Linear β-reduction[*]

Stefano Guerrini

LIPN, Institut Galilée, Université Paris Nord 13, Sorbonne Paris Cité

stefano.guerrini@univ-paris13.fr

Linear head reduction is a key tool for the analysis of reduction machines for λ-calculus and for game semantics. Its definition requires a notion of redex at a distance named primary redex in the literature. Nevertheless, a clear and complete syntactic analysis of this rule is missing. We present here a general notion of β-reduction at a distance and of linear reduction (i.e., not restricted to the head variable), and we analyse their relations and properties. This analysis rests on a variant of the so-called σ-equivalence that is more suitable for the analysis of reduction machines, since the position along the spine of primary redexes is not permuted. We finally show that, in the simply typed case, the proof of strong normalisation of linear reduction can be obtained by a trivial tuning of Gandy's proof for strong normalisation of β-reduction.

1 Introduction

Linear head reduction is a key tool for the analysis of reduction machines for λ-calculus and for game semantics. A detailed analysis of it, and more generally of a notion of reduction at a distance, has been given by Accattoli [1] in terms of proof nets and explicit substitutions. Linear head reduction is usually presented in terms of the so-called σ-equivalence introduced by Regnier in [8]. In the following, we introduce a variant of the σ-equivalence, which has the main advantage of leaving unchanged the order of primary redexes (a notion of β-redex that will be discussed later). Such a new equivalence is more suitable for the analysis of abstract reduction machines based on linear head reduction, as for instance Danos and Regnier's Pointer Abstract Machine (PAM) [3], which has been analysed in detail by the author and Pellitta in [5]. Indeed, most of the material that we shall present in this paper has been developed for formalising the results in [5].

The key tool of our approach is a notion of context which is indeed an implicit representation of environments mapping variables to their values. By means of these contexts, one can define a β-reduction at a distance and its linearised version. Both of these reduction rules preserve β-equivalence, and both of them are strongly normalising in the case of simply typed λ-calculus. However, the proof of strong normalisation is not at all evident. In fact, linear reduction does not erase any term, it just replaces one of the occurrences of a variable with a (larger) λ-term; in other words, the size of the reducing term always increases along the reduction. Surprisingly, this apparent difficulty can be trivially overcome by a small tuning of Gandy's proof for strong normalisation of β-reduction [4]. Just by changing a detail in the interpretation of variable occurrences—it suffices to increase by 1 their measure—we can adapt the measure used in Gandy's proof to the case of linear reduction. Moreover, the new measure obtained in this way simultaneously proves strong normalisation of β-reduction and of its linearised version.

As already remarked, linear reduction has been studied in detail by Accattoli [1] by means of linear logic proof nets. Such an approach has been inspired by the structural calculus introduced by Accattoli

and Kesner [2], a calculus with explicit substitutions and reduction rules at a distance. In the present paper, our goal is to analyse linear reduction directly on λ-calculus, without introducing explicit substitutions or without going down to the low level analysis of reduction that can be achieved by means of proof nets. As we shall see later, such a goal is achieved by introducing a variant of σ-equivalence, named E-equivalence, which is more suitable for investigating reduction machines based on pointers, as for instance the PAM. Moreover, the proof of strong normalisation that we shall give is much simpler then the one based on reducibility candidates required in the case of proof nets. Recently, Pedrot and Saurin [7] have proposed a call-by-need variant of $\lambda\mu$-calculus defined in terms of a notion of closure contexts. Such closure contexts correspond to the E-contexts introduced in the following by Definition 2, but extended to $\lambda\mu$-calculus too. We remark that most of the material that we shall present below is a by-product of the study of the PAM started in [5], and that one of the further developments mentioned in [5] is the extension of the PAM to the $\lambda\mu$-calculus; Pedrot and Saurin's call-by-need closure contexts seem to be the right tool for formalising such an extension.

2 Preliminaries

The set of the λ-terms Λ is defined by the abstract grammar $s,t ::= x \mid \lambda x.t \mid st$, where $x \in \mathcal{V}$ and λx is a binder for the variable x. The set of the free variables of a term t is denoted by $\mathsf{FV}(t)$. The key computational step of λ-calculus is β-contraction $(\lambda x.t)s \to_\beta t\{s/x\}$, where $t\{s/x\}$ denotes that every free occurrence of the variable x in t is replaced by s, provided that such a replacement does not cause any name clash of some free variable of s; otherwise, if this is not the case, one has to preliminarily apply a suitable sequence of variable renamings, or α-rules, to t. (The α-congruence is the least congruence induced by the α-rule $\lambda x.t = \lambda y.t\{y/x\}$, in which y replaces all the free occurrences of x in t and y does not occur in t.) As usual, $\to_\beta^*$ denotes the reflexive and transitive closure of the binary relation defined by the β-rule, and $=_\beta$ denotes the corresponding equivalence (closing by symmetry also). Such notations will extend to the other rewriting rules that we shall see in the paper.

In order to avoid the bureaucratic problems connected to α-congruence, we can assume to work modulo it, and that all the bound variables in the terms that we shall consider have the *distinct names property* (sometimes referred to as Baredrengt variable names convention). A term has the distinct names property if no free variable in it has the same name of a bound variable, and all the bound variables have distinct names. Remarkably, for every λ-term there is an α-congruent one which has the distinct name property. In this way, no name clash can arise by replacing s for x in t in the β-reduction of $(\lambda x.t)s$. However, even if correct, the resulting term $t\{s/x\}$ might not have the above distinct names property. In order to guarantee that $t\{s/x\}$ preserves the distinct names property of $(\lambda x.t)s$, we can assume to replace each occurrence of x with a *fresh* copy of s, in which every bound variable has a fresh name which has not been already used in the term or in another copy of s.

In the simply typed λ-calculus, every term has a type. The set of types is given by the abstract grammar $\tau,\sigma ::= o \mid \tau \to \sigma$, where the constant o is the unique *base type* and any type $\tau \to \sigma$ is said a *functional type*. The set $\Lambda^\to$ of the simply typed terms is the subset of Λ whose terms respect the following typing rules: (*i*) each variable x has a given type τ; (*ii*) if the variable x has type τ, and the term t has type σ, then $\lambda x.t$ has type $\tau \to \sigma$; (*iii*) if the term s has type $\tau \to \sigma$ and t has type τ, then st has type σ. We shall write $t : \tau$ or t^τ to denote that a term t has type τ. The β-rule preserves typing; namely, if $t \to_\beta^* s$ and $t : \tau$, then $s : \tau$.

A reduction strategy is a set of rules specifying how to reduce a λ-term. Roughly speaking, given a reducible term t, a reduction strategy is a function that selects the redex (or the redexes) of t that must

(or among which we can choose the redex to) be reduced at the next step. A reduction strategy defines a sub-rewriting system of β-reduction and, in some cases, if some β-reducible term t contains no valid redex for the given reduction strategy, it introduces new normal forms.

2.1 Head reduction

Let us say that a β-redex $(\lambda x.t)s$ is in *outermost head position*[1] in v when $v = \lambda y_1 \dots \lambda y_k.(\lambda x.t)su_1 \dots u_h$, and that v head reduces to $v' = \lambda y_1 \dots \lambda y_k.t\{s/x\}u_1 \dots u_h$, written $v \to_h v'$, by reducing its outermost head redex. A term v is in *head normal form* when $\lambda y_1 \dots \lambda y_k.xu_1 \dots u_h$, which in general is not a β-normal form, since $u_1, \dots, u_h$ may contain β-redexes. Indeed, the β-normal form of t, if it exists, can be found by head reducing a term t to its head normal form $\lambda y_1 \dots \lambda y_k.xu_1 \dots u_h$ (if any) first, and then by recursively applying the head reduction strategy to every u_i and to the subterms of their head normal forms.

2.2 Head contexts

As usual, a context C is a term with a hole $\square$ (a sort of dummy free variable occurring exactly once in the term) $C ::= \square \mid \lambda x.C \mid Ct \mid tC$. Given any term t, by $C[t]$ we denote the term obtained by replacing the hole of the context C with the term t, without performing any variable renaming; therefore, when the hole is under the scope of a λ-abstraction binding the variable x, any free occurrence of x in t is captured in $C[t]$, and becomes bound.

Definition 1 (H-context, head variable) A *head context*, or *H-context*, is a context whose hole appears in head position. More precisely, H-contexts are defined by the following grammar

$$H ::= \square \mid \lambda x.H \mid Ht.$$

A head context of a term t is any H-context H s.t. $t = H[s]$, for some term s, that we shall say to be in head position in t. In particular, for every λ-term t, there is a unique head context H of t (the *maximal head context* of t) and a unique variable $x = \mathsf{hv}(t)$ (the *head variable* of t) s.t. $t = H[x]$.   □

2.3 Spine

A spine λ-abstraction/application of a term t is any λ-abstraction/application in head position in t. The *spine* of $t = H[x]$, and of its head context H, is the sequence of its spine λ-abstractions/applications ordered from the head variable of t (the hole of H) to its root. A variable x bound by a spine abstraction is a *spine variable*, while the right subterm of a spine application is a *spine argument* of t. By $\mathsf{SV}(t)$ and $\mathsf{SV}(H)$ we denote the set of the spine variables of a term t and of a H-context H, respectively.

A H-context H_λ is a *λ-context* if its spine is formed of λ-abstractions only (equivalently, H_λ has no spine arguments). A H-context $H_@$ is a *@-context* if its spine is formed of applications only (equivalently, $H_@$ has no spine variables).

---

[1]Usually this is simply referred to as *head position*. In the following we shall however present a larger notion of head position, in which a β-redex may be in head position even if it is inside the body of a β-redex in head position. According to such a new notion of head position, the redex reduced by the head reduction is the outermost β-redex in head position.

3 β-reduction at a distance

3.1 Environment contexts

Definition 2 (E-context) An *environment context*, or *E-context*, is a particular H-context in which spine λ-abstractions and spine applications are balanced. E-contexts are defined by the grammar

$$E_{1,2} ::= \Box \mid E_1[\lambda x.E_2]t$$

An E-context E contains an equal number $\#_p E$ of spine variables and of spine arguments. For every E-context $E \neq \Box$, there is a unique pair (x,t) s.t. $E = E_1[\lambda x.E_2]t$, for some pair of E-contexts E_1, E_2. Therefore, every E-context defines a unique bijection between its spine variables and its spine arguments. Such a correspondence can be formalised in terms of environments. An *environment* $\eta = t_1/x_1, \ldots, t_k/x_k$ is an ordered sequence of variable substitutions t_i/x_i (where t_i is a term replacing the variable x_i). Given an environment η, we define $t\{\eta\} = t\{t_1/x_1, \ldots, t_k/x_k\} = t\{t_1/x_1\} \ldots \{t_k/x_k\}$.

Definition 3 The environment $\eta(E)$ associated to an E-context is inductively defined by

$$\eta(\Box) = \varepsilon \qquad \text{and} \qquad \eta(E_1[\lambda x.E_2]t) = \eta(E_2), t/x, \eta(E_1)$$

According to the above definition, every pair of matching spine argument/variable corresponds to a substitution t/x in $\eta(E)$. We remark that the order of the substitutions in an environment is relevant, since for $i < j$, the occurrences of x_j in the term t_i are replaced by the term t_j, while this is not the case for any occurrence of x_j in a term t_k with $k \geq j$. In particular, the order of the spine variables in $\eta(E)$ corresponds to the order in which they appear in E, assuming to move from the inner head position to the root. In other words, x precedes y in $\eta(E)$ iff the binder of x is in the scope of the binder of y.

Lemma 1 *Let E be an E-context. For every λ-term t, $E[t] \to_\beta^* t\{\eta(E)\}$.* $\Box$

3.2 Primary redexes and β-contraction at a distance

Any pair of matching spine argument/variable in an E-environment is as a sort of redex at a distance.

Definition 4 (Primary β-redex) A β-*redex at a distance* is a term $E[\lambda x.t]s$, where E is an E-context. A *primary β-redex* is a β-redex at a distance occurring in a head position. $\Box$

As a particular case, for $E = \Box$, any β-redex is a β-redex at a distance. β-redexes at a distance can be reduced as usual β-redexes, by defining the following generalisation at a distance of the β-rule

$$E[\lambda x.t]s \to_{\beta_d} E[t\{s/x\}]$$

and by taking the β_d-reduction as the closure by contexts of the above rule. The H-context $E[\lambda x.\Box]s$ of a β-redex at a distance is an E-context. Then, every pair t/x of matching spine argument/variable of an E-context (and therefore every substitution in $\eta(E)$) forms a primary redex. As a consequence, it is readily seen that $E[t] \to_{\beta_d}^* t\{\eta(E)\}$ for every E-context E and every term t. More generally, β-reduction at a distance is sound w.r.t. the usual β-reduction.

Proposition 1 *Let $t \to_{\beta_d}^* s$, then $t =_\beta s$. Moreover, s is a normal form for $\to_{\beta_d}$ iff it is a β-normal form.* $\Box$

4   Spine permutation equivalence of λ-terms

The *head canonical E-contexts* are a particular case of E-contexts in which every redex at a distance is also a β-redex. Head canonical E-contexts are defined by the grammar $E_c ::= \Box \mid (\lambda x.E_c)t$, and any head canonical E-context has the shape $(\lambda x_n. \ldots . (\lambda x_2.(\lambda x_1.\Box)t_1)t_2 \ldots)t_n$. An environment η can be seen as the explicit representation of a head canonical E-context $\mathscr{E}(\eta)$ in which the order of the β-redexes along the spine is the inverse of the substitution pairs in the environment

$$t_1/x_1, t_2/x_2, \ldots, t_n/x_n \qquad \overset{\mathscr{E}}{\longmapsto} \qquad (\lambda x_n. \ldots . (\lambda x_2.(\lambda x_1.\Box)t_1)t_2 \ldots)t_n$$

Which corresponds to the inductive definition $\mathscr{E}(\varepsilon) = \Box$, and $\mathscr{E}(t/x, \eta) = \mathscr{E}(\eta)[(\lambda x.\Box)t]$.

4.1   Surface E-equivalence

By Lemma 1, we have that $E_1 =_\beta E_2$, for every pair of E-contexts E_1 and E_2 s.t. $\eta(E_1) = \eta(E_2)$. We can then define the following equivalence.

Definition 5 (Surface E-equivalence on E-contexts) The *surface E-equivalence on E-contexts* is the least equivalence $\sim_E$ defined by

$$E_1[\lambda x.E_2]t \sim_E E_1[(\lambda x.E_2)t] \qquad\qquad \text{if } \mathsf{FV}(t) \cap \mathsf{SV}(E_1) = \emptyset$$
$$E_1[E_2] \sim_E E_1'[E_2'] \qquad\qquad \text{if } E_i \sim_E E_i' \text{ for } i = 1, 2$$

Such an equivalence captures exactly the equivalence classes of E-contexts $\{E \mid \mathscr{E}(\eta(E)) = E_c\}$, where E_c is head canonical, as formally stated by the following lemma.

Lemma 2 *For every E-context E, there is a unique canonical E-context $E_c \sim_E E$, which is also the unique normal form of the terminating rewriting system $\to_E$ obtained by orienting the E-equivalence rules of Definition 5 from the left to the right*

$$E_1[\lambda x.E_2]t \to_E E_1[(\lambda x.E_2)t] \qquad \textit{if } \mathsf{FV}(t) \cap \mathsf{SV}(E_1) = \emptyset$$
$$E_1[E_2] \to_E E_1'[E_2'] \qquad \textit{if } E_i \to_E E_i' \textit{ and } E_j = E_j', \textit{ with } i, j \in \{1, 2\}, \textit{ and } i \neq j$$

Moreover, $E_c = \mathscr{E}(\eta(E))$, and therefore $E \sim_E E'$ iff $\eta(E) = \eta(E')$. $\Box$

Example 1 Let $E = E_1[\lambda x.E_2]t$ with $E_1 = (\lambda y.\Box)s$ and $E_2 = \Box$. The E-context $E_c = (\lambda y.(\lambda x.\Box)t)s$ is the unique canonical E-context $\sim_E$-equivalent to $E = (\lambda y.\lambda x.\Box)st$. $\Box$

4.2   Canonical λ-terms

The E-equivalence can be extended to terms. In the corresponding head canonical forms, along the spine, one finds first all the unmatched spine abstractions, then the E-context formed of the primary redexes, and finally the unmatched spine arguments.

Definition 6 (Surface E-equivalence on terms) The *surface E-equivalence on terms* is the least equivalence defined by the E-equivalence rules on E-contexts of Definition 5, plus

$$H_\lambda[E[\lambda x.s]] \sim_E H_\lambda[\lambda x.E[s]] \qquad\qquad \text{if } x \notin FV(E)$$
$$H_\lambda[E[s]t] \sim_E H_\lambda[E[st]] \qquad\qquad \text{if } FV(t) \cap \mathsf{SV}(E) = \emptyset$$
$$H_\lambda[E_1[s]] \sim_E H_\lambda[E_2[s]] \qquad\qquad \text{if } E_1 \sim_E E_2$$

where H_λ is a λ-context, and E, E_1, E_2 are E-contexts. The equivalence naturally extends to H-contexts, by replacing $\Box$ for s in the above equations. $\Box$

Definition 7 (head canonical λ-term) Let us say that H is a *head canonical H-context* when

$$H = H_\lambda[E_c[H_@]]$$

where H_λ is a λ-context, $H_@$ is an @-context, and E_c is a head canonical E-context. The spine λ-abstractions of H_λ are the *head λ-abstractions* of H, while the spine arguments of $H_@$ are the *head arguments* of H. The λ-term t is *head canonical* when its maximal head context is head canonical. □

Summing up, any head canonical λ-term t has the shape

$$\begin{aligned}
t &= \lambda x_1.\ldots.\lambda x_n.E_c[z t_1 \ldots t_m] \\
&= \lambda x_1.\ldots.\lambda x_n.(\lambda y_1.(\ldots(\lambda y_p.z t_1 \ldots t_m)s_p)\ldots)s_1)
\end{aligned}$$

and we can define $\#_\lambda t = n$, $\#_@ t = m$, $\eta(t) = E_c$, and $\#_p t = \#_p E_c = p$.

Every H-context H, and then every λ-term $t = H[x]$, has a unique E-equivalent head canonical form H_c, or $H_c[x]$ for terms. Moreover, as shown by Theorem 1 below, H_c preserves the same relative positions of unmatched spine λ-abstractions, unmatched spine arguments, and primary redexes of H. (A spine λ-abstraction/argument is unmatched when it is not involved in a primary redex.) More precisely, the i-th head λ-abstraction of H_c is the i-th unmatched λ-abstraction on the spine of H, the i-th head argument of the head canonical form is the i-th unmatched spine argument on the spine of H, the i-th primary redex of H_c is the i-th primary redex on the spine of H.

Theorem 1 *For any H-context H, there is a unique head canonical context $H_c \sim_E H$. More precisely,*

1. *for every H-context H, there is a unique sequence of spine variables $x_1,\ldots x_n$, a unique sequence of spine arguments $t_1,\ldots,t_m$, and a unique sequence of E-contexts $E_0, E_1,\ldots,E_{n+m}$ s.t.*

$$\begin{aligned}
H &= E_0[H_1] & \\
H_i &= \lambda x_i.E_i[H_{i+1}] & \text{for } 1 \le i \le n \\
H_{n+j} &= E_{n+j}[H_{n+j+1}]t_{m-j+1} & \text{for } 1 \le j \le m \\
H_{n+m+1} &= \square &
\end{aligned}$$

that is

$$H = E_0[\lambda x_1.E_1[\lambda x_2.E_2[\ldots[\lambda x_n.E_n[E_{n+1}[\ldots[E_{n+m-1}[E_{n+m}t_1]t_2]\ldots]t_m]]\ldots]]]$$

2. *there is a unique head canonical context $H_c \sim_E H$, and $H_c = H_\lambda[E[H_@]]$ is equal to*

$$\begin{aligned}
H_\lambda &= \lambda x_1.\ldots.\lambda x_n.\square \\
H_@ &= \square t_1 \ldots t_m \\
E_c &= \widetilde{E}_0[\widetilde{E}_1[\ldots[\widetilde{E}_{n+m}]]]
\end{aligned}$$

that is

$$H \sim_E H_c = \lambda x_1.\ldots.\lambda x_n.E_c[\square t_1 \ldots t_m]$$

where $\widetilde{E}_i = \mathscr{E}(\eta(E_i)) \sim_E E_i$ is the unique head canonical E-context equivalent to E_i;

3. *the canonical context H_c of H is the unique normal form of the rewriting system $\to_E$ obtained by orienting from the left to the right the surface E-equivalences on terms of Definition 6. Namely,*

$$\begin{aligned}
H_\lambda[E[\lambda x.s]] &\to_E H_\lambda[\lambda x.E[s]] & \text{if } x \notin FV(E) \\
H_\lambda[E[s]t] &\to_E H_\lambda[E[st]] & \text{if } FV(t) \cap SV(E) = \emptyset \\
H_\lambda[E_1[s]] &\to_E H_\lambda[E_2[s]] & \text{if } E_1 \to_E E_2
\end{aligned}$$

plus the rules for E-contexts in Lemma 2. □

4.3   E-equivalence

The surface E-equivalence permutes the arguments on the spine of a term without modifying them. The E-equivalence is obtained by recursively applying the surface E-equivalence to spine arguments too. If we denote by $\arg(t,i)$ the i-th head spine argument of the term t (which corresponds to the i-th spine argument in the head @-context of its head canonical form) and by $\arg(t,-i)$ the spine argument of the i-th primary redex of t (which corresponds to the i-th spine argument in the head canonical E-context E_c of the head canonical form of t), we define $\simeq_E$ as the least equivalence s.t. $t_1 \simeq_E t_2$ if $t_1 \sim_E t_2$, and $\arg(t_1,i) \simeq_E \arg(t_2,i)$, for $1 \leq i \leq \#_@t_1 = \#_@t_2$ or $-\#_pt_2 = -\#_pt_1 \leq i \leq -1$.

4.4   σ-equivalence

The *σ-equivalence* [8] is the least congruence induced by

$$((\lambda x.u)v)w \simeq_\sigma (\lambda x.uw)v \qquad\qquad \text{with } x \notin \mathsf{FV}(w)$$
$$(\lambda x.\lambda y.u)v \simeq_\sigma \lambda y.(\lambda x.u)v \qquad\qquad \text{with } y \notin \mathsf{FV}(v)$$

The rewriting system obtained by orienting the latter σ-equivalences from the left to the right is terminating—its head canonical forms are the same already defined for the E-equivalence—but is not confluent. Indeed, the σ-equivalence contains the E-equivalence, but it equates head canonical forms E_1 and E_2 s.t. the environments $\eta(E_1)$ and $\eta(E_2)$ are equivalent modulo the following permutation rule $t_1/x_1, t_2/x_2 \sim t_2/x_2, t_1/x_1$ if $x_1 \notin \mathsf{FV}(t_2)$ and $x_2 \notin \mathsf{FV}(t_1)$.

Example 2 Let us take the λ-term $u = E[v] = (\lambda y.\lambda x.v)st$, where E is the E-context of Example 1. Its unique head E-canonical form is $(\lambda y.(\lambda x.v)t)s$, which can be also obtained by applying the first σ-rule. However, since by applying the second σ-rule, $u \rightarrow_\sigma (\lambda x.(\lambda y.v)s)t$ too, the λ-term u has two σ-equivalent canonical forms.   □

Summing up, the E-equivalence is a variant of the σ-equivalence which equates less terms then the latter one. The definition of the σ-equivalence is simpler and more elegant, and has a direct and nice interpretation in terms of linear logic proof nets. However, the better rewriting properties of the E-equivalence—canonical form uniqueness and preservation of primary redexes relative positions—makes it more suitable for a finer analysis of reduction machines requiring a reduction at a distance based on σ-equivalence, as for instance the PAM. The σ-equivalence can be recovered from the E-equivalence by adding the following permutation equivalence of primary redexes

$$E_1[(\lambda x_1.(\lambda x_2.E_2)t_2)t_1] \sim E_1[(\lambda x_2.(\lambda x_1.E_2)t_1)t_2]$$

if $x_1 \notin \mathsf{FV}(t_2)$ and $x_2 \notin \mathsf{FV}(t_1)$, to the E-equivalence of E-contexts.

5   Linear head reduction

5.1   Linear reduction

Let $(\lambda x.t)s$ be a redex s.t. the term t contains at least one occurrence of x. For any occurrence of x in t, we can take the context C obtained by replacing such an occurrence of x with $\square$. The following reduction rule $(\lambda x.C[x])s \multimap_\beta (\lambda x.C[s'])s$, where s' is a fresh copy of s, is a linearised variant of the usual β-rule in which, instead of removing the redex after replacing all the occurrences of the bound variable x, the

redex is kept and only one occurrence of x is replaced by a fresh copy of the argument s. Such a linear β-reduction can be extended to be applied at a distance too. We obtain then the *linear reduction rule* (at a distance)

$$E[\lambda x.C[x]]s \multimap E[\lambda x.C[s']]s$$

where s' is a fresh copy of s. When the term t in $E[\lambda x.t]s$ does not contain any occurrence of x, we can instead take the following *garbage rule* (which is just a degenerated case of β-reduction at a distance)

$$E[\lambda x.t]s \to_g E[t] \qquad \text{if } x \notin \mathsf{FV}(t)$$

Given a β-redex (at a distance), by iterating the linear β-reduction (at a distance), we can eventually obtain a redex (at a distance) to which apply the *garbage rule*. Therefore, β-reduction (at a distance) can be simulated by a sort of affine reduction $\to_a$ which is the union of linear and garbage reduction.

Proposition 2 *Let $\to_a = \multimap \cup \to_g$.*

   1. If $t \to_\beta^* s$, then $t \to_a^* s$. Moreover, there is s' s.t. $t \multimap^* s' \to_g^* s$.

   2. If $t \to_a^* u$, then $u =_\beta t$. Therefore, there is $t \to_\beta^* s$ s.t. $u \to_a^* s$. □

As a consequence of the above proposition, a term has a normal form for $\to_a$ iff it has a β-normal form; moreover, the two normal forms coincide. We also remark the second part of the first item of Proposition 2. This is a particular case of a more general property stating that garbage reductions can be always postponed; that is, for every $t \to_a^* s$, there is s' s.t. $t \multimap^* s' \to_g^* s$.

5.2   Linear head β-rule

A particular case of linear reduction arises when the occurrence to be replaced is the head variable.

Definition 8 (Linear head reduction) The *linear head reduction* is the least reduction which contains the linear head β-rule

$$E[\lambda x.H[x]]s \multimap_h E[\lambda x.H[s']]s$$

where s' is a fresh copy of s, and which is closed by head contexts. □

Linear head reduction is strongly related to head β-reduction, as shown by the following statements.

Proposition 3 *Let $t \multimap_h^* s$. There is $t \to_h^* s'$ s.t. $s \to_h^* s'$.* □

Corollary 1 *A term t has a linear head normal form iff it has a head normal form. Moreover, let s be the linear head normal form of t.*

   1. The head normal form of s is obtained by β-reducing all the primary redexes in s.

   2. The head normal form of s is the head normal form of t, indeed. □

6   Strong normalisation

All the rewriting systems defined above are strong normalising on simply typed λ-terms. The proof of strong normalisation is however not at all evident. In fact, since linear reduction does not erase the reducing redex—it just replaces the occurrence of a variable by a (larger) λ-term—the size of the reducing term increases at each step. Accattoli [1], in its analysis of proof nets linear reduction, proved strong normalisation by applying reducibility candidates. Here, we show that, surprisingly, the proof

of strong normalisation of linear reduction is simpler then one might have thought, as it can be easily obtained by a trivial tuning of the proof of strong normalisation originally proposed by Gandy for β-reduction [4]. In Gandy's proof, each type τ is interpreted as a well-founded ordered set $[\tau]$. In particular, any functional type $\tau \to \sigma$ is mapped into a set of increasing functions from $[\tau]$ to $[\sigma]$. A measure is then associated to every term by interpreting any $t : \tau$ as an element $[t] \in [\tau]$. Strong normalisation is a consequence of the fact that any β-reduction $t \to_\beta s$ sends $[t]$ to a lower element $[s]$.

The original measure defined for the analysis of β-reduction does not directly work for the case of linear reduction, since such a measure does not change along linear reduction (i.e., $[t] = [s]$, when $t \multimap s$). Indeed, Gandy's measure just counts the number of λ-abstractions erased along a β-reduction. However, by taking the successor of the usual interpretation of a variable occurrence, one obtains a new measure which counts the number of variable occurrences replaced by some λ-term. Such a new measure decreases along linear reduction, and allows to prove at the same time the strong normalisation of all the rewriting systems described in the present papers.

In the following, we shall follow the presentation of Gandy's proof given by Miquel [6]. Let us interpret the base type o as the strict partial order $(\mathbb{N}, <)$, and every functional type $\tau \to \sigma$ as the strict partial order of the increasing functions from the interpretation of τ to the interpretation of σ. Formally, for every type τ, let us inductively define $([\tau], \prec_\tau)$ by

$$[\tau \to \sigma] = \{ f \in [\tau] \to [\sigma] \mid \forall v, w \in [\tau] : v \prec_\tau w \Rightarrow f(v) \prec_\sigma f(w) \}$$

$$\forall f, g \in [\tau \to \sigma] \quad : \quad f \prec_{\tau \to \sigma} g \quad \text{iff} \quad \forall v \in [\tau] : f(v) \prec_\sigma g(v)$$

with $[o] = \mathbb{N}$ and $\prec_o = <$. We define then the binary operation $+_\tau : [\tau] \times \mathbb{N} \to [\tau]$ as

$$n +_o k = n + k \qquad\qquad f +_{[\tau \to \sigma]} k = (\lambda v \in [\tau].f(v) +_\sigma k)$$

for $n, k \in \mathbb{N}$ and $f \in [\tau \to \sigma]$. It is readily seen that $v +_\tau 0 = v$, that $(v +_\tau k) + h = v +_\tau + (k + h)$, and that $k < h$ implies $v +_\tau k \prec v +_\tau h$, for every $v \in [\tau]$ and $k, h \in \mathbb{N}$.

For every type τ, let us define $\tau_* \in [\tau]$ and $\tau^* : [\tau] \to \mathbb{N}$ by

$$o_* = 0 \qquad\qquad (\tau \to \sigma)_* = (\lambda v \in [\tau].\sigma_* +_\sigma \tau^*(v))$$
$$o^*(n) = n \qquad\qquad (\tau \to \sigma)^*(f) = \sigma^*(f(\tau_*))$$

for $n \in \mathbb{N}$ and $f \in [\tau \to \sigma]$. By induction, we can see that τ^* is increasing (that is, $\tau^*(v) < \tau^*(w)$, for all $v, w \in [\tau]$ s.t. $v \prec_\tau w$).

A *valuation* is a function ϕ associating an element of $[\tau]$ to every variable $x : \tau$. Given a valuation ϕ, a variable $x : \tau$, and a value $v \in [\tau]$, we shall denote by $\phi[x \mapsto v]$ a new valuation s.t. $\phi[x \mapsto v](x) = v$, and $\phi[x \mapsto v](y) = \phi(y)$, when $y \neq x$.

Given a valuation ϕ, any typed λ-term t^τ can be interpreted as an element $[t]_\phi \in [\tau]$ by application of the following inductive definition

$$[x : \tau]_\phi = \phi(x) +_\tau 1$$
$$[\lambda x.t : \tau \to \sigma]_\phi = \lambda v \in [\tau].[t]_{\phi[x \mapsto v]} +_\sigma (\tau^*(v) + 1)$$
$$[s^{\tau \to \sigma} t : \sigma]_\phi = [s : \tau \to \sigma]_\phi [t : \sigma]_\phi$$

For every valuation ϕ, we can also define the measure $\mu_\phi : \Lambda^\to \to \mathbb{N}$, by $\mu_\phi(t^\tau) = \tau^*[t]_\phi$.

Remark 1 The only difference w.r.t. the usual interpretation used in the proof of strong normalisation of β-reduction is the interpretation of variables. Indeed, one usually takes $[x : \tau]_\phi = \phi(x)$ (see [6]). With this choice, however, we would get $[t] = [s]$ when $t \multimap s$.   □

Lemma 3 *For every valuation ϕ, every $C[x^\tau] : \sigma$, and every $t : \tau$, we have that*

1. *$[C[x^\tau]]_{\phi[x \mapsto [t]_\phi]} \prec_\sigma [C[t]]_\phi$;*
2. *if $t \to_a s$, then $[s]_\phi \prec_\sigma [t]_\phi$ and $\mu_\phi(s) < \mu_\phi(t)$.* □

By the previous lemma, and the fact that there is at least a valuation (for instance, the valuation ϕ_0 defined by $\phi_0(x^\tau) = \tau_*$), we can eventually get the strong normalisation result.

Theorem 2 *The rewriting systems $\to_a$, $\multimap$, $\to_{\beta_d}$, $\to_\beta$, $\to_h$, and $\multimap_h$ are strongly normalising.* □

7 Conclusions

In the paper we have analysed linear β-reduction in terms of a notion of evaluation context, and we have seen how a simple adaptation of the semantic proof of strong normalisation for the simply typed λ-calculus allows to prove the same result for the linear case. The proof is surprisingly simple and its idea might be adapted to prove strong normalisations of other λ-calculi in which the β-rule is decomposed in more elementary steps, as for instance in the case of explicit substitution λ-calculi.

References

[1] Beniamino Accattoli (2013): *Linear Logic and Strong Normalization*. In Femke van Raamsdonk, editor: *24th International Conference on Rewriting Techniques and Applications (RTA 2013)*, LIPIcs 21, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, pp. 39–54, doi:10.4230/LIPIcs.RTA.2013.39.

[2] Beniamino Accattoli & Delia Kesner (2010): *The Structural λ-Calculus*. In Anuj Dawar & Helmut Veith, editors: *Computer Science Logic*, LNCS 6247, Springer Berlin Heidelberg, pp. 381–395, doi:10.1007/978-3-642-15205-4_30.

[3] Vincent Danos & Laurent Regnier (2004): *Head Linear Reduction*. Http://iml.univ-mrs.fr/ regnier/articles/pam.ps.gz.

[4] R. O. Gandy (1980): *Proofs of strong normalisation*. In J. P. Seldin & J. R. Hindley, editors: *To H. B. Curry: Essays in Combinatory Logic, Lambda Calculus, and Formalism*, Academic Press, pp. 457–477.

[5] Stefano Guerrini & Giulio Pellitta (2016): *Dissecting the PAM*. Submitted.

[6] Alexandre Miquel: *A combinatorial proof of strong normalisation for the simply typed -calculus*. Unpublished draft.

[7] Pierre-Marie Pédrot & Alexis Saurin (2016): *Classical By-Need*. In Peter Thiemann, editor: *Programming Languages and Systems. 25th European Symposium on Programming, ESOP 2016*, LNCS 9632, Springer, pp. 616–643, doi:10.1007/978-3-662-49498-1_24.

[8] Laurent Regnier (1994): *Une équivalence sur les lambda-termes*. Theoretical Computer Science 126(2), pp. 281–292, doi:10.1016/0304-3975(94)90012-4.