Development Separation in Lambda-Calculus

Hongwei Xi

Boston University

Work partly funded by NSF grant CCR-0229480

Motivation for the Research

- To facilitate inductive reasoning on λ-terms
- To facilitate the encoding of λ-calculus in proof systems
  - Proofs based on structural induction on λ-terms are often amenable to such encoding

Some Notations (1)

$$\lambda\text{-terms} \quad M, N \quad ::= \quad x \mid \lambda x.M \mid M(N)$$

- $\mathbf{FV}(M)$ denotes the set of free variables in M.
- A β-redex R is a λ-term of the form $(\lambda x.M_1)(M_2)$ and its contractum is $M_1[x := M_2]$.
- $M_1 \rightarrow_\beta M_2$ stands for 1-step β-reduction: M_2 is obtained from replacing a β-redex in M_1 with its contractum.
- $M_1 \rightarrow^n_\beta M_2$ stands for n-step β-reduction ($n \geq 0$).
- $M_1 \rightarrow^*_\beta M_2$ stands for multi-step β-reduction.
- $M_1 \xrightarrow{R}_\beta M_2$ stands for 1-step β-reduction where a particular occurrence of the redex R in M_1 is reduced.

Some Notations (2)

Given a finite β-reduction sequence $\sigma : M \to_\beta^* N$, we use $\sigma(M)$ for N and $|\sigma|$ for the number be β-reduction steps in σ.

- Given $\sigma_1 : M_1 \to_\beta^* M_2$ and $\sigma_2 : M_2 \to_\beta^* M_3$, we use $\sigma_1 + \sigma_2$ for the concatenation of σ_1 and σ_2.
- $[R_1] + [R_2] + \ldots + [R_n]$ stands for a β-reduction sequence of the following form:

$$M_1 \xrightarrow{R_1}_\beta M_2 \xrightarrow{R_2}_\beta \cdots \xrightarrow{R_n}_\beta M_{n+1}$$

So each finite σ is always of the form $[R_1] + [R_2] + \ldots + [R_n]$.

Residuals of β-redexes (1)

Let $\mathcal{R}$ be a set of β-redexes in M_1. Assume $M_1 \xrightarrow{R}_\beta M_2$ for $R = (\lambda x.M)(N)$. Then this β-reduction step affects β-redexes R' in $\mathcal{R}$ in the following ways:

- R' is R. Then R' has no residual in M_2.
- R' is in N. Then all copies of R' in $M[x := N]$ are residuals of R'.
- R' is in M. Then the λ-term $R'[x := N]$ in M_2 is the residual be R'.
- R' contains R. Then the residual of R' in M_2 is the term obtained from replacing R with its contractum in R'.
- R' and R are disjoint. Then R' is not affected and is its own residual.

Residuals of β-redexes (2)

- We use $\mathcal{R}/[R]$ for the set of residuals of all β-redexes in $\mathcal{R}$ after R is reduced.
- The residual relation is transitive. Given a β-reduction sequence $\sigma = [R_1] + [R_2] + \ldots + [R_n]$, we use $\mathcal{R}/\sigma$ for the residuals of all the β-redexes in $\mathcal{R}$ under σ, which is formally defined as:

$$\mathcal{R}/\sigma = (\ldots((\mathcal{R}/[R_1])/[R_2])\ldots)/[R_n]$$

Involvement

Given $M \to_\beta^* M_1$ and $\sigma : M_1 \to_\beta M_2 \to_\beta \cdots$, we say that a redex R in M is involved in σ if a residual of R in M_1 is reduced in σ.

Developments

Let R be a set β-redexes in M. A β-reduction sequence σ from M is a development of $\mathcal{R}$ if each β-redex reduced in σ is either in $\mathcal{R}$ or a residual of some β-redex in $\mathcal{R}$. Furthermore, σ is a complete development of $\mathcal{R}$ if $\mathcal{R}/\sigma$ is empty.

Development Separation

Given a λ-term M, $M[x \ldots, x]_x$ is a representation of M in which all the free occurrences of x in M are enumerated from left to right in $[x, \ldots, x]$. We write $M[N_1, \ldots, N_n]_x$ for the λ-term obtained by substituting N_i for the ith free occurrence of x in M for $i = 1, \ldots, n$.

Let $M = (\lambda x.M_1)(M_2)$, $\mathcal{R}_1$ be a set of β-redexes in M_1, $\mathcal{R}_2$ be a set of β-redexes in M_2 and $\mathcal{R} = \{M\} \cup \mathcal{R}_1 \cup \mathcal{R}_2$. For every development σ of $\mathcal{R}$ in which M is involved, $\sigma(M)$ is of the following form:

$$\sigma_1(M_1)[\sigma_{2,1}(M_2), \ldots, \sigma_{2,n}(M_2)]_x,$$

where σ_1 is a development of $\mathcal{R}_1$ from M_1 and $\sigma_{2,i}$ are developments of $\mathcal{R}_2$ from M_2 for $i = 1 \leq i \leq n$.

Canonical and Standard Developments

Let $\sigma = [R_1] + \ldots + [R_n]$ be a development of $\mathcal{R}$. If R_j are not residuals of any β-redexes containing R_i for all $1 \leq i < j \leq n$, then σ is a *canonical* development. Furthermore, if R_j are not residuals of any β-redexes to the left of R_i for all $1 \leq i < j \leq n$, then σ is a *standard* development.

Canonicalization of Developments

For every development $\sigma : M \rightarrow_{\beta}^{*} N$ of $\mathcal{R}$, there exists a canonical development $\mathbf{cad}(\sigma) : M \rightarrow_{\beta}^{*} N$ such that for every $R \in \mathcal{R}$, R is involved in σ if it is involved in $\mathbf{cad}(\sigma)$.

Proof of Canonicalization of Developments

The proof is by structural induction on M. We present an interesting case where $M = (\lambda x.M_1)(M_2)$ and M is involved in σ. By Development Separation, $\mathbf{sep}(\sigma)$ is of the following form:

$$[M] + \sigma_1[x := N] + \sigma_{2,1} + \ldots + \sigma_{2,n}$$

where σ_1 is a development from M_1 and $\sigma_{2,i}$ are developments from M_2 for $1 \leq i \leq n$. By induction hypothesis, we can defi ne $\mathbf{cad}(\sigma)$ as follows:

$$\mathbf{cad}(\sigma) = [M] + \mathbf{cad}(\sigma_1)[x := N] + \mathbf{cad}(\sigma_{2,1}) + \ldots + \mathbf{cad}(\sigma_{2,n})$$

Clearly, $\mathbf{cad}(\sigma)$ is canonical since both $\mathbf{cad}(\sigma_1)$ and $\mathbf{cad}(\sigma_{2,i})$ are canonical for $1 \leq i \leq n$. Assume that $R \in \mathcal{R}$ is involved in $\mathbf{cad}(\sigma)$. Then it can be readily verifi ed that R is involved in $\mathbf{sep}(\sigma)$. Hence, R is also involved in σ.

Standardization of Developments

For every development $\sigma : M \to_\beta^* N$ of $\mathcal{R}$, there exists a standard development $\mathbf{std}(\sigma) : M \to_\beta^* N$ such that for every $R \in \mathcal{R}$, R is involved in σ if it is involved in $\mathbf{std}(\sigma)$.

Proof: This follows from canonicalization of developments immediately: $\mathbf{std}(\sigma)$ is obtained by a proper reshuffling of $\mathbf{cad}(\sigma)$.

Standardization Lemma

Given $\sigma = \sigma_1 + \sigma_2$, where σ_1 is a standard development of $\mathcal{R}$ and σ_2 is a standard β-reduction sequence, we can construct a β-reduction sequence $\mathbf{std}_2(\sigma_1, \sigma_2)$ which standardizes σ.

Proof of Standardization Lemma (1)

By Development Separation, we have that the function $\mathbf{std}$ is defined on all developments. Let us define $\mathbf{std}_2(\sigma_1, \sigma_2)$ and prove that $\mathbf{std}_2(\sigma_1, \sigma_2)$ standardizes $\sigma_1 + \sigma_2$ by induction on $\langle |\sigma_2|, |\sigma_1| \rangle$, lexicographically ordered. Clearly, for σ_1, σ_2, $\mathbf{std}(\sigma_1, \emptyset)$ and $\mathbf{std}(\emptyset, \sigma_2)$ can be defined as σ_1 and σ_2, respectively. We now assume $\sigma_1 = [R_1] + \sigma_1'$ and $\sigma_2 = [R_2] + \sigma_2'$, and we have two cases.

Proof of Standardization Lemma (2)

- R_2 is a residual of some β-redex in $\mathcal{R}$ that is to the left of R_1. Hence, $\sigma_1 + [R_2]$ is a development. We define $\mathbf{std}_2(\sigma_1, \sigma_2)$ as follows:

$$\mathbf{std}_2(\sigma_1, \sigma_2) = \mathbf{std}_2(\mathbf{std}(\sigma_1 + [R_2]), \sigma_2')$$

  Assume that $R \in \mathcal{R}$ is involved in $\mathbf{std}_2(\sigma_1, \sigma_2)$. Then by induction hypothesis, R is involved in $\mathbf{std}(\sigma_1 + [R_2]) + \sigma_2'$. This implies that R is involved in $\sigma_1 + [R_2] + \sigma_2' = \sigma_1 + \sigma_2 = \sigma$.

- R_2 is not a residual of any β-redex in $\mathcal{R}$ that is to the left of R_1. Then we define $\mathbf{std}_2(\sigma_1, \sigma_2)$ as follows:

$$\mathbf{std}_2(\sigma_1, \sigma_2) = [R_1] + \mathbf{std}_2(\sigma_1', \sigma_2)$$

Standard β-Reduction Sequences

Given a β-reduction sequence σ of the following form:

$$M_1 \xrightarrow{R_1}_\beta M_2 \xrightarrow{R_2}_\beta \cdots \xrightarrow{R_n}_\beta M_{n+1}$$

we say that σ is standard if for all $1 \leq i < j \leq n$, R_j is not a residual of some β-redex to the left of R_i. We say that $\sigma_s : M \to_\beta^* N$ standardizes σ if σ_s is a standard β-reduction sequence and for every R in M that is involved in σ_s, R is also involved in σ.

Standardization Theorem

For every β-reduction sequence σ, we can construct a β-reduction sequence $\mathbf{std}_1(\sigma)$ that standardizes σ.

Proof: Let us define $\mathbf{std}_1$ as follows:

$$\mathbf{std}_1(\emptyset) = \emptyset \qquad \mathbf{std}_1([R] + \sigma) = \mathbf{std}_2([R], \mathbf{std}_1(\sigma))$$

By Standardization Lemma, $\mathbf{std}_1(\sigma)$ standardizes σ.

$\mu(M)$

If M is strongly normalizing, let $\mu(M)$ be the least natural number such that $|\sigma| \leq \mu(M)$ holds for each β-reduction sequence from M. Otherwise, let $\mu(M) = \infty$.

Proposition on μ

Assume $M \xrightarrow{R}_{\beta} M'$, where $R = (\lambda x.N_1)(N_2)$ is the leftmost β-redex in M. Then $\mu(M) \leq \mu(M') + \mu(N_2)$ holds.

The proof is straightforward.

Normalization Lemma

Given a λ-term M, let $\Lambda(M)$ be the longest leftmost β-reduction sequence from M, which may be of infinite length. Assume that $\sigma : M \to_\beta^* M'$ is a standard development. If $|\Lambda(M')| < \infty$, then $|\Lambda(M')| \leq |\Lambda(M)| < \infty$ holds.

Proof of Normalization Lemma (1)

The proof proceeds by induction on $\langle|\Lambda(M')|, |\sigma|\rangle$, lexicographically ordered. If M is in β-normal form, then $M' = M$ and we are done. We now assume $M \xrightarrow{R_l}_\beta M_l$, where R_l is the leftmost β-redex in M. Then $\Lambda(M) = [R_l] + \Lambda(M_l)$. We have two cases as follows.

Proof of Normalization Lemma (2)

- R_l is involved in σ. Since σ is standard, σ is of the form $M \xrightarrow{R_l}_\beta M_l \xrightarrow{\sigma'}^*_\beta M'$ for some standard development σ'. Since $|\sigma'| < |\sigma|$ holds, we have $|\Lambda(M')| \leq |\Lambda(M_l)| < \infty$ by induction hypothesis. Hence $|\Lambda(M')| \leq |\Lambda(M)| < \infty$ holds.
- R_l is not involved in σ. Then R_l has a residual R'_l in M', which also happens to be the leftmost β-redex in M'. Then $\sigma + [R'_l]$ is a development of $\mathcal{R} \cup \{R_l\}$. Hence $\mathbf{std}(\sigma + [R'_l]) = R_l + \sigma'$ for some $\sigma' : M_l \to^*_\beta M'_l$, which is a standard development of $\mathcal{R}/[R_l]$. Assume $M' \xrightarrow{R'_l}_\beta M'_l$. Then $|\Lambda(M'_l)| < |\Lambda(M')|$ holds. By induction hypothesis, we have $|\Lambda(M'_l)| \leq |\Lambda(M_l)| < \infty$. This yields that $|\Lambda(M')| = 1 + |\Lambda(M'_l)| \leq 1 + |\Lambda(M_l)| = |\Lambda(M)| < \infty$.

Normalization Theorem

If M can be reduced to a normal form, then $|\Lambda(M)| < \infty$ holds.

Proof: By Normalization Lemma, the theorem follows from straightforward induction on the length of σ.

Related Work

- Parallel β-reductions are complete developments. Therefore, it is not surprising that the work in [Takahashi] can also be done in our setting.
- On the other hand, Takahashi's method can clearly be used to establish various lemmas in this paper (after they are properly formulated in terms of parallel β-reductions). This can probably described as *separating parallel β-reductions (from other β-reductions)*.

Conclusion

We have demonstrated some interesting uses of the *development separation* lemma, proving by structural induction on λ-terms that developments are Church-Rosser and can be standardized.

The Church-Rosser theorem in λ-calculus follows immediately. Also, we have employed the technique of development separation in establishing structurally inductive proofs for the standardization theorem, the conservation theorem and the normalization theorem.