

# Implementing First-Class Continuations by Source to Source Translation

Arjun Guha

University of Massachusetts Amherst

# Continuations Recap

## Definition

The *active expression* in a program is the smallest part of the program that the language will evaluate next.

## Definition

The *continuation* of an active expression is the rest of the program (i.e., excluding the active expression).

Therefore, the continuation is an expression with a single “hole”, where the active expression used to be. We will write that hole as  $\square$ .

- Program: `console.log(2 + 3 * 4); console.log("hi")`
- Active expression: `3 * 4`
- Continuation `console.log(2 +  $\square$ ); console.log("hi")`

## Note

All programming languages have continuations. They are fundamental for evaluation.

## Capturing Continuations Recap

Languages with first-class continuations provide a primitive operator that can turn its own continuation into a value (“capturing the current continuation”).

**control** is a unary operator, which takes a unary function as an argument. When applied to a function  $f$ , **control**( $f$ ):

- 1 Turns its own continuation into a continuation value  $k$ , and
- 2 Applies the function to the continuation value  $f(k)$  in the empty continuation.

Using first-class continuations, we can simulate a synchronous `sleep` function:

```
1 function sleep(n) {
2   function sleeper(k) {
3     setTimeout(k, n);
4   }
5   control(sleeper);
6 }
7
8 console.log("Three");
9 sleep(1000);
10 console.log("Two");
11 sleep(1000);
12 console.log("One");
13 sleep(1000);
14 console.log("Liftoff!");
```

**Today's topic:** How to implement first-class continuations by source-to-source translation.

This material is based on the following paper:

Samuel Baxter, Rachit Nigam, Joe Gibbs Politz, Shriram Krishnamurthi, and Arjun Guha.  
**Putting in All the Stops: Execution Control for JavaScript.** *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, 2018

**Today's meta-topic:** How to conduct research on a full-fledged programming language, without losing your sanity.

# Several Languages Compile to JavaScript

- 1 JavaScript is nearly universal
- 2 Web-based demos are convenient
- 3 Compiling to JavaScript is easy

A few compilers for languages with interesting control operators:



blocking I/O



generators



concurrency

These compilers (and a few others) effectively implement first-class continuations in JavaScript.

Many other compilers drop support for control operators that JavaScript does not directly support.

*There are several ways to implement first-class continuations. Which approach will work best for your programming language?*

A very hard question to answer:

- 1 Cross-language performance evaluation is difficult: many other factors could affect performance
- 2 Changing the implementation of continuations is difficult: tends to affect many parts of the compiler and runtime system
- 3 The answer may be browser-dependent
- 4 The answer may be program-dependent (e.g., is continuation capture rare, or the norm?)

Samuel Baxter, Rachit Nigam, Joe Gibbs Politz, Shriram Krishnamurthi, and Arjun Guha.  
**Putting in All the Stops: Execution Control for JavaScript.** *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, 2018

- 1 Stopify adds first-class continuations to JavaScript. In contrast, prior work supports first-class continuations in other languages that compile to JavaScript.<sup>1</sup>
- 2 Stopify implements some nice features atop continuations:
  - 1 Stopping and resuming the running program
  - 2 Running long computations without locking up the browser
  - 3 Breakpoints and single-stepping (at higher cost)
  - 4 Simulated synchronous I/O and support for web browser events
  - 5 Polyfills for native, higher-order functions that support continuation capture

## This is a common kind of pitch:

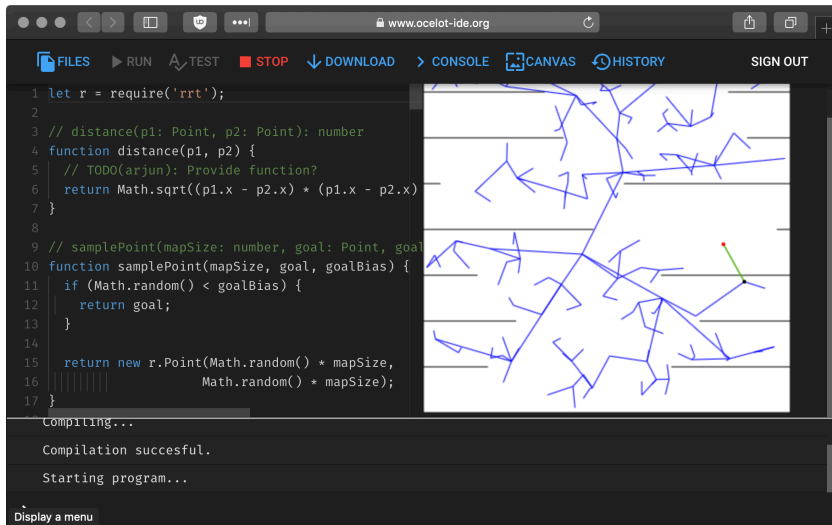
Stopify implements first-class continuations for JavaScript *once and for all*.

- Language-specific solutions are often faster in practice.
- Results suggest that Stopify is competitive or faster than language-specific approaches.
- Stopify uses a few tricks (cheats?) to achieve this result. We will see how later.

---

<sup>1</sup>There are a few earlier implementations, but they are significantly less complete than Stopify. See paper for details.

# Stopify in the Classroom



The screenshot shows the Ocelot IDE interface. The top bar includes navigation icons and the URL [www.ocelot-ide.org](http://www.ocelot-ide.org). The main area is split into a code editor on the left and a canvas on the right. The code editor contains the following JavaScript code:

```
1 let r = require('rrt');
2
3 // distance(p1: Point, p2: Point): number
4 function distance(p1, p2) {
5   // TODO(arjun): Provide function?
6   return Math.sqrt((p1.x - p2.x) * (p1.x - p2.x))
7 }
8
9 // samplePoint(mapSize: number, goal: Point, goalBias: number): Point
10 function samplePoint(mapSize, goal, goalBias) {
11   if (Math.random() < goalBias) {
12     return goal;
13   }
14
15   return new r.Point(Math.random() * mapSize,
16     Math.random() * mapSize);
17 }
```

The canvas on the right displays a visualization of a pathfinding algorithm. It shows a complex, blue, jagged path starting from a red dot and ending at a green dot. The path is composed of many small segments, suggesting a random search process like RRT. The background is white with horizontal lines.

Below the code editor, the console shows the following output:

```
Compiling...
Compilation succesful.
Starting program...
```

At the bottom left, there is a button labeled "Display a menu".

- Used by 200-300 students each semester in a sophomore software engineering class
- Goal is *not* to teach web programming
- Uses Stopify for synchronous I/O, long computations, stop button, etc.

## Stopify Implementation: High-level idea

- The Stopify compiler instruments each function to run in three modes:
  - *Normal mode*: function runs normally
  - *Capture mode*: function saves its own stack frame and returns to its caller, which does the same
  - *Restore mode*: function restores its local variables and jumps to the last statement it was executing
- The Stopify runtime system manages the transitions between each mode
- The **control** operator sets the mode to capture
- Applying a continuation value clears the stack and sets the mode to restore



## Example

```
1 var saved = "nothing yet";
2
3 function save(k) {
4   saved = k;
5 }
6
7 function G(y) {
8   console.log("Enter G");
9   var z = control(save);
10  return y + z;
11 }
12
13 function F(x) {
14   console.log("Enter F");
15   var s = G(x + 10);
16   console.log(s);
17 }
18
19 F(0);
```

- How should we represent the continuation value captured at line 9?
- The continuation value must reconstruct the calls to F and G, without re-printing the output.

# Continuation Representation

```
1 let saved = "nothing yet";
2
3 function save(k) {
4   saved = k;
5 }
6
7 function G(y) {
8   L1: console.log("Enter G");
9   L2: var z = control(save);
10  L3: return y + z;
11 }
12
13 function F(x) {
14   L4: console.log("Enter F");
15   L5: var s = G(x + 10);
16   L6: console.log(s);
17 }
18
19 F(0);
```

```
1 let saved = [
2   { f: F, l: L5, vars: { x: 0 } },
3   { f: G, l: L2, vars: { y: 10 } }
4 ];
```

- In capture mode, each function pushes its own frame into the array
- In restore mode, each function stores its own stack frame and then applies the next function in `kValue`

# Restoring Continuations

Let's pretend we have the `goto` control operator in JavaScript:

```
1 function G(y) {
2   if (mode === 'restore') {
3     y = cont[0].vars.y;
4     var l = cont[0].l;
5     cont.shift();
6     goto l;
7   }
8   L1: console.log("Enter G");
9   L2: let z = control(save);
10  L3: return y + z;
11 }
12
13 function F(x) {
14   if (mode === 'restore') {
15     x = cont[0].vars.x;
16     var l = cont[0].l;
17     cont.shift();
18     goto l;
19   }
20   L4: console.log("Enter F");
21   L5: let s = G(x + 10);
22   L6: console.log(s);
23 }
24
25 F(0);
```

```
1 let saved = [
2   { f: F, l: L5, vars: { x: 0 } },
3   { f: G, l: L2, vars: { y: 10 } }
4 ];
```

- `isRestoring` and `cont` are variables set by the Stopify runtime system
- The `control` operator is implemented as a function

```
1 function control(receiver) {
2   if (mode === 'restore') {
3     mode = 'normal';
4     return /* continuation arg */;
5   }
6   else if (mode === 'normal') {
7     // Start capturing
8   }
9   else {
10    // impossible case
11  }
12 }
```

# Capturing Continuations

- To capture a continuation, the program has to save its stack
- The **control** operator throws an exception that contains an array
- Each function will catch the exception and add itself to the array

```
1 function G(y) {
2   if (mode === 'restore') {
3     ...
4   }
5   L1: console.log("Enter G");
6   try {
7     L2: var z = control(save);
8   }
9   catch (exn) {
10    exn.stack.push({
11      f: G, l: L2,
12      vars: { y: y }
13    });
14    throw exn;
15  }
16  L3: return y + z;
17 }
```

```
18
19 function F(x) {
20   if (mode === 'restore') {
21     ...
22   }
23   L4: console.log("Enter F");
24   try {
25     L5: var s = G(x + 10);
26   }
27   catch (exn) {
28     exn.stack.push({
29       f: F, l: L5,
30       vars: { x: x } });
31     throw exn;
32   }
33 }
```

```
1 let saved = [
2   { f: F, l: L5, vars: { x: 0 } },
3   { f: G, l: L2, vars: { y: 10 } }
4 ];
```

```
1 function control(receiver) {
2   if (mode === 'restore') {
3     ...
4   }
5   else if (mode === 'normal') {
6     throw {
7       stack: [],
8       receiver: receiver
9     };
10  }
11  else {
12    // impossible case
13  }
14 }
```

# Capture/Restore with Arbitrary Expressions

```
1 function P(f, g, x) {  
2   return f(g(x));  
3 }
```

Same approach:

```
1  
2 function P(f, g, x) {  
3   if (mode === 'restore') {  
4     ...  
5   }  
6   try {  
7     L1: return(f(g(x)));  
8   }  
9   catch (exn) {  
10    exn.stack.push(...);  
11    throw exn;  
12  }  
13 }
```

Two serious issues:

- During capture: we cannot (easily) determine if the exception was thrown by `f` or `g`.
- During restore: `goto L1` always applies `g` first, and there is no way to skip it and apply `f`

(Automatically) rewrite `P` to the following equivalent function:

```
1 function P(f,g,x) {  
2   let t = g(x);  
3   return f(t);  
4 }
```

## A Normal Form

Cormac Flanagan, Amr Sabry, Bruce F. Duba, and Matthias Felleisen. The Essence of Compiling with Continuations. PLDI 1993.

My other experiences with ANF:

- Type system design
- Program verification
- Program analysis

# After A Normalization

P after A Normalization:

```
1 function P(f,g,x) {
2   let t = g(x);
3   return f(t);
4 }
```

Correctly instrumented version of P:

```
1 function P(f, g, x) {
2   if (mode === 'restore') {
3     f = cont[0].vars.f;
4     g = cont[0].vars.g;
5     x = cont[0].vars.x;
6     let l = cont[0].l;
7     goto l;
8   }
9   L1: let t;
10  try {
11    L2: t = g(x);
12  }
13  catch (exn) {
14    exn.stack.push({ f: P, l: L2,
15                  vars: { f: f, g: g, x: x } });
16    throw exn;
17  }
18  L3: return f(t);
19 }
```

- Note that P does not capture/restore the application  $f(t)$
- Therefore, when capture starts within  $f$ , the next frame saved in `cont` will be frame of the caller of P
- During restore, the caller of P will restore the  $f$ -frame (skipping the P-frame)
- This is safe, because the P-frame is “useless”: it is merely `return f(t)`
- This lets you fudge *proper tail calls*

# JavaScript A Normalization Pipeline

Doing this correctly was half the work.

Multi-step process:

- 1 Desugar a few JavaScript features (e.g., `var x, y;`)
- 2 Add `try` around loop/if bodies
- 3 Eliminate all loops except **while** loops
- 4 Eliminate **continue**
- 5 Eliminate **switch**
- 6 Eliminate short-circuiting boolean expressions (e.g., `f() || g()`)
- 7 A Normalize

Approach:

- 1 Straightforward assumptions and guarantees between each step
- 2 Some steps are not fundamental, but simplify later stages
- 3 Guiding principle: grammar of ANF-JavaScript
- 4 Only one optimization: do not A-normalize expressions without applications (e.g., `x + 12`)
- 5 Reusable artifact

# Implementing Continuations by Source to Source Translation

- 1 A Normalize JavaScript (7+ steps)
- 2 Box assignable variables that are captured by nested functions and may be captured in continuations

Analogous problem in compilers:

```
1 function F(x) {  
2   function g(y) {  
3     return x + y;  
4   }  
5   return g;  
6 }  
7 let addTen = F(10);  
8 addTen(5); // is x on the stack?
```

- 3 Allow `goto` to function applications
- 4 Add capture/restore code to each function
- 5 One optimization: avoid capturing code when it is obvious that the called function will not capture its continuation ( $\approx 30\%$  performance improvement)



Everything discussed so far is JavaScript-neutral, and is applicable to other languages.

# Arity Mismatch Errors

- The *arity* of a function is the number of arguments it takes.
- An *arity mismatch error* occurs when a function receives the wrong number of arguments. In statically typed languages (e.g., Java), arity mismatch errors occur before the program is run. In dynamically typed languages, arity mismatch errors occur while the program is running.
- JavaScript does not have arity mismatch errors.

```
1 function F(x) {  
2   console.log(x);  
3 }  
4  
5 F(1, 2); // prints 1. The 2 is ignored.  
6 F(); // prints undefined.
```

- This is not a problem for continuations

# The “arguments” Object

All the arguments of a function are available in the `arguments` object.

```
1 function F(x) {  
2   console.log(x);  
3   for (let i = 0; i < arguments.length; i++) {  
4     console.log(arguments[i]);  
5   }  
6 }  
7  
8 F(1,2,3,4); // prints 1 1 2 3 4
```

## Problem

This is a problem: we may lose the extra arguments during capture.

# Implicit Method Calls

## Trivia

Give value for `x`, such that `"Hello " + x` is an infinite loop.  
Hint: The same idea holds in Java too.

```
1 let x = {  
2   toString() {  
3     while (true) { }  
4   }  
5 }
```

## Problem

```
1 let x = {  
2   toString() {  
3     control(...)  
4   }  
5 }
```

- Problem: functions with arity mismatches
- Solution: save and restore the `arguments` object<sup>3</sup>
- Performance problem: `arguments` object materialization
- Problem: implicit method calls,
- Solution: make the implicit calls explicit

```
1 let t1 = x.toString(); // and valueOf  
2 let t2 = "Hello " + t1;
```

- Performance problem: lots of extra instrumentation
- Problem: JavaScript supports getters and setters:

```
1 let a = o.b; // may call the getter b
```

- Solution: make the implicit calls explicit
- Performance problem: lots of extra instrumentation

---

<sup>3</sup>This does not work in the general case.


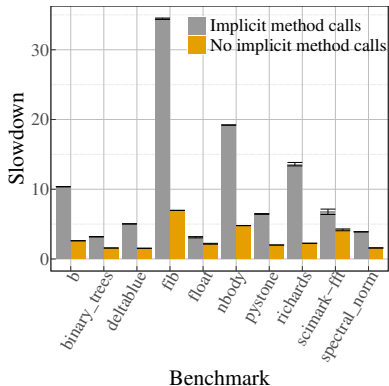
## Solution: Ignore Problems When Possible

A compiler that produces JavaScript emits code into a sub-language of JavaScript.

Compiler	Impl	Args	Getters	Eval	(#) Benchmarks
PyJS	X	M	X	X	(16) PyPy Benchmarks, Shootout
ScalaJS	+	X	X	X	(18) Shootout
scheme2js	X	V	X	X	(13) Larceny
ClojureScript	+	M	X	X	(8) Shootout
dart2js	+	X	T	T	(15) Ton80, Shootout
Emscripten	X	V	X	X	(13) JetStream, Shootout
BuckleScript	X	X	X	X	(15) OPerf-Micro, Shootout
JSweet	+	M	X	X	(9) SciMark, Shootout
JavaScript	✓	✓	✓	✓	(19) Kraken, Shootout
Pyret	X	X	X	T	(21) Pyret

- A ✓ or X indicates that a JavaScript feature is used in full or completely unused. The other symbols indicate restricted variants of the feature.
- Identifying the right sub-language can improve performance dramatically










# Example: Performance of PyJS



Performance relative to unmodified PyJS on a suite of 10 Python benchmarks run 10 times each. Each graph shows how an option setting affects running time or latency. Error bars show the 95% confidence interval.

Speedups are a fluke, due to PyJS generating really bad code.

# Performance Evaluation



CDFs of Stopify's slowdown on nine languages. The median slowdown is in the legend.



- Skulpt: implementation of Python in JavaScript that does execution control
- PyJS: implementation of Python in JavaScript without execution control
- Issues: (1) Skulpt can only run 8 of 16 benchmarks, (2) PyJS and Skulpt pass/fail different portions of the CPython test suite, (3) probably other threats to validity

Benchmark	$\mu$	95% CI
anagram	0.25	$\pm 0.01$
binary-trees	0.27	$\pm 0.01$
fib	0.25	$\pm 0.00$
gcbench	0.08	$\pm 0.01$
nbody	0.25	$\pm 0.00$
pystone	0.37	$\pm 0.01$
schulze	1.25	$\pm 0.08$
spectral-norm	0.36	$\pm 0.01$

Slowdown relative to Skulpt. (Stopify is faster when  $\mu < 1$ .)

# Stopify + Pyret

- Joe Politz, one of Pyret's lead developers, co-authored this work
- Pyret: mostly-functional programming language that compiles to JS, self-hosting compiler, does proper tail calls, blocking I/O, REPL, animations, graceful termination, etc.
- Five years of engineering, thousands of users, still has some issues
- Our approach: modify the last phase of the Pyret compiler and portions of the runtime system to invoke Stopify, instead of doing capture/restore itself (lots of code deleted)<sup>4</sup>

From Pyret's runtime system (several bug fixes over the years):

```
1 function eachLoop(fun, start, stop) {
2   var i = start;
3   function restart(_) {
4     var res = thisRuntime.nothing;
5     if (--thisRuntime.GAS <= 0) { res = thisRuntime.makeCont(); }
6     while(!thisRuntime.isContinuation(res)) {
7       if (--thisRuntime.RUNGAS <= 0) { res = thisRuntime.makeCont(); }
8       else {
9         if(i >= stop) {
10          ++thisRuntime.GAS;
11          return thisRuntime.nothing;
12        } else {
13          res = fun.app(i);
14          i = i + 1;
15        } } }
16   res.stack[thisRuntime.EXN_STACKHEIGHT++] =
17     thisRuntime.makeActivationRecord("eachLoop", restart, true, [], []);
18   return res;
19 }
20 return restart();
21 }
```


---

<sup>4</sup>This is the right way to use Stopify.

# Stopify + Pyret

After Stopify (which generates equivalent instrumentation):

```
1 function eachLoop(fun, start, stop) {  
2   for (var i = start; i < stop; i++) { fun.app(i); }  
3   return thisRuntime.nothing;  
4 }
```



- Use existing infrastructure when possible (in this case, Babel)
- When working on a full-scale language, identify a small fragment in which the essence of the problem and solution can be formulated
- The translation may be problem-dependent (e.g., type-checking vs translation)
- ANF is likely to help you
- Take testing and CI seriously