

CPSC 311: Lazy evaluation (DRAFT) ("lec-lazy")

Joshua Dunfield

University of British Columbia

November 21, 2016

1 Evaluation strategies: review and update

1.1 Review

Earlier in the course, we came up with two strategies for evaluating function application ($\text{App } e1 \ e2$): the *expression strategy*, in which the function argument $e2$ is substituted for x in the body of ($\text{Lam } x \ eB$), and the *value strategy*, in which the function argument $e2$ is evaluated immediately and the resulting value $v2$ is substituted for x in eB .

Under the expression strategy, we evaluate $e2$ as many times as $(\text{Id } x)$ is evaluated; under the value strategy, we evaluate $e2$ exactly once. The value strategy is usually more efficient than the expression strategy, but the expression strategy is more efficient if $(\text{Id } x)$ is not evaluated. For example, in

$$(\text{App } (\text{Lam } x \ (\text{Num } 0)) \ e2)$$

there is no $(\text{Id } x)$ in the body of the Lam , so (under the expression strategy) the argument $e2$ will never be evaluated.

■ **Exercise 1.** Give a slightly larger (and hopefully less contrived) example of a function application for which the expression strategy is more efficient than the value strategy. (Hint: apply a function of *two* arguments, that is, $(\text{Lam } x \ (\text{Lam } y \ \dots))$.)

1.2 Update for environment-based evaluation

We have mostly used the value strategy, and we stayed with that strategy in developing the environment-based evaluation rule for App , Env-app ; for clarity, we'll now call that rule Env-app-value :

$$\frac{\text{env} \vdash e1 \Downarrow (\text{Clo } \text{env}_{\text{old}} \ (\text{Lam } x \ eB)) \quad \text{env} \vdash e2 \Downarrow v2 \quad x=v2, \text{env}_{\text{old}} \vdash eB \Downarrow v}{\text{env} \vdash (\text{App } e1 \ e2) \Downarrow v} \text{Env-app-value}$$

To facilitate experimenting, I'm leaving the value-strategy application App alone, but adding an expression-strategy application App-expr .

To implement the expression strategy, we can use

$$\frac{\text{env} \vdash e1 \Downarrow (\text{Clo } \text{env}_{\text{old}} \ (\text{Lam } x \ eB)) \quad x=(\text{Clo } \text{env } e2), \text{env}_{\text{old}} \vdash eB \Downarrow v}{\text{env} \vdash (\text{App-expr } e1 \ e2) \Downarrow v} \text{Env-app-expr}$$

§1 Evaluation strategies: review and update

As we did for Lam, we are using a closure to save the current environment so that when `(ld x)` is evaluated, we can evaluate `e2` under `env` rather than some later environment.

For historical reasons, a closure that is used in this way is called a *thunk*.

The rule Env-app-expr (for App-expr) is implemented in `env-lazy.rkt`.

Examples:

```
; for app-expr, do we evaluate the argument twice?
(unparse (interp (parse '{App-expr {Lam x {+ x x}} {+ 1 2}})))

; does it work when the argument is a Lam?
(unparse (interp (parse '{App-expr {Lam f {App f 5}} {Lam y {+ y y}}}))

; does it work when the argument has a free variable (z)?
(unparse (interp (parse '{Let z 100
                        {App-expr {Lam f {App f 5}}
                        {Lam y {+ y z}}})))))

; are we still doing lexical scope?
(unparse (interp (parse '{Let z 100
                        {App-expr {Lam f {Let z 444 {App f 5}}
                        {Lam y {+ y z}}}))}))
```

2 Lazy evaluation

Environments make it possible to implement a third evaluation strategy, which I think is better than the expression strategy. Whether it's better than the value strategy is unclear.

It's useful to pause and relate my terminology in this course to terminology that you may come across elsewhere. I invented my own terminology because I think it's less confusing, but you should be aware of the more common usage:

invention	CPSC 311 name	“formal” names	“popular” names	other names
1950s	value strategy	call-by-value, CBV	eager evaluation, strict evaluation	[applicative order]
1960	expression strategy	call-by-name, CBN	by name	[normal order]
1971–76	lazy evaluation	call-by-need	lazy evaluation	

Brackets, e.g. “[applicative order]”, indicate that there is no consensus that those terms are being used accurately, and that many people will (probably correctly) object and reserve those terms for different concepts. I mention them because you may come across them, and they aren't *entirely* wrong: applicative order is *more like* the value strategy than it is like the expression strategy.

Also, people often say “evaluation order” rather than “evaluation strategy”. But I prefer “strategy”, at least in 311, because it's not just the *order* in which expressions are evaluated: it's also whether they're evaluated at all.

2.1 Overview

From a distance, lazy evaluation looks like the expression strategy:

- function arguments are not evaluated immediately;
- function arguments are only evaluated when used.

The difference from the expression strategy is in what happens when the argument is evaluated, *if* it is evaluated:

- The expression strategy evaluates the argument, but doesn't remember the result. So if it sees $(\text{Id } x)$ again, it evaluates the argument again.
- Lazy evaluation remembers the result of evaluating the argument. If it sees $(\text{Id } x)$ a second (or third...) time, it returns the result without evaluating it again.

We can't use the environment to remember the result, however, because the environment is only passed *up* the evaluation derivation, not "threaded through". But earlier, we added support for mutable references (boxes), which live in a store that *is* threaded through! So we can use the store to remember our results.

(By the way, this is roughly how lazy evaluation actually works in languages that use it, like Haskell.)

2.2 Rules

We want to put the result of evaluating an argument in the store; since the store is what survives leaving a lexical scope, we also need to remember in the store *whether* we have evaluated that argument. The environment will *refer* to the store, using a location.

I left out the store from the above rules, so let's put that in `Env-app-expr` and then rewrite `Env-app-expr` to be lazy.

$$\frac{\text{env}; \mathbf{S} \vdash e1 \Downarrow (\text{Clo env}_{\text{old}} (\text{Lam } x \text{ eB})); \mathbf{S1} \quad x = (\text{Clo env } e2), \text{env}_{\text{old}}; \mathbf{S1} \vdash eB \Downarrow v; \mathbf{S2}}{\text{env}; \mathbf{S} \vdash (\text{App-expr } e1 \text{ } e2) \Downarrow v; \mathbf{S2}} \text{Env-app-expr}$$

I'll use "Lazy-thk" for the thunk we're creating.

$$\frac{\text{env}; \mathbf{S} \vdash e1 \Downarrow (\text{Clo env}_{\text{old}} (\text{Lam } x \text{ eB})); \mathbf{S1} \quad \ell \text{ fresh for } \mathbf{S1} \quad x = (\text{Lazy-ptr } \ell), \text{env}_{\text{old}}; \ell \triangleright (\text{Lazy-thk env } e2), \mathbf{S1} \vdash eB \Downarrow v; \mathbf{S2}}{\text{env}; \mathbf{S} \vdash (\text{App-lazy } e1 \text{ } e2) \Downarrow v; \mathbf{S2}} \text{Env-app-lazy}$$

In the second premise of `Env-app-lazy`:

1. We bind x to $(\text{Lazy-ptr } \ell)$. Since the environment isn't mutable, x will *always* be bound to $(\text{Lazy-ptr } \ell)$ for the entire time that x is in scope.
2. We extend the store with a new location ℓ containing $(\text{Lazy-thk env } e2)$.

§2 Lazy evaluation

Note that we haven't evaluated e_2 yet—and if evaluating e_B does not evaluate $(\text{Id } x)$, we never will.

When we evaluate $(\text{Id } x)$, we will look it up (Env-id) and evaluate what we find. The environment (not the store!) has $x=(\text{Lazy-ptr } \ell)$, so we find $(\text{Lazy-ptr } \ell)$ and evaluate that.

We need a rule for the case where we haven't evaluated the argument yet. In that case, looking up ℓ in the store will give a Lazy-thk . We evaluate e_2 under the environment that Env-app-lazy saved, resulting in a value v , and use update-loc to replace ℓ with v :

$$\frac{\text{lookup-loc}(S, \ell) = (\text{Lazy-thk env}_{\text{arg}} e_2) \quad \text{env}_{\text{arg}}; S \vdash e_2 \Downarrow v; S1 \quad \text{update-loc}(S1, \ell, v) = S2}{\text{env}; S \vdash (\text{Lazy-ptr } \ell) \Downarrow v; S2} \text{Env-lazy-ptr}$$

We also need a rule for the case where we already applied Env-lazy-ptr to ℓ . In that case, looking up ℓ in the store will *not* give a Lazy-thk , but some value—the v that we got while applying Env-lazy-ptr .

$$\frac{\text{lookup-loc}(S, \ell) = v \quad v \neq (\text{Lazy-thk } \dots \dots)}{\text{env}; S \vdash (\text{Lazy-ptr } \ell) \Downarrow v; S} \text{Env-lazy-ptr-done}$$

```
(unparse (interp (parse '{App-lazy {Lam x {+ x x}} {- 10 1}})))
```

2.3 Ideology

Evaluation strategy, like typing, is one of the most enduring controversies in programming language design. The controversy began the moment there was more than one evaluation order:

The first call-by-name language, Algol 60, also supported call-by-value. It seems that call-by-value was the language committee's preferred default, but Peter Naur, the editor of the Algol 60 report, independently reversed that decision—which he said was merely one of a “few matters of detail”. A committee member, F.L. Bauer, said this showed that Naur “had absorbed the Holy Ghost after the Paris meeting. . . there was nothing one could do. . . it was to be swallowed for the sake of loyalty.” (From Dunfield (2015), “Elaborating evaluation-order polymorphism”; quotations from Wexelblat (1981), *History of Programming Languages I.*)

(There was also some argument about whether Naur had independently decided to include recursion in Algol-60, but my reading is that he didn't do that—the committee had agreed to support recursion, but may have had arguments over details.)

Later developments continued to be full of ideology. Lazy evaluation (under the name call-by-need) was introduced in a 1971 PhD thesis, and first implemented (twice, mostly independently, I believe) in 1976. One of the 1976 papers has the rather opinionated title “CONS Should Not Evaluate its Arguments”, only softened slightly in the paper itself: “we have uncovered a critical class of elementary functions which probably should never be treated as strict: the functions which allocate or *construct* data structures.” I suspect that not all ML programmers would agree. Nor would users (at least, designers) of Lisp, the language used in the 1976 paper: neither Lisp nor its descendants Scheme and Racket are lazy.

These controversies are (partly) grounded in legitimate disagreements about design tradeoffs between the value strategy (eager evaluation) and lazy evaluation:

- The value strategy trades speed in one, possibly uncommon case—the case where the argument isn't used—for simplicity.

§2 Lazy evaluation

- Lazy evaluation trades simplicity for speed, but also trades space for speed: the many thunks that have to be created take up space (and slow down garbage collection) even if they are never evaluated. Reasoning about how much space is used is difficult; for example, I believe that the Haskell community relies on *space profilers* to debug “space leaks” that result from thunks being built that are never needed.
- Lazy evaluation sometimes trades actual improvement for apparent improvement: if the expression whose evaluation is being avoided is simple, it would be faster to evaluate it without creating a thunk—even if the argument is never used.

There are certainly cases where lazy evaluation is superior, but opponents of lazy evaluation argue that these cases can be handled by explicit programmer-controlled laziness instead. That is, laziness should be an option that must be asked for explicitly, rather than the default.

2.4 Function application vs. the whole language

The rules we’ve developed add lazy function application, without changing any other language constructs. Languages with built-in laziness usually don’t stop there. For example, Haskell is entirely lazy: if you add two expressions with `+`, the addition won’t be performed unless the result is “demanded” (such as by printing the value to the user).