Evaluation and convergence in the computational calculus

Claudia Faggian, Giulio Guerrieri, Riccardo Treglia

Evaluation and convergence in the computational calculus

Claudia Faggian[*]   Giulio Guerrieri[†]   Riccardo Treglia[‡]

Abstract

In Moggi's computational calculus, reduction is the contextual closure of the rules obtained by orienting three monadic laws. In the literature, evaluation is usually defined as the closure under weak contexts (no reduction under binders): $\mathsf{E} = \langle\rangle \mid \mathsf{let}\, x := \mathsf{E}\, \mathsf{in}\, M$.

We show that, when considering all the monadic rules, weak reduction is non-deterministic, non-confluent, and normal forms are not unique. However, when interested in returning a value (convergence), the only necessary monadic rule is β, whose evaluation is deterministic.

The *computational λ-calculus*, noted λ_c, was introduced by Moggi [Mog88, Mog89, Mog91] as a meta-language to describe computational effects in programming languages. Since then, computational λ-calculi have been developed as foundations of programming languages, formalizing both functional and effectful features [WT03, BHM02, PP03, LPT03, DGL17], in a still active line of research.

To model effectful features at a semantic level, Moggi used the categorical notion of *monad*. A monad can be equivalently presented as a Kleisli triple satisfying three identities [Mog91, Mac97]. At an operational level, Moggi [Mog88] internalized these identities into the syntax of λ_c, giving rise to three conversion rules—called *monadic laws*—that are added to the usual β and η rules.

Nowadays the literature is rich of computational calculi that refine Moggi's λ_c. Such calculi are presented in at least three different fashions: fully equational systems [LPT03, PP02] (all conversion rules are unoriented identities); hybrid systems where β (and η, if considered) are oriented rules while the monadic laws are identities on terms [DGL17]; reduction systems where every rule is oriented [SW97]. Here we follow the latter approach, which brings to the fore operational aspects of reduction and evaluation which seem to have been neglected in the literature.

Indeed, in the literature of calculi with effects [LPT03, DGL17], *evaluation* is usually *weak*, that is, it is not allowed in the scope of the binders (λ or let). This is the way evaluation is implemented by functional programming languages such as Haskell and OCaml. Moreover, only β and $\mathsf{let}.\beta$ are considered. However,

---

[*]Université de Paris, IRIF, CNRS, F-75013 Paris, France

[†]University of Bath, Department of Computer Science, Bath, UK

[‡]Università di Torino, Department of Computer Science, Turin, Italy

in Moggi's λ_c and in Sabry and Wadler's [SW96], the reduction is full, that is, reduction is the compatible closure of all the monadic rules. When considering *all the rules*, we observe—quite unexpectedly—that evaluation (*i.e.* weak reduction) is *non-deterministic, non-confluent, and normal forms are not unique*.

Reduction and Evaluation. Here we focus on a computational λ-calculus which is standard in the literature, namely Sabry and Wadler's λ_{ml^*} [SW96]. This is a neat and compact refinement of Moggi's untyped λ_c [Mog88]—the relation between the two calculi is formalized by a reflection [SW96], which guarantees that there exists a kernel of λ_c which is isomorphic to λ_{ml^*}.

λ_{ml^*}—which we display in Figure 1— has a two sorted syntax that separates *values* (*i.e.* variables and abstractions) and *computations*. The latter are either *let*-expressions (aka explicit substitutions, capturing monadic binding), or applications (of values to values), or coercions $[V]$ of values V into computations (corresponding to the `return` operator in Haskell).

$$
\begin{array}{llll}
\textit{Values:} & V, W & ::= & x \mid \lambda x.M \\
\textit{Computations:} & M, N & ::= & [V] \mid \text{let}\, x := M \,\text{in}\, N \mid VW
\end{array}
$$

Reduction rules:

$$
\begin{array}{rrcll}
(\beta) & (\lambda x.M)V & \mapsto_\beta & M[V/x] & \\
(\eta) & \lambda x.Vx & \mapsto_\eta & V & x \notin \text{fv}(V) \\
(\text{let}.\beta) & \text{let}\, x := [V] \,\text{in}\, N & \mapsto_{\text{let}.\beta} & N[V/x] & \\
(\text{let}.\eta) & \text{let}\, x := M \,\text{in}\, [x] & \mapsto_{\text{let}.\eta} & M & \\
(\text{let}.\textit{ass}) & \text{let}\, y := (\text{let}\, x := L \,\text{in}\, M) \,\text{in}\, N & \mapsto_{\text{let}.\textit{ass}} & \text{let}\, x := L \,\text{in}\, (\text{let}\, y := M \,\text{in}\, N) & x \notin \text{fv}(N)
\end{array}
$$

Figure 1: λ_{ml^*}: Syntax and Reduction

- The *reduction rules* for λ_{ml^*} (in Figure 1) are the usual β (and η) rules from Plotkin's *call-by-value* λ-calculus [Plo75], plus the *oriented* version of the three *monadic laws*: $\text{let}.\beta$, $\text{let}.\eta$, $\text{let}.\textit{ass}$. The last two, in particular, correspond to the monadic laws of *identity* and *associativity*, respectively.

- *Reduction* $\rightarrow$ is the *contextual closure* of the reduction rules.

Following standard practice, we define *evaluation* $\xrightarrow{}_{\text{w}}{}_{ml^*}$ (aka *sequencing*) as the closure of the rules under *evaluation context* E:

$$
\text{E} ::= \langle\,\rangle \mid \text{let}\, x := \text{E} \,\text{in}\, N \qquad \text{evaluation context}
$$

Informally, the operational understanding of weak reduction is that evaluating $\text{let}\, x := M \,\text{in}\, N$ amounts to first evaluate M until it returns a value, that is, until a computation of the form $[V]$ is reached. Then V is passed to N by substituting V

2

for x in N, thanks to the rule let.β.

Despite the prominent role that weak reduction has in the literature of calculi with effects, its reduction properties are somehow surprising. While full reduction $\rightarrow_{ml^*}$ is confluent, the closure of the rules under *evaluation context* turns out to be *non-deterministic*, *non-confluent*, and its *normal forms* are *not unique*.

Note that such issues only come from the monadic rules let.η and let.*ass* (*identity* and *associativity*), not from β or let.β. It is worth to clarify that while the literature on computational λ-calculi often adopts weak reduction (see for instance, [LPT03, DGL17], where a big-step variant is used), the rules let.*ass* and let.η are usually dealt with as *unoriented* identities—the only oriented rules being β and let.β.

(Non-)Confluence. In λ_{ml^*}, the reduction $\rightarrow_{ml^*}$ is confluent, but weak reduction $\underset{w}{\rightarrow}_{ml^*}$ is not. We now give some examples. For every $\gamma \in \{\beta, \eta, \text{let}.\beta, \text{let}.\eta, \text{let}.ass\}$, the *weak γ-reduction* $\underset{w}{\rightarrow}_\gamma$ is the closure of the rule $\mapsto_\gamma$ under *weak contexts* E.

Example 1 (Non-confluence)**.** *Let M be a computation in normal form, for instance $M = xx$.*

$$
\begin{array}{ccc}
\text{let } y := (\text{let } x := zz \text{ in } M) \text{ in } [y] & \xrightarrow[\text{w}]{\text{let.}\eta} & \text{let } x := zz \text{ in } M \\
\quad\downarrow{\scriptstyle\text{let.ass}}\,\text{w} & & \\
\text{let } x := zz \text{ in } (\text{let } y := M \text{ in } [y]) & &
\end{array}
$$

Both $\text{let } x := zz \text{ in } M$ *and* $\text{let } x := zz \text{ in } (\text{let } y := M \text{ in } [y])$ *are normal for* $\underset{w}{\rightarrow}_{ml^*}$ *(in the latter, the let.η-redex* $\text{let } y := M \text{ in } [y]$ *cannot be fired by weak reduction), but they are distinct.*

Example 2 (Non-confluence)**.** *Let $R = P = Q = L = zz$ and:*

$$
M := \overline{\text{let } z = \underline{(\text{let } x = (\text{let } y = L \text{ in } Q) \text{ in } P)} \text{ in } R}
$$

There are two weak let.ass-redexes, the overlined one and the underlined one. So,

$$
M \underset{w}{\rightarrow}_{\text{let.ass}} \text{let } x := (\text{let } y := L \text{ in } Q) \text{ in } (\text{let } z := P \text{ in } R)
$$

$$
\underset{w}{\rightarrow}_{\text{let.ass}} \text{let } y := L \text{ in } (\text{let } x := Q \text{ in } (\text{let } z := P \text{ in } R)) =: M'
$$

$$
M \underset{w}{\rightarrow}_{\text{let.ass}} \text{let } z := (\text{let } y := L \text{ in } (\text{let } x := Q \text{ in } P)) \text{ in } R
$$

$$
\underset{w}{\rightarrow}_{\text{let.ass}} \text{let } y := L \text{ in } (\text{let } z := (\text{let } x := Q \text{ in } P) \text{ in } R) =: M''
$$

Both M' are M'' are normal for $\underset{w}{\rightarrow}_{ml^*}$ *(in M'', the let.ass-redex* $\text{let } z := (\text{let } x := Q \text{ in } P) \text{ in } R$ *is under the scope of a* let *and so cannot be fired by weak reduction), but they are distinct.*

3

Example 3.

Non-determinism—but confluence—of $\rightarrow_{let.\eta}$. Let $M = yy$ and $N = zz$:

$$
\begin{array}{ccc}
\mathsf{let}\,x := (\mathsf{let}\,y := (\mathsf{let}\,z := N\,\mathsf{in}\,[z])\,\mathsf{in}\,M)\,\mathsf{in}\,[x] & \xrightarrow[\mathsf{w}]{\;let.\eta\;} & \mathsf{let}\,y := (\mathsf{let}\,z := N\,\mathsf{in}\,[z])\,\mathsf{in}\,M \\[4pt]
\Big\downarrow{\scriptstyle let.\eta}\,\mathsf{w} & & \Big\downarrow{\scriptstyle let.\eta}\,\mathsf{w} \\[4pt]
\mathsf{let}\,x := (\mathsf{let}\,y := N\,\mathsf{in}\,M)\,\mathsf{in}\,[x] & \xrightarrow[\mathsf{w}]{\;let.\eta\;} & \mathsf{let}\,y := N\,\mathsf{in}\,M
\end{array}
$$

Summing up the situation:

1. $\underset{\mathsf{w}}{\rightarrow}_\beta$ and $\underset{\mathsf{w}}{\rightarrow}_{let.\beta}$ and $\underset{\mathsf{w}}{\rightarrow}_{\beta,let.\beta} := \underset{\mathsf{w}}{\rightarrow}_\beta \cup \underset{\mathsf{w}}{\rightarrow}_{let.\beta}$ are deterministic.

2. $\underset{\mathsf{w}}{\rightarrow}_{let.\eta}$ is non-deterministic, but it is confluent.

3. $\underset{\mathsf{w}}{\rightarrow}_{let.ass}$ is non-deterministic, non-confluent and normal forms are not unique.

4. $\underset{\mathsf{w}}{\rightarrow}_{let.ass} \cup \underset{\mathsf{w}}{\rightarrow}_{let.\beta} \cup \underset{\mathsf{w}}{\rightarrow}_\beta$ is non-deterministic, non-confluent and normal forms are not unique.

5. $\underset{\mathsf{w}}{\rightarrow}_{ml^*}$ is non-deterministic, non-confluent and normal forms are not unique.

(Non-)Factorization. Another remarkable aspect making the reduction theory for λ_{ml^*} (and for other computational λ-calculi) tricky to study is the lack of *factorization*, which is the simplest possible form of *standardization*.

In Plotkin's call-by-value λ-calculus [Plo75] (which can be seen as the restriction of λ_{ml^*} where the reduction is generated only by the β-rule), weak reduction satisfies factorization, that is any reduction sequence can be *reorganized* as weak steps followed by non-weak steps:

$$
\rightarrow_\beta^* \;\subseteq\; \underset{\mathsf{w}}{\rightarrow}_\beta^* \cdot \underset{\neg\mathsf{w}}{\rightarrow}_\beta^* \tag{1}
$$

But in λ_{ml^*} (and similar computational λ-calculi), weak factorization *does not hold*. The problem is here the $let.\eta$ rule, as shown by the following counterexample, due to van Oostrom [vO20].

Example 4 (Non-factorization [vO20]). *Consider*

$$
M := \mathsf{let}\,y := (zz)\,\mathsf{in}\,(\mathsf{let}\,x := [y]\,\mathsf{in}\,[x]) \;\underset{\neg\mathsf{w}}{\rightarrow}_{let.\eta}\; \mathsf{let}\,y := (zz)\,\mathsf{in}\,[y] \;\underset{\mathsf{w}}{\rightarrow}_{let.\eta}\; (zz) =: N
$$

Weak steps are not possible from M, so it is impossible *to factorize the reduction form M to N as $M \underset{\mathsf{w}}{\rightarrow}_{ml^*}^* \cdot \underset{\neg\mathsf{w}}{\rightarrow}_{ml^*}^* N$.*

A bridge between Evaluation and Reduction. On the one hand, computational λ-calculi such as λ_{ml^*} have an unrestricted *non-deterministic reduction* that generates the equational theory of the calculus, studied for foundational and semantic purposes. On the other hand, *weak reduction* has a prominent role in the literature of computational λ-calculi, because it models an ideal programming language. Indeed, when restricted to closed terms (which are the terms corresponding to programs), normal forms of weak reduction coincide with values; and when restricted to β and let.β steps, weak reduction is deterministic and corresponds to an abstract machine, implementing a programming language. It is then natural to wonder what is the relation between reduction and evaluation.

In Plotkin's call-by-value λ-calculus [Plo75], the following *convergence result* provides a bridge between reduction and evaluation: if a term M β-reduces to a value, then M only needs *weak β-reduction* to reach a value.

$$M \to_\beta^* V \text{ (for some value } V) \iff M \xrightarrow[\mathsf{w}]{}_\beta^* V' \text{ (for some value } V') \qquad (2)$$

In λ_{ml^*}, despite several drawbacks of weak reduction, we can still prove a *convergence* result similar to (2) relating reduction and evaluation: to reach a value in λ_{ml^*}, weak β-steps and weak let.β-steps suffice.

Theorem 5 (Convergence). *Let M be a computation in λ_{ml^*} and let $\to_{ml^*_-} := \to_{ml^*} \smallsetminus \to_\eta$.*

$$M \to_{ml^*_-}^* [V] \text{ (for some value } V) \iff M \xrightarrow[\mathsf{w}]{}_{\beta, \, let.\beta}^* [V'] \text{ (for some value } V') \qquad (3)$$

Because of the issues which we have presented, this result is non-trivial. We obtain it via the analysis of a calculus recently introduced by de'Liguoro and Treglia's, namely the *computational core* $\lambda_\copyright$ [dT20]. $\lambda_\copyright$ has the same issues, but a different syntax, which is more closely related to *calculi inspired by linear logic* [Sim05, EG16, GM19, FG21], whose properties and tools we can then use.

The analysis of the reduction theory of $\lambda_\copyright$ is carried-out in [FGdT21]. We then transfer the convergence of $\lambda_\copyright$ to that of λ_{ml^*}, via a rather sophisticated analysis of the translation.

Proof (sketch). The connection with calculi based on linear logic comes from the fact that the β-reduction of $\lambda_\copyright$ is the same as the $\beta_!$-reduction in Simpson's *linear λ-calculus* [Sim05], which has been studied in the literature of linear logic [EG16, GM19, FG21] also with the name *bang calculus*. There, a unary operator ! marks resources that can be duplicated or erased during reduction. The operator ! corresponds exactly to the coercion of a value V into a computation $[V]$ (the `return` operator in Haskell). This close relation allows us to rely on results and tools coming from the reduction theory of the bang calculus to study reduction and evaluation of β and let.β.

We obtain the proof of convergence for λ_{ml^*} in two phases:

1. We first study $\lambda_{\circledcirc}$, whose reduction is simpler to study thanks to the facts we mentioned above. Weak factorization and convergence of $\lambda_{\circledcirc}$ are analysed in [FGdT21], via several steps which progressively simplify the reduction sequences.

   - $\lambda_{\circledcirc}$ enjoys *surface factorization*, *i.e.* any reduction sequence can be rearranged by first performing *surface* steps (which are not in the scope of any !) and then non-surface steps.

   - The steps corresponding to the *identity law can be postponed* after all the other steps. So any reduction sequence can be rearranged to have an initial sequence **s** of surface steps which do not include (the $\lambda_{\circledcirc}$ version of) let.η.

   - For such initial surface sequence **s**, we have a *weak factorization* result, which extends (1) (recall that the calculus here is richer than Plotkin's CbV).

   - Finally, in the initial weak sequence, the steps which correspond to the *associativity law can also be postponed*, yielding in $\lambda_{\circledcirc}$ a convergence result analogous to (3).

2. We then transfer to λ_{ml^*} the convergence result of $\lambda_{\circledcirc}$. The difficulty here is that the translation between the two calculi does not directly preserve weak reduction: a more sophisticated analysis is needed.

Conclusion. We have analyzed evaluation (according to the standard definition of evaluation context) in a well-established formalization of the computational calculus, namely Sabry and Wadler's λ_{ml^*} [SW96]. Notice that the properties of non-confluence and non-factorization of evaluation which we presented actually hold in any calculus in which the monadic rules are oriented. We find this fact quite surprising, and worth to be explicitly stated. To our knowledge, it does not appear in the literature.

We are able to show that a convergence result analogous to (3) holds in λ_{ml^*}. Convergence in λ_{ml^*} relates full reduction to evaluation, and provides a theoretical justification to the following facts:

1. functional programming languages with computational effects use weak reduction as evaluation mechanism; indeed, weak reduction is enough to *return values*.

2. in computational λ-calculi, when interested in returning a value, the only *rules* of interest *for weak reduction* are β and let.β—which are deterministic and do not have unpleasant rewriting properties—while the rules let.*ass* and let.η can be safely considered as unoriented identities—they are necessary to the equational theory but are external to the reduction.

6

References

[BHM02] Nick Benton, John Hughes, and Eugenio Moggi. Monads and effects. In *Applied Semantics, International Summer School, APPSEM 2000*, volume 2395 of *Lecture Notes in Computer Science*, pages 42–122. Springer, 2002.

[DGL17] Ugo Dal Lago, Francesco Gavazzo, and Paul B. Levy. Effectful Applicative Bisimilarity: Monads, Relators, and Howe's Method. In *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017*, pages 1–12. IEEE Computer Society, 2017.

[dT20] Ugo de'Liguoro and Riccardo Treglia. The untyped computational λ-calculus and its intersection type discipline. *Theor. Comput. Sci.*, 846:141–159, 2020.

[EG16] Thomas Ehrhard and Giulio Guerrieri. The bang calculus: an untyped lambda-calculus generalizing call-by-name and call-by-value. In *Proceedings of the 18th International Symposium on Principles and Practice of Declarative Programming (PPDP 2016)*, pages 174–187. ACM, 2016.

[FG21] Claudia Faggian and Giulio Guerrieri. Factorization in call-by-name and call-by-value calculi via linear logic. In *Foundations of Software Science and Computation Structures - 24th International Conference, FOSSACS 2021*, volume 12650 of *Lecture Notes in Computer Science*, pages 205–225. Springer, 2021.

[FGdT21] Claudia Faggian, Giulio Guerrieri, Ugo de'Liguoro, and Riccardo Treglia. On reduction and normalization in the computational core. *CoRR*, abs/2104.10267, 2021. Submitted to Math. Struct. Comp. Sci., special issue of IWC 2020.

[GM19] Giulio Guerrieri and Giulio Manzonetto. The bang calculus and the two Girard's translations. In *Proceedings Joint International Workshop on Linearity & Trends in Linear Logic and Applications (Linearity-TLLA 2018)*, volume 292 of *EPTCS*, pages 15–30, 2019.

[LPT03] Paul Blain Levy, John Power, and Hayo Thielecke. Modelling environments in call-by-value programming languages. *Information and Computation*, 185(2):182 – 210, 2003.

[Mac97] Saunders MacLane. *Categories for the Working Mathematician*. Graduate Texts in Mathematics. Springer, 2 edition, 1997.

[Mog88] Eugenio Moggi. Computational Lambda-calculus and Monads. Report ECS-LFCS-88-66, University of Edinburgh, Edinburgh, Scotland, October 1988.

[Mog89]   Eugenio Moggi. Computational lambda-calculus and monads. In *Proceedings of the Fourth Annual Symposium on Logic in Computer Science (LICS '89)*, pages 14–23. IEEE Computer Society, 1989.

[Mog91]   Eugenio Moggi. Notions of computation and monads. *Inf. Comput.*, 93(1):55–92, 1991.

[Plo75]   Gordon D. Plotkin. Call-by-name, call-by-value and the lambda-calculus. *Theor. Comput. Sci.*, 1(2):125–159, 1975.

[PP02]   Gordon D. Plotkin and John Power. Notions of computation determine monads. In *Foundations of Software Science and Computation Structures, 5th International Conference, FOSSACS 2002*, volume 2303 of *Lecture Notes in Computer Science*, pages 342–356. Springer, 2002.

[PP03]   Gordon D. Plotkin and John Power. Algebraic operations and generic effects. *Applied Categorical Structures*, 11(1):69–94, 2003.

[Sim05]   Alex Simpson. Reduction in a linear lambda-calculus with applications to operational semantics. In Jürgen Giesl, editor, *Term Rewriting and Applications*, pages 219–234, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.

[SW96]   Amr Sabry and Philip Wadler. A reflection on call-by-value. In Robert Harper and Richard L. Wexelblat, editors, *Proceedings of the 1996 ACM SIGPLAN International Conference on Functional Programming, ICFP 1996, Philadelphia, Pennsylvania, USA, May 24-26, 1996*, pages 13–24. ACM, 1996.

[SW97]   Amr Sabry and Philip Wadler. A reflection on call-by-value. *ACM Trans. Program. Lang. Syst.*, 19(6):916–941, 1997.

[vO20]   Vincent van Oostrom. Private communication via electronic mail, 2020.

[WT03]   Philip Wadler and Peter Thiemann. The marriage of effects and monads. *ACM Trans. Comput. Log.*, 4(1):1–32, 2003.