Church–Rosser Made Easy

Dexter Kozen
Department of Computer Science
Cornell University
Ithaca, New York 14853-7501, USA

For Jurek, on the Occasion of his Sixtieth Birthday

Abstract

The Church–Rosser theorem states that the λ-calculus is confluent under α- and β-reductions. The standard proof of this result is due to Tait and Martin-Löf. In this note, we present an alternative proof based on the notion of *acceptable orderings*. The technique is easily modified to give confluence of the $\beta\eta$-calculus.

1 Introduction

A fundamental result in the λ-calculus is *confluence*: if $e \to e_1$ and $e \to e_2$ by some arbitrary sequences of reductions, then there exists an e_3 such that $e_1 \to e_3$ and $e_2 \to e_3$. This is result is originally due to Alonzo Church and J. Barkley Rosser in 1936 [2] and is known as the *Church–Rosser theorem*.

The standard proof of this result, as presented by Barendregt [1] (see also [5, 6]) is due to Tait and Martin-Löf. The Tait–Martin-Löf proof is based on an auxiliary reduction relation defined by formal rules and is very amenable to machine verification. Several implementations in automated deduction systems have been reported, along with various improvements and simplifications [5, 6]. A good overview is given in [5].

Besides the Tait–Martin-Löf proof, Barendregt [1, Chp. 11] presents another proof based on the idea of *developments*. This proof involves tracing a set of occurrences of redexes in a term through a sequence of reductions. It is somewhat more transparent than the Tait–Martin-Löf proof, but is longer and unfortunately does not readily generalize to the $\beta\eta$-calculus.

The same result in the presence of η-reductions was first proved by Curry and Feys [3] and later improved and generalized by Hindley [4]. These proofs show that every reduction sequence can be transformed to one in a special normal form.

In this note we offer a short alternative treatment based on the notion of *acceptable orderings*. We prove confluence under α- and β-reductions in Section 3. A slight modification of the proof admits η-reductions as well, and we present this modification in Section 4.

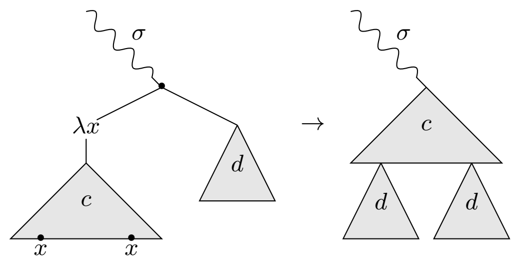

Figure 1: A β-reduction at σ

2 Preliminaries

2.1 λ-Terms as Labeled Trees

We view λ-terms as finite labeled trees. A *tree* is a nonempty prefix-closed subset of $\mathbb{N}^*$. A λ-*term* is a partial function e whose domain $\operatorname{dom} e \subseteq \mathbb{N}^*$ is a finite tree such that $e(\sigma)$ is either an abstraction operator λx, in which case σ has one child $\sigma 0$; the application operator, in which case σ has a left child $\sigma 0$ and a right child $\sigma 1$; or a variable, in which case σ is a leaf.

The *subterm of* e *rooted at* $\sigma \in \operatorname{dom} e$ is the term $e \restriction \sigma = \lambda\tau.\, e(\sigma\tau)$ (λ is used here as a meta-operator). Note that if σ is a prefix of τ, then $e \restriction \tau$ is a subterm of $e \restriction \sigma$.

We adopt the *Barendregt variable convention*: α-equivalent terms are considered identical. The β-reduction rule then takes the following form. Suppose σ is a β-redex in e, say $e \restriction \sigma = (\lambda x.\, c)\, d$. This is replaced at σ by the corresponding contractum consisting of the term c with d substituted for all free occurrences of x, renaming bound variables as necessary to avoid capture (Fig. 1).

2.2 Acceptable Orderings

If σ and τ are both β-redexes in e and σ is a proper prefix of τ, then $e \restriction \tau$ is a proper subterm of $e \restriction \sigma$. If we reduce σ before reducing τ, then τ will in general no longer be a redex; indeed, it may no longer even exist in the resulting tree. However, if we reduce τ first, then σ is still a redex in the resulting tree, although the subterm at σ may have changed.

More generally, if $A \subseteq \operatorname{dom} e$ is a set of β-redexes in e, and we reduce them in some order consistent with the subterm relation—that is, we reduce $\sigma \in A$ only if all proper extensions $\sigma\tau \in A$ have already been reduced—then every redex in A will still be available when it is time to reduce it, and it will be possible to reduce all of them. Moreover, the actual order does not matter, as long as it is consistent with the subterm relation.

Formally, we say that a linear ordering $\sigma_1, \ldots, \sigma_n$ of the elements of A is *acceptable* if $\sigma_i = \sigma_j \tau$ implies $i \le j$; in other words, the sequence $\sigma_1, \ldots, \sigma_n$ is a subsequence of some total extension of the partial order $\{(\sigma\tau,\, \sigma) \mid \sigma,\ \tau \in \mathbb{N}^*\}$ (or, if you like, a topological sort of A with respect to the edges $(\sigma\tau,\, \sigma)$).

Acceptable orderings of A are not unique, but this does not matter: it is easily proved inductively that all acceptable orderings give reduction sequences of the same length, namely the cardinality of A, and the resulting final terms are the same up to α-equivalence. Let us call this final term $\theta_A(e)$, as it depends only on e and A and not on the order of reductions.

For $\sigma \in \mathbb{N}^*$, let $\sigma\!\downarrow = \{\sigma\tau \mid \tau \in \mathbb{N}^*\}$. For $\sigma \in \operatorname{dom} e$, $\sigma\!\downarrow \cap \operatorname{dom} e$ represents the set of subterms of $e \restriction \sigma$.

For $A, X \in \mathbb{N}^*$, write $A \rhd X$ if there exists a σ such that $A \subseteq \sigma\!\downarrow$ and $\sigma\!\downarrow \cap X = \varnothing$. If $A \rhd X$, then A and X are disjoint, and there exists an acceptable ordering of $A \cup X$ such that all elements of A come before all elements of X.

3 Confluence under β-Reductions

We start by proving confluence under β-reductions in some special cases, building up to the general result in Theorem 3.6.

Lemma 3.1 *Let A and B be two sets of redexes of e such that all elements of A are prefix-incomparable to all elements of B. Then $\theta_B(\theta_A(e))$ and $\theta_A(\theta_B(e))$ both exist and are equal. This gives the confluent diagram illustrated in Fig. 2.*

Proof. Both $\theta_B(\theta_A(e))$ and $\theta_A(\theta_B(e))$ represent the reduction of the redexes in $A \cup B$ in different acceptable orders, thus both terms are equal to $\theta_{A\cup B}(e)$. $\square$

$$\begin{tikzcd}
 & e \arrow[dl, "B"'] \arrow[dr, "A"] & \\
\theta_B(e) \arrow[dr, "A"] & & \theta_A(e) \arrow[dl, "B"'] \\
 & \begin{array}{c}\theta_B(\theta_A(e))\\ = \theta_A(\theta_B(e))\end{array} &
\end{tikzcd}$$

Figure 2

Lemma 3.2 *Let σ be a redex of e, and let A be a set of redexes of e such that $A \subseteq \sigma\!\downarrow$. Then there exists a set B of redexes of $\theta_\sigma(e)$ such that $B \subseteq \sigma\!\downarrow$ and*

$$\theta_C(\theta_A(e)) \;=\; \theta_B(\theta_\sigma(e)),$$

where $C = \{\sigma\}$ if $\sigma \notin A$ and $C = \varnothing$ if $\sigma \in A$. This gives the confluent diagram illustrated in Fig. 3.

$$\begin{tikzcd}
 & e \arrow[dl, "A"'] \arrow[dr, "\sigma"] & \\
\theta_A(e) \arrow[dr, "C"] & & \theta_\sigma(e) \arrow[dl, "B"'] \\
 & \begin{array}{c}\theta_C(\theta_A(e))\\ = \theta_B(\theta_\sigma(e))\end{array} &
\end{tikzcd}$$

Figure 3

Proof. Suppose first that $\sigma \notin A$. Let $e \restriction \sigma = (\lambda x.\, c)\, d$. The set A may contain redexes in c and d. Reducing σ first, a copy of d replaces each free occurrence of x in c (see Fig. 1). If we then reduce the redexes in these copies of d in some acceptable order, then reduce the remaining redexes in c in some acceptable order, this yields the same result as reducing the redexes in d and c in some acceptable order before reducing σ, then reducing σ.

Formally, take $B = \{\sigma\gamma_i \mid 1 \le i \le m\} \cup \{\sigma\delta_i\tau_j \mid 1 \le i \le k,\ 1 \le j \le n\}$, where $A = \{\sigma 00\gamma_i \mid 1 \le i \le m\} \cup \{\sigma 1\tau_j \mid 1 \le j \le n\}$ and the free occurrences of x in c are located at $\{\sigma 00\delta_1, \ldots, \sigma 00\delta_k\}$. The elements of A of the form $\sigma 00\gamma_i$ represent the redexes in c, which after reducing σ become the elements of B of the form $\sigma\gamma_i$. The elements of A of the form $\sigma 1\tau_j$ represent the redexes in d, which after reducing σ become the elements of B of the form $\sigma\delta_i\tau_j$ representing the corresponding redexes in the copies of d that replaced the free occurrences of x in c. In Fig. 1, $k = 2$.

If $\sigma \in A$, then it must appear last in any acceptable ordering of A. By the previous argument, there exists $B \subseteq \sigma\!\downarrow$ such that $\theta_\sigma(\theta_{A-\{\sigma\}}(e)) = \theta_B(\theta_\sigma(e))$, therefore $\theta_\varnothing(\theta_A(e)) = \theta_\sigma(\theta_{A-\{\sigma\}}(e)) = \theta_B(\theta_\sigma(e))$. □

Lemma 3.3 *Let A and X be sets of redexes of e such that $A \rhd X$. There exists a set B of redexes of $\theta_X(e)$ such that*

$$\theta_X(\theta_A(e)) \;=\; \theta_B(\theta_X(e)).$$

Proof. This follows easily by induction on the cardinality of X using Lemmas 3.1 and 3.2. Starting with $X_0 = X$ and $B_0 = A$, construct a sequence of sets X_i and B_i by taking the elements of X one at a time in some acceptable order, maintaining the invariant $B_i \rhd X_i$. Fig. 4 illustrates the case $X = \{\tau_1, \tau_2, \tau_3\}$. □

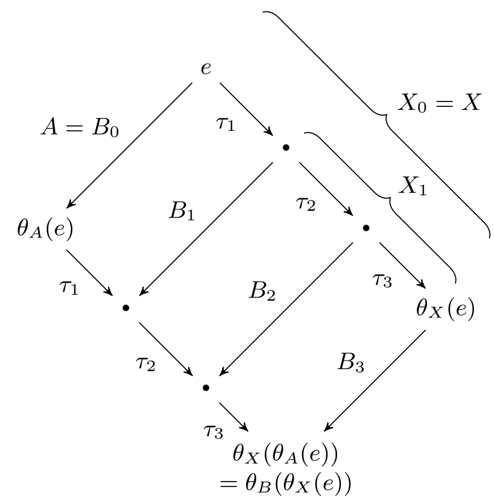

Figure 4

Lemma 3.4 *Let A be an arbitrary set of redexes of e, and let σ be a redex of e. Then there exist redex sets C of $\theta_A(e)$ and B of $\theta_\sigma(e)$ such that*

$$\theta_C(\theta_A(e)) \;=\; \theta_B(\theta_\sigma(e)).$$

This gives the confluent diagram of Fig. 3 (the same diagram as for Lemma 3.2, but with a different interpretation of the symbols).

Proof. Partition A into $A_1 = \sigma\!\downarrow \cap A$ and $A_2 = A - A_1$. Then $A_1 \rhd A_2$. By Lemma 3.2, there exist a set $B_1 \subseteq \sigma\!\downarrow$ of redexes of $\theta_\sigma(e)$ and $C_1 \subseteq \{\sigma\}$ such that

$$\theta_{C_1}(\theta_{A_1}(e)) \;=\; \theta_{B_1}(\theta_\sigma(e)). \tag{3.1}$$

Take $B = B_1 \cup A_2$. Since $B_1 \subseteq \sigma\!\downarrow$, $C_1 \subseteq \sigma\!\downarrow$, and $\sigma\!\downarrow \cap A_2 = \varnothing$, we have $B_1 \rhd A_2$ and $C_1 \rhd A_2$. By Lemma 3.3, there exists a set C of redexes of $\theta_{A_2}(\theta_{A_1}(e)) = \theta_A(e)$ such that

$$\theta_C(\theta_{A_2}(\theta_{A_1}(e))) \;=\; \theta_{A_2}(\theta_{C_1}(\theta_{A_1}(e))). \tag{3.2}$$

Then

$$\begin{aligned}
\theta_C(\theta_A(e)) &= \theta_C(\theta_{A_2}(\theta_{A_1}(e))) && \text{since } A_1 \rhd A_2 \\
&= \theta_{A_2}(\theta_{C_1}(\theta_{A_1}(e))) && \text{by (3.2)} \\
&= \theta_{A_2}(\theta_{B_1}(\theta_\sigma(e))) && \text{by (3.1)} \\
&= \theta_B(\theta_\sigma(e)) && \text{since } B_1 \rhd A_2.
\end{aligned}$$

□

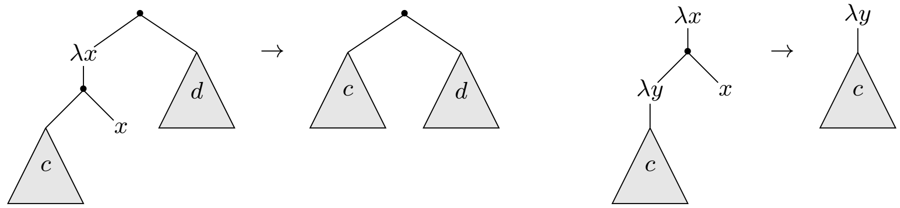

(a) (b)

Figure 5: Overlapping redexes: (a) a β-η overlap; (b) an η-β overlap

Lemma 3.5 *Let $e \to e'$ by some arbitrary sequence of β-reductions, and let A be a set of redexes of e. Then there exists a set B of redexes of e' such that $\theta_A(e) \to \theta_B(e')$.*

Proof. This follows in a straightforward fashion by induction on the length of the reduction $e \to e'$ by composing the reductions of Lemma 3.4. □

Theorem 3.6 (Church–Rosser Theorem) *Let $e \to e_1$ and $e \to e_2$ by some arbitrary sequences of β-reductions. Then there exists an e_3 such that $e_1 \to e_3$ and $e_2 \to e_3$.*

Proof. Lemma 3.5 gives a confluent diagram for each step in the reduction sequence $e \to e_1$, and these can be composed to get a confluent diagram for the entire sequence. □

4 Accommodating η

The η-reduction rule is $\lambda x.cx \to c$, where c contains no free occurrences of x. We show in this section that a minor modification of the argument of Section 3 gives confluence under β- and η-reductions.

The main concern is that due to *overlapping redexes*, it is no longer true in general that any set of redexes $A \subseteq \mathrm{dom}\, e$ can be completely reduced simply by reducing them in acceptable order. There are two problematic situations, as illustrated in Fig. 5.

Consider the configuration of Fig. 5(a). There is a β-redex at the root whose left child is an η-redex. If the η-reduction is performed first, the root is no longer a β-redex in general. However, a key observation is that we can perform either the β-reduction at the root or the η-reduction at the left child, and the resulting contractum is the same, as shown.

Similarly, Fig. 5(b) shows an η-redex at the root whose only child is a β-redex. As with (a), performing the β-reduction at the child may destroy the η-redex at the root. However, if we perform either reduction, the resulting contractum is the same (up to α-equivalence), as shown.

The solution is simply to disallow redex sets A containing either of these two configurations. Equivalently, A may not contain both σ and $\sigma 0$ for any σ. We will call a redex set $A \subseteq \mathrm{dom}\, e$ *overlap-free* if this property holds. Any overlap-free set of redexes can be fully reduced in acceptable order.

The entire development of Section 3 now goes through with minor modification. The formal statements of Lemmas 3.1–3.5 and Theorem 3.6 are modified as follows:

Lemma 4.1 *Let A and B be two overlap-free sets of redexes of e such that all elements of A are prefix-incomparable to all elements of B. Then $\theta_B(\theta_A(e))$ and $\theta_A(\theta_B(e))$ both exist and are equal.*

Lemma 4.2 *Let σ be a redex of e, and let A be an overlap-free set of redexes of e such that $A \subseteq \sigma\!\downarrow$. Then there exists an overlap-free set B of redexes of $\theta_\sigma(e)$ such that $B \subseteq \sigma\!\downarrow$ and*

$$\theta_C(\theta_A(e)) \;=\; \theta_B(\theta_\sigma(e)),$$

where $C = \{\sigma\}$ if both $\sigma, \sigma 0 \notin A$, $C = \varnothing$ if either $\sigma \in A$ or $\sigma 0 \in A$.

Lemma 4.3 *Let A and X be sets of redexes of e such that $A \rhd X$ and $A \cup X$ is overlap-free. There exists an overlap-free set B of redexes of $\theta_X(e)$ such that*

$$\theta_X(\theta_A(e)) \;=\; \theta_B(\theta_X(e)).$$

Lemma 4.4 *Let A be an arbitrary overlap-free set of redexes of e, and let σ be a redex of e. Then there exist overlap-free redex sets C of $\theta_A(e)$ and B of $\theta_\sigma(e)$ such that*

$$\theta_C(\theta_A(e)) \;=\; \theta_B(\theta_\sigma(e)).$$

Lemma 4.5 *Let $e \to e'$ by some arbitrary sequence of β- and η-reductions, and let A be an overlap-free set of redexes of e. Then there exists an overlap-free set B of redexes of e' such that $\theta_A(e) \to \theta_B(e')$.*

Theorem 4.6 (Church–Rosser Theorem for the $\beta\eta$-calculus) *Let $e \to e_1$ and $e \to e_2$ by some arbitrary sequences of β- and η-reductions. Then there exists an e_3 such that $e_1 \to e_3$ and $e_2 \to e_3$.*

Lemma 4.2 for the case of σ a β-redex is the same as in Lemma 3.2, with the extra observation that B cannot contain overlapping redexes if A did not. For the case of σ an η-redex, if $A = \{\sigma 00\gamma_i \mid 1 \le i \le m\}$, we take $B = \{\sigma\gamma_i \mid 1 \le i \le m\}$. In both cases, if $\sigma \in A$ or $\sigma 0 \in A$, we can take $C = \varnothing$, otherwise $C = \{\sigma\}$.

For Lemma 4.4, we can assume without loss of generality that $A \cup \{\sigma\}$ is overlap-free; for if $\tau \in A$ and either $\tau = \sigma 0$ or $\sigma = \tau 0$, we can just replace σ with τ in the proof, as $\theta_\sigma(e) = \theta_\tau(e)$. We can then conclude that the $B_1 \cup A_2$ and $C_1 \cup A_2$ constructed in the proof are overlap-free. All else is the same as in Lemma 3.4.

The proofs of Lemmas 4.1, 4.3, 4.5, and Theorem 4.6 go through essentially unchanged.

Acknowledgements

It is an honor and a privilege to acknowledge Jurek Tiuryn on his birthday. I am grateful for my many years of association with him, and I look forward to many more. Whether collaborating on research in Warsaw, painting an apartment in New York, building a playground in Ithaca, or installing electricity in an old farmhouse in Karczmisko, it is always a pleasure to be in his company. One could wish for no finer colleague and friend.

Thanks to Andrew Myers and Vasilis Syrgkanis for valuable discussions. This work was supported by NSF grant CCF-0635028.

References

[1] H. P. Barendregt. *The Lambda Calculus, Its Syntax and Semantics*. North-Holland, 2nd edition, 1984.

[2] Alonzo Church and J. Barkley Rosser. Some properties of conversion. *Trans. Amer. Math. Soc.*, 39:472–482, 1936.

[3] Haskell B. Curry and R. Feys. *Combinatory Logic*, volume I. North Holland, 1958.

[4] R. Hindley. Standard and normal reductions. *Trans. Amer. Math. Soc.*, 241:253–271, July 1978.

[5] Robert Pollack. Polishing up the Tait–Martin-Löf proof of the Church–Rosser theorem. In *Proc. De Wintermöte 95*. Department of Computing Science, Chalmers University, Göteborg, Sweden, January 1995.

[6] Masako Takahashi. Parallel reductions in λ-calculus (revised version). *Information and Computation*, 118(1):120–127, April 1995.