

First-Class Continuations: What and Why

Arjun Guha

University of Massachusetts Amherst

Plan

Today's topic: An introduction to *first-class continuations*

Tomorrow's topic: How to implement first-class continuations by source-to-source translation

Today's meta-topic: The *expressiveness* of programming language features

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

We will use basic JavaScript for all code samples. Even if you don't know the language, it should be straightforward to follow. Please ask questions if you have trouble with the syntax.

Please interrupt and ask questions!

We are going to see a few code examples that are very challenging to understand.

Ask questions of the form, "*What if we change the example to do X instead?*" We will change the example and run it to see how it behaves.

Control Operators

Definition

A *control operator* is a programming language construct that changes the normal flow of execution in a program.

First-class continuations are a kind of control operator.

Some other control operators that are more familiar:

- *Conditionals*: **if** and **switch**
- *Loops*: **for** and **while**
- *Structured jumps*: **break** and **continue**
- *Exceptions*: **try catch** and **try finally**

Recent versions of JavaScript support some additional control operators:

- *Asynchronous functions*: **async function** and **await**
- *Generator functions*: **function***, **yield**, and **yield***

None of this is JavaScript specific: you can find all the control operators above in several other programming languages.

Why New Control Operators?

- The right control operator can make programs much easier to read and write. We will show how, using JavaScript's asynchronous functions and generator functions.
- First-class continuations make it possible to build several control operators as a library (i.e., without building them into the language).

Definition

A *synchronous* or *blocking* I/O operation blocks execution while performing I/O and then resumes execution with the result of the I/O operation.

For example, a function with the following type would synchronously download an image:

```
1 loadImage(url: string) => Image
```

We could then display the image:

```
1 // drawImage(image: Image) => void
2 function drawImage(image) {
3   document.body.appendChild(image);
4 }
5
6 let image = loadImage("example.jpg");
7 drawImage(image);
```

In JavaScript, (almost) all I/O is asynchronous. It is impossible to write a synchronous `loadImage` function in JavaScript.

Definition

Asynchronous or *nonblocking* I/O performs the operation in the background and applies *callback function* when the I/O result is available.

We can implement an asynchronous function to load images:

```
1 loadImage(url: string, callback: Image => void) => void
```

So, to load two images in sequence ([link](#)):

```
1 function callback2(image2) {  
2   drawImage(image);  
3 }  
4  
5 function callback1(image1) {  
6   drawImage(image);  
7   loadImage("example2.jpg", callback2);  
8 }  
9  
10 loadImage("example1.jpg", callback1);
```

We need a new callback function to wait for the result of each new asynchronous operation.

Definition

An *async function* can start an asynchronous task, suspend its own execution, and then resume with the result of the task when the task completes.

([link](#))

```
1 function loadImage(url) {  
2   // implementation omitted. Turns callbacks into Promises.  
3 }  
4  
5 async function F() {  
6   let image1 = await loadImage("example1.jpg");  
7   drawImage(image1);  
8   let image2 = await loadImage("example2.jpg");  
9   drawImage(image2);  
10 }
```

- Within an async function, we use the **await** keyword to call another async function.
- We can read the program from top to bottom (contrast to the callback-based approach).

Definition

A *generator function* is a special kind of function that suspends its execution when it produces a value for the caller. The caller may then resume the generator function to make it produce the next value (if any).

Ordinary functions *do not* work this way: they run to completion and cannot be suspended.

```
1 function* makeThreeGen() {  
2   yield 1;  
3   yield 2;  
4   yield 3;  
5 }
```

```
1 let gen = makeThreeGen();  
2 console.log(gen.next().value); // displays 1  
3 console.log(gen.next().value); // displays 2
```


Unbounded Generators

The following example helps illustrate that the **yield** statement truly suspends execution of the generator function:

```
1 function* genNats() {
2   let i = 0;
3   while (true) {
4     yield i;
5     i = i + 1;
6   }
7 }
8
9 let gen = genNats();
10
11 // Displays 0, 1, 2, 3, 4, 5, 6, 7, 8, 9
12 for (let i = 0; i < 10; i++) {
13   console.log(gen.next().value);
14 }
```

If **yield** did not suspend the generator function, the loop in the generator would run forever.

Some More Control Operators

There are several other kinds of control operators that JavaScript does not have:

- *Sampling executions from probability distributions*: [WebPPL](#) implements a compiler from a small subset of JavaScript with extensions to support probabilistic programming.
- *Backtracking search*: [Tau Prolog](#) implements a Prolog interpreter in JavaScript.
- *(Green) Threads*: [Concurrent JavaScript](#)

Question

Should JavaScript implement these natively too?

Using first-class continuations, we can implement each control operator above as a small library. In essence, first-class continuations subsume a wide variety of other control operators and language features.

We need to understand three concepts before we get to first-class continuations:

- 1 Environments
- 2 Active Expressions
- 3 Continuations

This will be a quick, informal introduction via examples. For a rigorous, formal approach, see [Semantics Engineering with PLT Redex](#).

Definition

The *environment* of a program maps variable names to values.

```
1 let x = 100;  
2 let y = 200;  
3 console.log(x + y);
```

At the last line, the environment holds the values of x and y .

```
1 function F(x) {  
2   console.log(x);  
3 }  
4  
5 F(100);  
6 F(200);
```

Each application evaluates the body of F in a different environment.

Definition

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

Examples:

```
1 console.log(2 + 3 * 4);
```

```
1 console.log(2 + 12);
```

```
1 console.log(14);
```

Two things to note:

- A statement or the whole program can be the active expression.
- Values cannot be active expressions.

```
1 if (5 > 10) {  
2   console.log("oops");  
3 }  
4 else {  
5   console.log("hi");  
6 }
```

```
1 if (false) {  
2   console.log("oops");  
3 }  
4 else {  
5   console.log("hi");  
6 }
```

```
1 console.log("hi");
```

Continuations

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 .

For example:

Program	Active Expression	Continuation
<code>console.log(2 + 3 * 4)</code>	<code>3 * 4</code>	<code>console.log(2 + \square)</code>
<code>console.log(2 + 12)</code>	<code>2 + 12</code>	<code>console.log(\square)</code>
<code>console.log(14)</code>	<code>console.log(14)</code>	\square

Note

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

To run a program, we:

- 1 Identify the active expression and the continuation,
- 2 Evaluate the active expression,
- 3 Plug the result of evaluation into the hole that the active expression,
- 4 Repeat, until the program does not have an active expression.

Another Example of Continuations

Step 1	Program Active Expression Continuation	<code>if (10 > 5) { log("o" + "k"); } else { log("oops"); }</code> <code>10 > 5</code> <code>if (□) { log("o" + "k"); } else { log("oops"); }</code>
Step 2	Program Active Expression Continuation	<code>if (true) { log("o" + "k"); } else { log("oops"); }</code> <code>if (true) { log("o" + "k"); } else { log("oops"); }</code> <code>□</code>
Step 3	Program Active Expression Continuation	<code>log("o" + "k");</code> <code>"o" + "k"</code> <code>log(□);</code>
Step 4	Program Active Expression Continuation	<code>log("ok");</code> <code>log("ok");</code> <code>□</code>

First-Class Continuations

Definition

In programming languages that support *first-class continuations*, it is possible to turn a continuation into a value.

A continuation value stored in a variable:

```
1 let k = (□ + 100);
```

An array of continuation values:

```
1 let arr = [□ + 100, □ + 200];
```

A continuation passed as an argument to a function:

```
1 f(□ + 100);
```

Note

The □ notation is pseudocode. It is not possible to directly write a continuation in a program, but it helps understand what is happening under the hood. We will resolve this in a moment.

Applying a Continuation Value

A continuation value can be applied to an argument. (Applying a continuation value looks like applying a unary function.) When a continuation k is applied to a value v :

- 1 We discard the current continuation, and
- 2 The continuation k is restored and the hole gets filled with v .

Step 1	Program Environment Active Expression Continuation	<code>let k = (\square + 1); log(100 + k(10))</code>
Step 2	Program Environment Active Expression Continuation	<code>log(100 + k(10));</code> <code>k = (\square + 1)</code> <code>k(10)</code> <code>log(100 + \square);</code>
Step 3	Program Environment Active Expression Continuation	<code>10 + 1; <i>This is k with 10 plugged in.</i></code> <code>k = (\square + 1)</code> <code>10 + 1</code> <code>\square <i>Continuation from Step 2 has been discarded</i></code>
Step 4	Result Environment	<code>11;</code> <code>k = (\square + 1)</code>

Note

If you have seen continuations before, you may know about `call/cc`. We are going to introduce an operator that is similar, but not identical to `call/cc`.

Capturing Continuations

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

We are going to introduce a new unary operator called **control**, 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.

The following example is fairly simple, since F does not use k ([link](#)):

Step 1	Program Environment Active Expression Continuation	function $F(k)$ { $\log(200)$; }; $\log(100 + \mathbf{control}(F))$;
Step 2	Program Environment Active Expression Continuation	$\log(100 + \mathbf{control}(F))$; function $F(k)$ { $\log(200)$; }; control (F) $\log(100 + \square)$;
Step 3	Program Environment Active Expression Continuation	$\log(200)$; let $k = \log(100 + \square)$; function $F(k)$ { $\log(200)$; }; $\log(200)$; \square

Applying a Captured Continuation (Example 1)

In the following example, the **throw** statement is never reached ([link](#)):

```
1 function F(k) {  
2   k(200);  
3   throw "bad";  
4 }  
5 log(100 + control(F));
```

Step 1	Program Environment Active Expression Continuation	function F(k) { k(200); throw "bad"; }; log(100 + control (F)); function F(k) { k(200); throw "bad"; }; □; log(100 + control (F));
Step 2	Program Environment Active Expression Continuation	log(100 + control (F)); function F(k) { k(200); throw "bad"; }; control (F) log(100 + □);
Step 3	Program Environment Active Expression Continuation	k(200); throw "bad" let k = log(100 + □); k(200) □; throw "bad";
Step 4	Program Environment Active Expression Continuation	log(100 + 200); 100 + 200 log(□);

...

Applying Captured Continuations (Example 2)

Trace the execution of this program:

```
1 let i = 0;
2 let saved = "nothing";
3 function handler(k) {
4   saved = k;
5   saved("Start");
6 }
7
8 log(control(handler));
9 if (i < 3) {
10  i = i + 1;
11  saved(i);
12 }
```

Building New Control Operators with First-Class Continuations

- Code that uses first-class continuations directly is usually very hard to read.
- But, we can use first-class continuations to build other control operators that are more straightforward.

A Countdown Program

Programming Challenge

Write a program that counts down seconds, displaying: "Three", "Two", "One", "Liftoff!".

Solution ([link](#)):

```
1 function callback3() {
2   console.log("Liftoff!");
3 }
4
5 function callback2() {
6   console.log("One");
7   setTimeout(callback3, 1000);
8 }
9
10 function callback1() {
11   console.log("Two");
12   setTimeout(callback2, 1000);
13 }
14
15 console.log("Three");
16 setTimeout(callback1, 1000);
```

WTF

We cannot write a blocking sleep function in JavaScript (without busy-waiting).

Application: The Sleep Operator

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 }
```

The countdown program, refactored:

```
1 console.log("Three");  
2 sleep(1000);  
3 console.log("Two");  
4 sleep(1000);  
5 console.log("One");  
6 sleep(1000);  
7 console.log("Liftoff!");
```

We can take any asynchronous function and simulate synchronous execution following this recipe. This is effectively what **`async function`** and **`await`** do.

Application: Cooperative Threads

Definition

Cooperative threads run several logical threads on a single physical thread, thus there is no parallelism and only one thread is running at a time (other threads are suspended in background). Moreover, the running thread has to explicitly yield control of the physical thread for another thread to start running.

Key ideas in the implementation:

- 1 A global array of continuation values, where each continuation value is a suspended thread.
- 2 To yield, we capture the continuation of the active thread, add it to the array, and apply one of the other continuations.

```
1 let threads = [ ];
2
3 function createThread(f) {
4   function threadFunc() {
5     f();
6     start();
7   }
8   threads.push(threadFunc);
9   yieldThread();
10 }
```

```
1 function start() {
2   if (threads.length > 0) {
3     let nextThread = threads.shift();
4     nextThread();
5   }
6 }
7
8 function yieldThread() {
9   function switcher(k) {
10    threads.push(k);
11    let kOther = threads.shift();
12    kOther("resumed");
13  };
14
15  control(switcher);
16 }
```

Using Cooperative Threads

([link](#))

```
1 function threadA() {
2   for (let i = 0; i < 10; i++) {
3     yieldThread();
4     console.log(i);
5   }
6 })
7
8 function threadB() {
9   for (let i = 100; i < 110; i++) {
10    yieldThread();
11    console.log(i);
12  };
13 })
14
15 createThread(threadA);
16 createThread(threadB);
17 start(); // needed to activate other threads
```

- I apologize if your head hurts
- Control operators can make certain kinds of programs much easier to write (e.g., generator functions and async functions)
- First-class continuations are a powerful primitive for building new control operators
- All examples online (including some that I have not covered)
- Tomorrow: How to implement continuations by source-to-source translation