results or freeze the browser tab [23, 28, 29, 31–33, 49–54].

Current Pyret Implementation The final phase of the Pyret compiler is a single pass that both (1) translates Pyret expressions to JavaScript and (2) produces JavaScript that is instrumented to save and restore the stack. The core Pyret runtime system, which is 6,000 lines of JavaScript, also serves two roles: it (1) implements Pyret's standard library and (2) and carefully cooperates with compiled Pyret code. Since Pyret encourages functional programming, its standard library has several higher-order functions, which are implemented in JavaScript to improve performance. These JavaScript functions are instrumented by hand, since they may appear on the stack when Pyret needs to pause a user-written function.

Due to this close coupling, the Pyret developers inform us that the compiler and runtime system are difficult to maintain. For example, the Pyret runtime has a function called `eachLoop` that merely applies a functional argument to a range of numbers. However, most of its implementation is dedicated to saving and restoring the stack (figure 11a). In fact, this function has been rewritten several times to fix bugs in its instrumentation [34, 35, 55].

Pyret with Breakout We applied Breakout to Pyret, which simplifies both the compiler and runtime system. However, we also find that Breakout (1) largely maintains Pyret's performance, (2) supports Pyret's Web IDE, including animations and the REPL, and (3) exposes new bugs in Pyret.

For the other nine languages, we use Breakout as a blunt instrument: we leave an existing compiler and runtime system unchanged and simply apply Breakout to the output JavaScript. We apply Breakout more carefully to Pyret. First, we strip out the portions of the compiler that manage the stack and modify it to apply Breakout to its output. This shrinks the last phase of the compiler from 2,100 LOC to 1,500 LOC (nearly 30% smaller). Second, 900 LOC of Pyret's runtime system involve stack management. We remove the stack management logic, which shrinks the code to 350 LOC (60% smaller), and we modify Pyret's build process to apply Breakout to this code. For example, Breakout lets us remove *all* the instrumentation from `eachLoop` (figure 11). We envision that compiler authors who choose to use Breakout themselves will use it in this way to instrument only the necessary fragments of the language's libraries.

For the other nine languages, we use Breakout to only support long-running programs, stepping, and graceful termination. However, Pyret requires many more features (REPL, animations, blocking I/O). All these features are implemented in JavaScript and use an internal function to pause execution while they do their work. We made these features work by simply replacing Pyret's pause function's implementation with one in terms of Breakout's `pause`.

Finally, Breakout exposed new and nontrivial bugs—both acknowledged by the language's authors—in Pyret's stack saving mechanism [citation elided for anonymity]. Two functions in the Pyret runtime system wrongly assumed that they would never appear on the stack during continuation capture. By simplifying the runtime system, Breakout made these errors evident.

Performance Figure 12 compares Breakout-based Pyret against existing Pyret. On Chrome and Safari, the median slowdown is 1.1x and a number of benchmarks are *faster* with Breakout than with Pyret' own implementation of continuations. The median slowdown on Firefox (2.0x) and Edge (3.7x) is disappointing but, as mentioned in §6.1, these are browsers we did not use in developing Breakout.

Unfortunately, we also find that some of our benchmarks have more significant slowdowns (up to 20x). However, all these benchmarks share the following attribute: they are deeply recursive programs that require more stack space than the browser provides. Breakout supports deep stacks (§5.2), but our implementation performs far more poorly than Pyret's. We expect to address this problem in the near future: Breakout should be able to output exactly the kind of code that Pyret does to implement deep stacks.

Future work There are several more ways to improve our performance on Pyret. For example, Pyret performs several optimizations in its last phase that are entangled with its implementation of continuations. We omit these optimizations in our prototype compiler. Once ported to Breakout, they can be applied to other languages too. Furthermore, there are several ways to simplify Pyret's libraries now that we have Breakout. For example, many functions use expensive abstractions, such as callbacks and continuations, to avoid appearing on the stack during continuation capture. We could rewrite these functions using loops and apply Breakout, which may improve performance.

7 Related Work

Web Workers Web Workers [76] are essentially isolated processes for JavaScript and a Web IDE can use them to terminate a user's program [6]. However, unlike Breakout, Workers do not provide richer execution control (e.g., pausing or breakpointing) or deep stacks. Unlike Breakout, they also have a limited interface: they cannot directly access the DOM, and can communicate only through special shared datatypes [44] or by message passing.

WebAssembly WebAssembly [21] is a new low-level language in Web browsers. As of this writing, garbage collection, threads, tail calls, and host bindings (e.g., access to the DOM) are in the feature proposal stage [1]. Therefore, WebAssembly currently does not provide enough features to prototype a Breakout-like solution, nor are we aware of any multi-language demonstrations comparable to §6. Nevertheless, as it matures, Web IDEs may want to switch to it. For that to happen, WebAssembly will need to support the kind of execution control that Breakout provides for JavaScript.

Runtime Systems for the Web There are a handful of compilers that have features that overlap with Breakout, such as continuations [36, 69, 79], tail calls [36, 56, 69, 79], and graceful termination [56, 69, 79]. GopherJS [18] supports goroutines using mechanisms related to Breakout too. These compilers commingle language translation and working around the platform's limitations. Breakout frees compiler authors to focus on language translation. In addition, Breakout adds continuations to JavaScript itself and supports features such as exceptional control flow, constructors, ES6 tail calls, and more. Furthermore, Breakout supports a family of continuation implementation strategies and we show that the best strategy varies by browser.

Doppio [74] and Whalesong [79] provide features like threads and more in the browser by using a bytecode interpreter, i.e., by not using the JavaScript stack. Furthermore, these are bytecode interpreters for other platforms (JVM and Racket, respectively). Thus, existing compilers and libraries would have to change significantly to use them. Breakout has no such restrictions.

Browsix [57] acts as an "operating system" for processes in Web Workers. Therefore, it inherits Web Workers' restrictions: workers cannot share JavaScript values and cannot interact with the Web page. It also does not provide deep stacks. Breakout allows code to run in the main browser thread, enabling access to the DOM and allowing execution control for IDEs.

Continuations on Uncooperative Platforms Many past projects have investigated implementing continuations in other platforms that do not natively support them, from C to .NET [3, 22, 36, 43, 48, 68, 78]. They use a variety of strategies ranging from CPS with trampolines, to C's `setjmp` and `longjmp` to *effectively* provide tail calls (by using the same amount of net space), to a precursor of what we describe in §3. However, none of these goes in depth beyond providing continuation support, and in particular do not provide the other contributions discussed in §1.

Blocking Calls to Nonblocking APIs Pivot [39] is a Web isolation framework that provides blocking RPCs between iframes. It rewrites a frame's code into a generator that invokes nonblocking calls when iterated. However, Pivot only supports interrupts *between* frames rather than for code inside a single frame, and does not provide the other features discussed in §1. Breakout makes arbitrary JavaScript programs interruptible, and does not suffer the disadvantages of generators discussed in §3.

8 Conclusion

We have presented Breakout, a JavaScript-to-JavaScript compiler that enriches its source with execution control. This enables Web IDEs to work around the limitations of the JavaScript execution model. We present a new compilation strategy to support continuations, identify sub-languages of JavaScript to improve performance, and improve the responsiveness/performance tradeoff. We evaluate Breakout by showing that it smoothly composes with ten compilers for a diverse set of programming languages, enriching their output without imposing any maintenance burdens, and enabling the creation of better and more responsive IDEs.

References

[1] WebAssembly: Features to Add after the MVP. https://github.com/WebAssembly/design/blob/71c97d9081ac6d64497d586e8f45c8145ed3f86e/FutureFeatures.md. Accessed Nov 10 2017.

[2] Monica Anderson, Lee Rainie, Aaron Smith, Maeve Duggan, Dana Page, Margaret Porteus, and Shannon Greenwood. 2015. *Technology Device Ownership: 2015*. Technical Report. Pew Research Center.

[3] Henry G. Baker. 1995. CONS Should Not CONS Its Arguments, Part II: Cheney on the M.T.A. *ACM SIGPLAN Notices* 30, 9 (Sept. 1995), 17–20.

[4] Carl Bruggeman, Oscar Waddell, and Kent R. Dybvig. 1996. Representing Control in the Presence of One-shot Continuations. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*.

[5] Sebastian Burckhardt, Manuel Fahndrich, Peli de Halleux, Sean McDirmid, Michal Moskal, Nokolai Tillmann, and Jun Kato. 2013. It's Alive! Continuous Feedback in UI Programming. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*.

[6] Benjamin Canou, Roberto Di Cosmo, and Grégoire Henry. 2017. Scaling up Functional Programming Education: Under the Hood of the OCaml MOOC. In *ACM International Conference on Functional Programming (ICFP)*.

[7] Teams on CodePen. https://codepen.io/pro/teams/. Accessed Nov 10 2017.

[8] Codio. https://codio.com. Accessed Nov 10 2017.

[9] The Computer Language Benchmarks Game. https://benchmarksgame.alioth.debian.org. Accessed Nov 10 2017.

[10] Confusing behavior with exceptions inside goroutines. https://github.com/gopherjs/gopherjs/issues/225. Accessed Nov 10 2017.

[11] Evan Czaplicki and Stephen Chong. 2013. Asynchronous Functional Reactive Programming for GUIs. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*.

[12] defer + go routine/channel + blocking = possible crash. https://github.com/gopherjs/gopherjs/issues/381. Accessed Nov 10 2017.

[13] Deferring runtime.Gosched() causes error. https://github.com/gopherjs/gopherjs/issues/426. Accessed Nov 10 2017.

[14] Rip Empson. Treehouse Lands $7M From Kaplan, Social+Capital To Help You Learn To Code. https://techcrunch.com/2013/04/09/treehouse-lands-7m-from-kaplan-socialcapital-to-help-you-learn-to-code/. Accessed Nov 10 2017.

[15] execLimit not triggering on large xrange loops. https://github.com/skulpt/skulpt/issues/711. Accessed Nov 10 2017.

[16] fix for goroutine blocking bookkeeping. https://github.com/gopherjs/gopherjs/issues/209. Accessed Nov 10 2017.

[17] Cormac Flanagan, Amr Sabry, Bruce F. Duba, and Matthias Felleisen. 1993. The Essence of Compiling with Continuations. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*.

[18] GopherJS. https://github.com/gopherjs/gopherjs. Accessed Nov 10 2017.

[19] Philip J. Guo, Jeffery White, and Renan Zanelatto. 2015. Codechella: Multi-user program visualizations for real-time tutoring and collaborative learning. In *IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*.

[20] Mark Guzdial and Barbara Ericson. CS Principles: Big Ideas in Programming. http://interactivepython.org/runestone/static/StudentCSP/index.html. Accessed Nov 10 2017.

[21] Andreas Haas, Andreas Rossberg, Derek L. Schuff, Ben L. Titzer, Michael Holman, Dan Gohman, Luke Wagner, Alon Zakai, and JF Bastien. 2017. Bringing the Web Up to Speed with WebAssembly. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*.

[22] R. Hieb, R. Kent Dybvig, and Carl Bruggeman. 1990. Representing Control in the Presence of First-class Continuations. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*.

[23] Wayne Iba. too much recursion in 2048 example code. https://github.com/brownplt/pyret-lang/issues/555.

[24] List of Languages that Compile to JS. https://github.com/jashkenas/coffeescript/wiki/list-of-languages-that-compile-to-js. Accessed Nov 10 2017.

[25] JetStream 1.1. http://browserbench.org/JetStream/. Accessed Nov 10 2017.

[26] Khan Academy: Computer programming. https://www.khanacademy.org/computing/computer-programming. Accessed Nov 10 2017.

[27] Kraken. https://wiki.mozilla.org/Kraken. Accessed Nov 10 2017.

[28] Shriram Krishnamurthi. impossible(?) to kill infinite loop w/ reactor. https://github.com/brownplt/pyret-lang/issues/839.

[29] Shriram Krishnamurthi. calling plotting functions inside a reactor makes program unstoppable. https://github.com/brownplt/pyret-lang/issues/1089.

[30] Larceny Benchmarks. http://www.larcenists.org/benchmarks2009all.html. Accessed Nov 10 2017.

[31] Benjamin S. Lerner. An infinite loop hangs big-bang. https://github.com/brownplt/pyret-lang/issues/508.

[32] Benjamin S. Lerner. Responsiveness for rendering huge data at the REPL. https://github.com/brownplt/code.pyret.org/issues/37.

[33] Benjamin S. Lerner. Bignums considered harmful. https://github.com/brownplt/pyret-lang/issues/1118.

[34] Benjamin S. Lerner. eachLoop was simply broken: it did not restore the stack properly. https://github.com/brownplt/pyret-lang/commit/812d1c73ddb0d2aee8bef4ebc0d44a811fafd908.

[35] Benjamin S. Lerner. Fix broken eachLoop. https://github.com/brownplt/pyret-lang/commit/b7f9c9016754ed2a86e852980708822e67ff1fed.

[36] Florian Loitsch. 2007. Exceptional Continuations in JavaScript. In *Workshop on Scheme and Functional Programming*.

[37] Connie Loizos. Zach Sims Of Codecademy On Running A Company That (Still) Doesn't Charge Users. https://techcrunch.com/2016/01/25/zach-sims-of-codecademy-on-running-a-company-that-still-doesnt-charge-users/. Accessed Nov 10 2017.

[38] Ingrid Lunden. Codecademy, the free online coding school, raises another $30M led by Naspers. https://techcrunch.com/2016/07/12/codecademy-the-free-online-coding-school-raises-another-30m-led-by-naspers/. Accessed Nov 10 2017.

[39] James Mickens. 2014. Pivot: Fast, Synchronous Mashup Isolation Using Generator Chains. In *IEEE Security and Privacy (Oakland)*.

[40] Brad Miller and David Ranum. Problem Solving with Algorithms and Data Structures using Python. http://interactivepython.org/runestone/static/pythonds/index.html. Accessed Nov 10 2017.

[41] Brad Miller and David Ranum. How to Think Like a Computer Scientist: Interactive Edition. http://interactivepython.org/runestone/static/thinkcspy/index.html. Accessed Nov 10 2017.

[42] Miscompilation of functions w/ defer statements. https://github.com/gopherjs/gopherjs/issues/493. Accessed Nov 10 2017.

[43] Neil Mix and Dan Grisby. Narrative JavaScript. https://sourceforge.net/projects/narrativejs. Accessed Nov 10 2017.

[44] Inc. Mozilla. SharedArrayBuffer. https://developer.mozilla.org/en-US/docs/Web/JavaScript/Reference/Global_Objects/SharedArrayBuffer.

[45] OPerf Micro. https://www.typerex.org/operf-micro.html. Accessed Nov 10 2017.

[46] Krzysztof Palacz. 2008. The Lively Kernel Application Framework. In *International Conference on Scalable Vector Graphics*.

[47] Sarah Perez. Online Learning Service Pluralsight Acquires Code School For $36 Million. https://techcrunch.com/2015/01/26/online-learning-service-pluralsight-acquires-code-school-for-36-million/. Accessed Nov 10 2017.

[48] Greg Pettyjohn, John Clements, Joe Marshall, Shriram Krishnamurthi, and Matthias Felleisen. 2005. Continuations from generalized stack inspection. In *ACM International Conference on Functional Programming (ICFP)*.
[49] Joe Gibbs Politz. Make module loading stack safe. https://github.com/brownplt/pyret-lang/issues/145.
[50] Joe Gibbs Politz. Printing large values is not stack-safe. https://github.com/brownplt/pyret-lang/issues/146.
[51] Joe Gibbs Politz. Stopping (indefinitely) stopped programs. https://github.com/brownplt/pyret-lang/issues/163.
[52] Joe Gibbs Politz. Too much recursion. https://github.com/brownplt/pyret-lang/issues/213.
[53] Joe Gibbs Politz. Caja interferes with stack management/events during big-bang. https://github.com/brownplt/pyret-lang/issues/512.
[54] Joe Gibbs Politz. Stack management fails on shallow, but long-lasting, recursion. https://github.com/brownplt/pyret-lang/issues/596.
[55] Joe Gibbs Politz. more fixing of eachLoop. https://github.com/brownplt/pyret-lang/commit/84445476f0741f05251cf22b3cc06969e85e1ab9.
[56] Joe Gibbs Politz, Benjamin S. Lerner, and Shriram Krishnamurthi. Pyret. https://www.pyret.org/. Accessed Jul 5 2017..
[57] Bobby Powers, John Vilk, and Emery D. Berger. 2017. Browsix: Bridging the Gap Between Unix and the Browser. In *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*.
[58] Roldan Pozo and Bruce Miller. SciMark 2.0. http://math.nist.gov/scimark2/. Accessed Nov 10 2017.
[59] Programs can incorrectly terminate with yieldLimit set to suspend at regular intervals. https://github.com/skulpt/skulpt/issues/723. Accessed Nov 10 2017.
[60] PyPy Benchmarks. https://bitbucket.org/pypy/benchmarks. Accessed Nov 10 2017.
[61] PythonRoom. https://pythonroom.com. Accessed Nov 10 2017.
[62] Charles Severance. Programming for Everybody (Getting Started with Python). https://www.coursera.org/learn/python. Accessed Nov 10 2017.
[63] Charles Severance. Python Data Structures. https://www.coursera.org/learn/python-data. Accessed Nov 10 2017.
[64] Natasha Singer. 2017. Amid Stiff Competition, Apple Devices Lose Luster in American Schools. *The New York Times* (2017), B2. Mar 3 2017. https://www.nytimes.com/2017/03/02/technology/apple-products-schools-education.html.
[65] Dorai Sitaram and Matthias Felleisen. 1990. Control Delimiters and Their Hierarchies. *LISP and Symbolic Computation* 3, 1 (May 1990), 67–99.
[66] Skulpt. http://www.skulpt.org. Accessed Nov 10 2017.
[67] Terry Tang, Scott Rixner, and Joe Warren. 2014. An Environment for Learning Interactive Programming. In *ACM Technical Symposium on Computer Science Education (SIGCSE)*.
[68] David Tarditi, Peter Lee, and Anurag Acharya. 1992. No Assembly Required: Compiling Standard ML to C. *ACM Letters on Programming Languages and Systems (LOPLAS)* 1, 2 (June 1992), 161–177.
[69] Eric Thivierge and Marc Feeley. 2012. Efficient Compilation of Tail Calls and Continuations to JavaScript. In *Workshop on Scheme and Functional Programming*.
[70] Tight loop in goroutine never yields. https://github.com/gopherjs/gopherjs/issues/698. Accessed Nov 10 2017.
[71] Ton80. https://github.com/dart-lang/ton80. Accessed Nov 10 2017.
[72] Trinket. https://trinket.io. Accessed Nov 10 2017.
[73] Tynker. https://www.tynker.com. Accessed Nov 10 2017.
[74] John Vilk and Emery D. Berger. 2014. Doppio: Breaking the Browser Language Barrier. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*.
[75] Vocareum. https://www.vocareum.com. Accessed Nov 10 2017.
[76] W3C. Web Workers. https://www.w3.org/TR/workers/.
[77] Why is my browser freezing when I submit an exercise? https://help.codecademy.com/hc/en-us/articles/220803187. Accessed Nov 10 2017.
[78] James Wright. JWACS. http://chumsley.org/jwacs/index.html. Accessed Nov 10 2017.
[79] Danny Yoo and Shriram Krishnamurthi. 2013. Whalesong: Running Racket in the Browser. In *Dynamic Languages Symposium (DLS)*.
[80] Danny Yoo, Emmanuel Schanzer, Shriram Krishnamurthi, and Kathi Fisler. 2011. WeScheme: The Browser is Your Programming Environment. In *Conference on Innovation and Technology in Computer Science Education (ITiCSE)*.

A Limitations of Web IDEs

Most Web IDEs have serious limitations. Due to design of the Web platform, most IDEs can neither run programs that take more than a few seconds nor interrupt nonterminating programs. For example, Codecademy, which claims over 25 million users [37], has a Web IDE for JavaScript. However, the IDE does not have a "stop" button and the only way to interrupt an infinite loop is to refresh the page and lose recent changes. Similarly, CodeSchool, which claims over 40,000 active users [47], has a Web IDE for Python. Unlike Codecademy, the CodeSchool IDE imposes a hard timeout on all computations. For example, it cannot run the following Python program, which counts up to only 1 billion:

```
j = 0
while (j < 1000000000):
  j = j + 1
```

Instead, the IDE aborts with the message, *"TimeoutError: The executor has timed out."* The same problem affects Khan Academy's Web IDE for JavaScript [26], which terminates loops with more than five million iterations, printing *"a while loop is taking too long to run."* These kinds of IDEs are only suitable for trivial programming tasks.

Unfortunately, these problems aren't restricted to online coding schools. Codepen is a Web IDE that users to collaboratively develop Web applications (i.e., HTML, CSS, and JavaScript). This feature makes it popular at several organizations [7]. The CodePen IDE also terminates long-running loops, but continues running the program at the next statement after the loop. For example, given the following JavaScript program—

```
var i = 0; while (i++ < 1000000000); alert(i);
```

—CodePen first prints the warning, *"An infinite loop (or a loop taking too long) was detected, so we stopped its execution. Sorry!"* and then displays the wrong value for `i`!

These kinds of problems also affect Web IDEs developed by the programming languages community. For example, Elm [11] is a statically-typed, Haskell-inspired programming language with a time-traveling debugger that are step through a program's DOM events. However, infinite loops, such as the following, crash the Elm debugger:

```
import Window
loop : a -> a
loop x = loop x
main = plainText (loop "")
```

Python Tutor [19], which is a popular resource for visualizing program execution, cannot handle long-running Python programs either, terminating with a message for the user that such programs are not supported. The Lively Kernel [46], which is a Smalltalk-inspired environment for the browser, also crashes when given an infinite loop. Reason, which is a new OCaml-based language developed at Facebook, lets users try the language online, but stops execution after a short timeout. In contrast, many other languages with online "playgrounds" simply crash the browser.

B Comparison of Estimation Methods

Section §5.1 discusses how different implementation strategies of the `estimateElapsed` method incur a tradeoff between runtime overhead and high variance in latency. We implement three strategies for estimating the elapsed time: (1) assuming a fixed execution rate and monitoring a *countdown* timer, (2) calling `Date.now` to measure the *exact* amount of time elapsed, and (3) *approximating* the time by dynamically updating an estimate of the program's velocity. Figure 13 shows a more detailed evaluation of each of the three strategies. We evaluate 10 Python benchmarks, running 10 iterations of each `estimateElapsed` implementation in Chrome, targeting a latency of 100 milliseconds. Each column shows the mean (μ) and standard deviation (σ) of the latency for each of the 10 benchmarks.

Figure 14 shows the slowdown of *countdown* and *exact* relative to *approximate* on the same 10 Python benchmarks. While figure 13 shows that *approximate* has a higher variance in latency when compared to *exact*, it imposes less overhead on the runtime of programs; *exact* has a median slowdown of 1.47x over *approximate*.

C Native Performance

Tables 1 and 2 show the mean slowdown of our all benchmarks on all platforms relative to various native compilers. We use gcc to compile C++ and the Gambit Scheme interpreter to run Scheme code.

D Detailed Evaluation

Figure 15 presents the slowdown incurred by Breakout on each benchmark in our benchmark suite.

Benchmark	Countdown ($\mu \pm \sigma$)	Approximate ($\mu \pm \sigma$)	Exact ($\mu \pm \sigma$)
anagram	343.7 ± 225.1 ms	108.3 ± 88.91 ms	112.8 ± 29.31 ms
b	122.4 ± 27.35 ms	87.75 ± 39.33 ms	114.3 ± 26.19 ms
binary-trees	146.3 ± 32.55 ms	89.18 ± 24.71 ms	108 ± 3.92 ms
deltablue	386.4 ± 85.99 ms	97.31 ± 21.06 ms	109.2 ± 1.359 ms
fib	67.63 ± 8.398 ms	98.44 ± 12.64 ms	109.3 ± 0.8464 ms
nbody	201.6 ± 39.57 ms	100 ± 29.88 ms	109.4 ± 0.8818 ms
pystone	315.1 ± 112.1 ms	89.1 ± 24.46 ms	109.5 ± 1.414 ms
raytrace-simple	237.2 ± 44.22 ms	98.84 ± 15.28 ms	109.4 ± 0.6615 ms
richards	214.2 ± 54.36 ms	95.57 ± 23.14 ms	109.5 ± 1.152 ms
spectral-norm	182.4 ± 62.91 ms	77.58 ± 27.61 ms	109.7 ± 1.249 ms

Figure 13. A comparison of the three time estimation strategies on a subset of Python benchmarks. Results show the mean (μ) and standard deviation (σ) of the latency between yield intervals. For μ, closer to 100ms is better. For σ, lower is better.

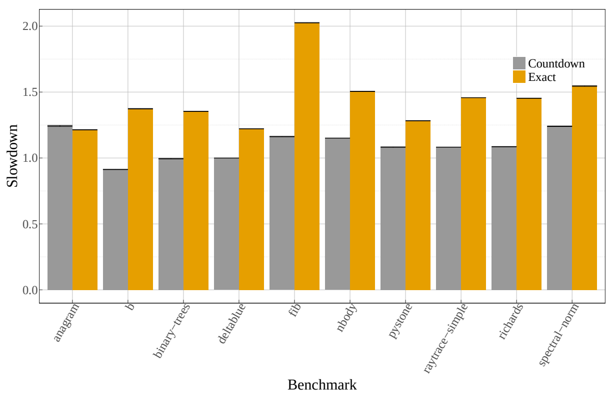

Figure 14. CDFs of Breakout's slowdown on nine languages. The median slowdown is in the legend.

Language	Benchmark	Chrome	Safari	Firefox	ChromeBook	Edge
ocaml	almabench	0.81	0.76	0.70	3.57	2.04
ocaml	bdd	3.28	2.96	4.63	13.58	4.68
ocaml	binarytrees	4.40	3.16	7.89	16.75	8.13
ocaml	boyer	16.49	504.73	23.60	52.41	105.33
ocaml	fannkuch-redux	1.70	1.58	1.70	6.14	18.52
ocaml	fft	6.32	0.93	1.02	24.73	3.33
ocaml	fibonnaci	2.57	1.84	2.61	9.73	3.39
ocaml	format-main	441.60	474.45	768.32	1357.42	2989.60
ocaml	hamming	3.54	5.48	5.45	13.87	11.12
ocaml	kahan-sum	8.63	1.87	3.29	32.77	7.90
ocaml	kb	33.54	2349.82	65.90	128.87	211.62
ocaml	list-main	6.38	7.39	8.07	23.89	30.31
ocaml	mandelbrot	10.25	33.95	10.90	32.81	70.42
ocaml	nbody	1.46	1.90	1.52	7.99	1.31
ocaml	nucleic	1.27	1.31	1.46	4.74	2.65
ocaml	nullable-array	3.17	1.10	1.94	12.39	4.11
ocaml	sequence	8.94	3.03	2.89	37.78	20.81
ocaml	spectral-norm	0.37	0.40	0.38	1.73	0.68
c++	bigfib	2.04	4.15	5.99	6.88	2.09
c++	binary-trees	1.32	1.68	1.42	4.98	1.81
c++	container	1.26	1.33	1.60	4.95	1.03
c++	dry	1.84	1.24	1.16	9.24	1.31
c++	fannkuch-redux	6.25	7.05	16.40	21.96	12.12
c++	fasta	4.00	4.24	5.15	10.55	7.16
c++	float-mm	3.09	4.77	4.30	11.08	2.83
c++	gcc-loops	9.55	10.47	12.30	37.95	7.49
c++	mandelbrot	2.10	2.55	2.87	9.39	2.11
c++	meteor	2.51	4.98	2.86	8.02	1.68
c++	nbody	1.68	2.03	1.90	8.32	1.36
c++	quicksort	1.96	1.82	2.29	4.45	3.04
c++	spectral-norm	0.97	1.10	1.00	4.52	1.57
c++	towers	3.60	4.43	3.71	14.09	6.58
c++	zlib	5.65	4.29	4.01	20.39	2.70
dart-dart2js	binary-tree	1.88	2.13	6.26	7.13	2.01
dart-dart2js	fannkuch-redux	1.12	1.12	0.93	3.87	1.46
dart-dart2js	fasta	1.09	1.10	2.78	3.22	1.21
dart-dart2js	fastaredux	1.01	4.15	2.42	3.81	7.41
dart-dart2js	knucleotide	2.67	2.69	4.85	10.14	6.14
dart-dart2js	nbody	1.21	1.08	1.20	2.92	1.24
dart-dart2js	pidigits	114.66	135.23	1355.63	456.82	353.36
dart-dart2js	regexdna	0.85	0.90	0.86	0.98	0.87
dart-dart2js	regexredux	0.87	0.82	0.95	1.12	0.85
dart-dart2js	revcomp	0.96	1.02	1.07	1.27	1.00
dart-dart2js	spectral-norm	0.97	1.05	0.98	4.56	0.97
java	scimark-fft	1.12	1.00	1.14	5.13	1.09
java	scimark-LU	8.84	1.27	2.09	29.84	7.26
java	scimark-SOR	4.20	1.61	1.18	16.26	1.83
java	scimark-SparseMatmult	2.06	2.52	1.43	9.09	1.12
java	shootout-binarytrees-2	3.53	2.00	7.99	12.39	3.85
java	shootout-fannkuchredux-2	1.35	1.37	1.19	4.57	1.73
java	shootout-mandelbrot	0.88	0.96	1.29	2.94	0.98
java	shootout-nbody	1.16	2.23	1.05	6.95	3.08
java	shootout-spectralnorm	0.72	0.78	0.73	3.85	0.74

Table 1. Slowdown of benchmarks when compiled to JavaScript, with respect to native running time

Language	Benchmark	Chrome	Safari	Firefox	ChromeBook	Edge
scheme	ack	0.8810	0.7359	0.7627	2.240	0.9019
scheme	array1	82.4034	161.2437	2121.2269	320.395	191.8739
scheme	boyer	0.6016	0.6254	0.7252	2.240	0.6602
scheme	browse	7.1673	6.6996	6.8696	25.924	11.3096
scheme	conform	2.3057	1.7673	4.1315	8.239	3.4126
scheme	dderiv	65.5191	36.5954	49.0229	176.168	113.9771
scheme	destruc	0.5097	0.7048	0.5599	1.329	0.7485
scheme	diviter	60.3971	46.9779	46.7647	114.441	141.9338
scheme	divrec	86.2290	75.8855	85.0992	209.496	235.3511
scheme	earley	0.1061	0.0467	0.0742	0.363	0.0676
scheme	fft	0.0049	0.0031	0.0048	0.017	0.0048
scheme	fib	0.0276	0.0173	0.0289	0.116	0.0362
scheme	lattice	0.0225	0.0206	0.0339	0.093	0.0758
scheme	mazefun	0.0461	0.0336	0.0782	0.159	0.0702
scheme	mbrot	0.0051	0.0058	0.0054	0.016	0.0061
scala	BinaryTrees	23.0045	1.8599	61.4117	90.990	23.4266
scala	Fannkuchredux	1.1633	0.8504	0.9947	3.992	7.4693
scala	Fasta	0.5407	0.4974	0.7701	1.941	0.8087
scala	FastaRedux	2.2203	2.1897	1.7571	7.959	3.0657
scala	Knucleotide	0.3587	0.9917	0.7376	1.746	1.0378
scala	LongMicro	1.8858	2.1136	2.9796	6.710	17.7090
scala	Mandelbrot	7.0547	1.2604	4.5806	4.005	9.7964
scala	Meteor	1.8643	1.2461	2.1164	7.107	5.8780
scala	Pidigits3	2.6430	3.1616	3.0117	8.846	3.3983
scala	Regexdna	0.0503	0.0435	0.0445	0.120	0.0439
scala	Regexredux	0.0477	0.0441	0.0507	0.133	0.0454
scala	Spectralnorm	5.5964	1.8199	3.4148	13.132	8.2967
scala	Tracer	0.0589	0.0547	0.0843	0.211	0.0906
clojurescript	binarytrees	0.4604	0.3932	0.5008	1.827	1.2243
clojurescript	collatz	0.8922	0.5847	1.1740	2.942	1.2976
clojurescript	fannkuch	0.4941	0.5406	0.5896	1.899	1.8805
clojurescript	mandelbrot	0.2635	0.2793	0.2790	0.999	0.2943
clojurescript	memuse	0.1672	0.1380	0.1458	0.728	0.4296
clojurescript	meteor	1.2536	1.4421	1.5129	4.833	1.7126
clojurescript	nbody	0.4584	0.7781	0.4772	2.111	0.9165
clojurescript	pidigits	0.1167	0.1383	0.1592	0.452	0.3907

Table 2. Slowdown of benchmarks when compiled to JavaScript, with respect to native running time

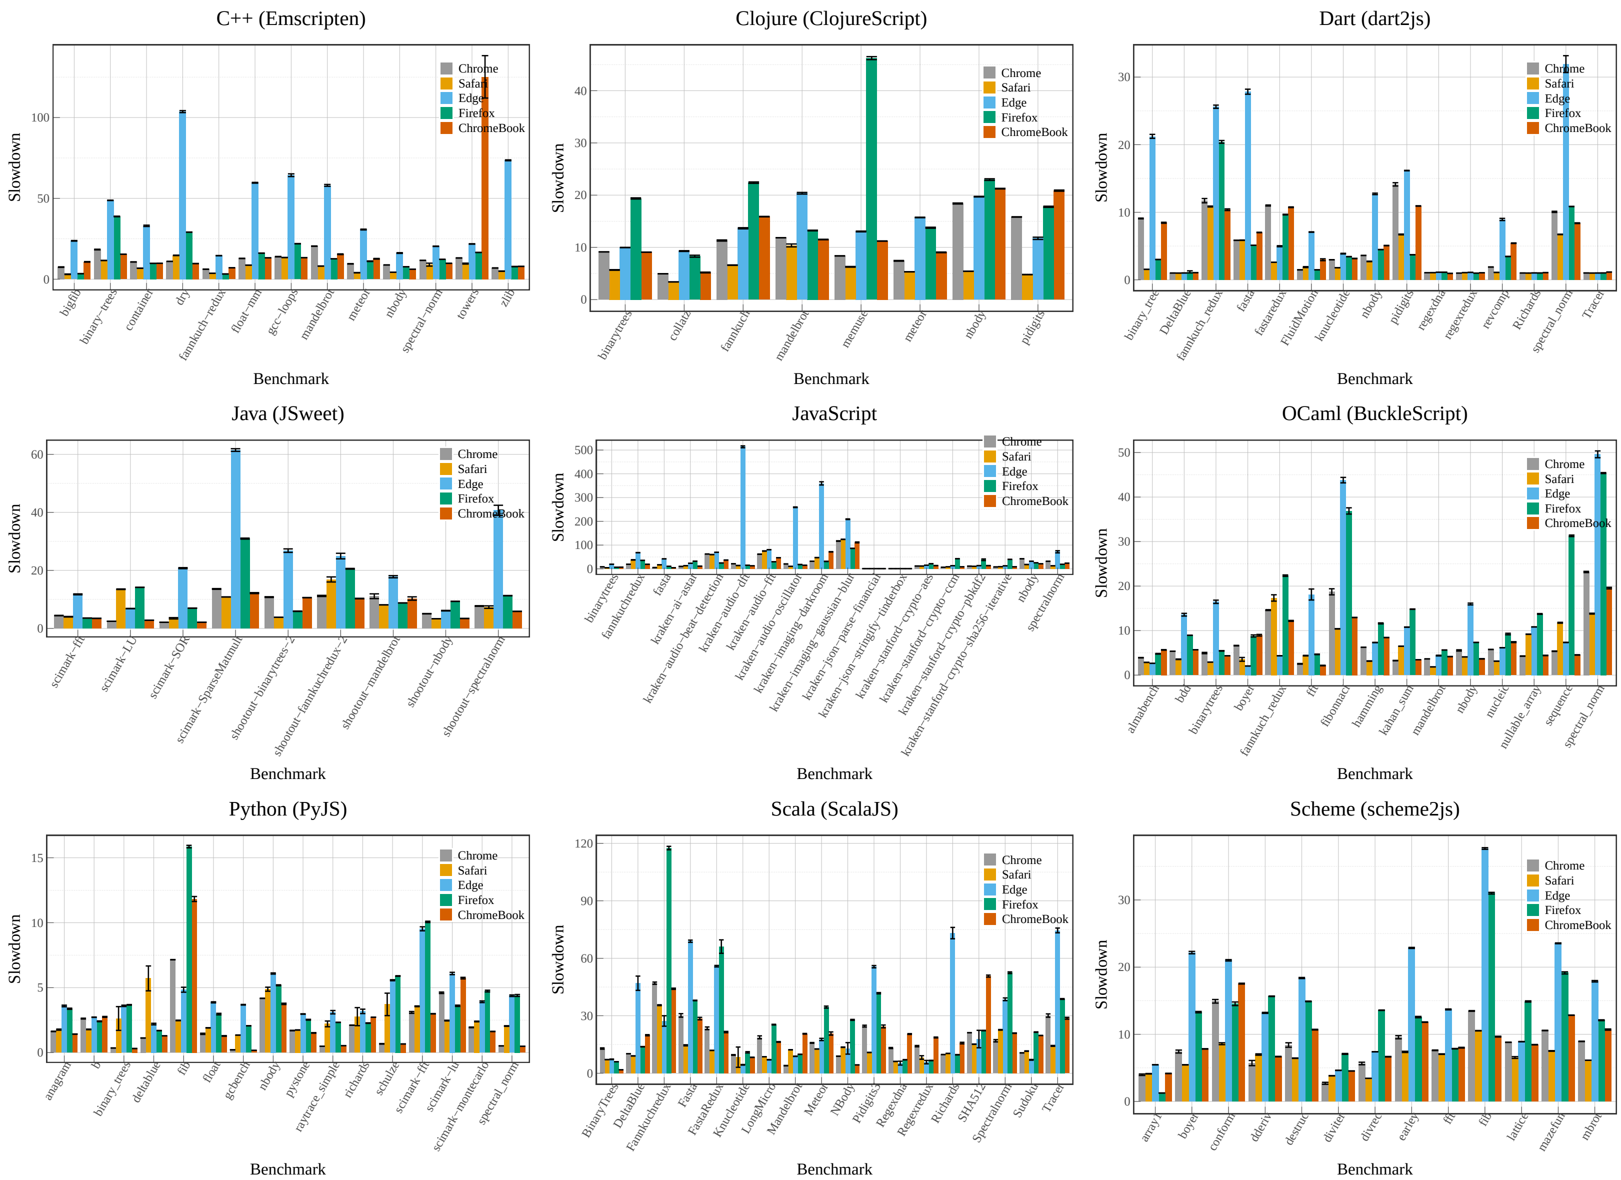

Figure 15. Slowdown incurred by Breakout on each benchmark.