New Notions of Reduction and Non-Semantic Proofs of β-Strong Normalization in Typed λ-Calculi*

A. J. Kfoury
kfoury@cs.bu.edu
Dept. of Computer Science
Boston University

J. B. Wells
jbw@cs.bu.edu
Dept. of Computer Science
Boston University

December 19, 1994

Boston University Computer Science Department
Technical Report 94-014

Abstract

Two new notions of reduction for terms of the λ-calculus are introduced and the question of whether a λ-term is β-strongly normalizing is reduced to the question of whether a λ-term is merely normalizing under one of the new notions of reduction. This leads to a new way to prove β-strong normalization for typed λ-calculi. Instead of the usual semantic proof style based on Girard's "candidats de réductibilité", termination can be proved using a decreasing metric over a well-founded ordering in a style more common in the field of term rewriting. This new proof method is applied to the simply-typed λ-calculus and the system of intersection types.

*This work is partly supported by NSF grant CCR–9113196.

1 Introduction

1.1 Background and Motivation.

The problem of strong normalization of β-reduction (β-SN) has been considered for various typed λ-calculi for over 25 years. Tait's proof that all λ-terms typable in the simply-typed λ-calculus (actually, Gödel's system T) are β-SN can be seen as the basis of Girard's proof of the β-SN property for higher-order typed λ-calculi, specifically, for systems **F** and $\mathbf{F}_\omega$ [Tai67, Gir71, Gir72]. Girard's method, using the so-called reducibility candidates ("candidats de réductibilité"), has been the paradigm for all later β-SN proofs for system **F** and other extensions of the simply-typed λ-calculus such as the system of positive-recursive types and the system of intersection types. Although later proofs of the β-SN property have their respective merits, they are essentially variations on Girard's original proof, simplifying or reformulating or cleaning up many of the concepts. All of these proofs rely heavily on semantics (model theory). Underlying them all is the concept of "reducibility candidate" (or "saturated set" or "type set") which is a set of strongly normalizable λ-terms (typed in some papers, untyped in others) satisfying certain closure conditions. Perhaps the least transparent part of this approach is the choice of closure conditions, which indeed vary from one proof to another, sometimes in subtle ways. (Gallier's paper [Gal90] is a useful comparative study of all proofs published until 1990.) In a very recent proof [MKO94], although the closure conditions are formally eliminated, they are still present in the guise of an evaluation function. In any case, there is a certain amount of difficulty in understanding the semantic definitions.

In this paper, we deal with the system of intersection types, so we briefly review the background of this system here. This type system was introduced by Coppo and Dezani just before 1980 [CDC80, CDCV81]. There are two important variants of the system, one of which is an extension of the other. We deal with the more basic system here which does not mention the ω type constant, for which the β-SN property has been established in papers by Pottinger [Pot80] and Leivant [Lei86] by extending Girard's original method. This system has the interesting property that the set of λ-terms typable in this system is exactly the set of λ-terms which are β-SN.

1.2 Contributions of This Paper.

For the last 25 years or so, proving the β-SN property for typed λ-calculi has been done one way—one basic method with many variations and refinements. Although there have undoubtedly been attempts to prove the β-SN property in other ways, no such results have been seen in the literature. We wish to break this trend by presenting a completely different method for proving the β-SN property. Our method is strictly proof-theoretic, relying on simple combinatorial properties of β-reduction and type-inference systems. The method consists of two parts:

1. The question of whether a λ-term is β-strongly normalizing is reduced to the question of whether the λ-term is normalizing under a new notion of reduction, $\star$-reduction.

2. The $\star$-normalization of all λ-terms typable in certain typed λ-calculi is established. For didactic reasons, we first apply our method to the simply-typed λ-calculus. The proof for simple types then extends in a simple way to the system of intersection types.

For the first part, in Section 3, we define new notions of reduction. One notion is $\star$-reduction, mentioned above, which is itself based on another new notion called γ-reduction. These notions of reduction have several implications regarding the pure untyped λ-calculus that go beyond the scope of the present report. Enough of their properties are included here in order to present our method for proving the β-SN of typed λ-calculi. Essentially, γ-reduction is a simple size-preserving transformation that reorganizes λ-bindings in a λ-term without changing the "meaning" of the λ-term. γ-reduction can be seen as "raising" a λ-abstraction outside of an enclosing β-redex. $\star$-reduction combines a β-reduction step of an I-redex followed by reduction to γ-normal form to bypass K-redexes. This behavior leads to the results that $\star$-reduction preserves the β-SN property and every $\star$-normal form is β-SN. From this, it can be seen that $\star$-normalization implies β-strong normalization. Thus, to prove the β-SN property, that *every* possible β-reduction sequence must terminate, it is sufficient to show the $\star$-normalization property, that there is *some* $\star$-reduction sequence that terminates.

For the second part, in Section 4, we show that if a λ-term M is typable in the simply-typed λ-calculus (or in the intersection-type discipline in Section 5) then we can devise a $\star$-reduction strategy from M and attach a particular well-founded partial ordering to it that guarantees the reduction strategy must terminate (normalize)—thus implying β-SN by the first part. The required reduction strategy is very simple: just reduce innermost I-redexes.

The new proof method for proving β-SN results which we present is important for several reasons. First, it is completely different from the previous methods. Nowhere does the new method involve concepts related to reducibility candidates, closure conditions, interpretation with respect to a model, validity, or soundness. Instead, the new method is very much in the style of decreasing-metric termination proofs found in the term rewriting literature. Second, we feel the new method compares well in understandability with previous proof methods. Using a decreasing metric on a well-founded order is a simple-to-understand way to show termination. In terms of proof length, this presentation includes all details; the reader is not asked to fill anything in. Although some recent proofs utilizing the reducibility candidate method have been very short, they do not seem to be any more intuitive. It would seem more appropriate to compare the length of this paper to the length of early β-SN results using the semantic methods, since this method has not yet had 25 years of work on simplifying the proof. We feel that the end result is a transparent proof, but we let the reader be the final judge on this.

1.3 Future Work.

The two typed λ-calculi for which we have carried our method all the way through give us reason for optimism regarding applying the method to other typed λ-calculi. In a sense, simple types and intersection types correspond respectively to a minimal and a maximal type discipline for which β-SN holds; every type system in use includes the simple types, while the intersection-type discipline can derive a type for every β-SN λ-term. Hence, given any other higher-order typed λ-calculus for which we are interested in proving β-SN—such as System **F** or some of its restrictions or extensions—there are two ways of proceeding. The direct way is to attach a well-founded partial ordering to a $\star$-reduction sequence, guaranteeing its termination. The indirect way is to translate an arbitrary derivation in the given type system into a derivation in the intersection-type discipline for the same untyped λ-term, without making use of the already-known fact that the type system types only β-SN λ-terms. We are particularly interested in system **F** (and certain restrictions and extensions of **F**) and in the positive-recursive-type discipline. In subsequent reports we wish to examine the β-SN property for these systems.

2 The Untyped λ-Calculus

In this section we present our definitions, notation, and nomenclature for standard concepts of the untyped λ-calculus.

2.1 λ-Terms.

The set of all λ-terms Λ is built from the countably infinite set of λ-term variables $\mathcal{V}$ using application and λ-abstraction as specified by the usual grammar $\Lambda ::= \mathcal{V} \mid (\Lambda\,\Lambda) \mid (\lambda\mathcal{V}.\,\Lambda)$. Small Roman letters (e.g. x, y, z) are used as metavariables ranging over $\mathcal{V}$ and capital Roman letters as metavariables ranging over Λ. When writing λ-terms, application associates to the left so that $MNP \equiv (MN)P$ and the scope of "$\lambda x.$" extends as far to the right as possible. We assume at all times that every λ-term M obeys the restriction that no variable is λ-bound more than once and no variable occurs both λ-bound and free in M. We assume α-conversion is used when necessary to make this happen.

As usual, $\mathrm{FV}(M)$ and $\mathrm{BV}(M)$ denote the free and λ-bound variables of a λ-term M. The expression $M[x := N]$ denotes the result of substituting N for all free occurrences of x in M, renaming λ-bound variables in M as necessary to maintain our assumptions and to avoid capturing free variables of N. A *context* $C[\;\;]$ is a λ-term with one hole (sometimes more than one) and if M is a λ-term then $C[M]$ denotes the result of inserting M into the hole in $C[\;\;]$, *including* the capture of free variables in M by the λ-bound variables of $C[\;\;]$. Unless specified otherwise, a context has only one hole. If M and N are λ-terms, then $M \equiv N$ means that M and N are identical after allowing α-conversion. $N \subset M$ denotes that N is a proper subterm of M and $N \subseteq M$ includes the possibility that $N \equiv M$.

2.2 Reduction.

Our notation on reduction generally follows Barendregt's [Bar84, § 3.1, p. 50–59] with some minor differences. A *reduction relation* $\mathcal{R}$ is a set of pairs of λ-terms. If $(M, N) \in \mathcal{R}$, then we say that M is an $\mathcal{R}$*-redex* and N is its *contractum*. If $C[\;\;]$ is a context and $(M, N) \in \mathcal{R}$, then $(C[M], C[N])$ is in the contextual closure of $\mathcal{R}$ and we say that $C[M]$ $\mathcal{R}$-reduces to $C[N]$ via the redex M and we write this as $C[M] \xrightarrow[\mathcal{R}]{M} C[N]$. If M $\mathcal{R}$-reduces to N by some unspecified redex, we write this as $M \xrightarrow[\mathcal{R}]{} N$. If $M \xrightarrow[\mathcal{R}]{} N$, we say that N is a $\mathcal{R}$*-reduct* of M. The transitive, reflexive closure of "$\xrightarrow[\mathcal{R}]{}$" is written as "$\twoheadrightarrow_{\mathcal{R}}$".

A $\mathcal{R}$*-normal form* is a λ-term containing no $\mathcal{R}$-redexes which therefore does not $\mathcal{R}$-reduce to any other term. If $M \twoheadrightarrow_{\mathcal{R}} N$ and N is a $\mathcal{R}$-normal form, we say that M $\mathcal{R}$*-normalizes* and that N is a $\mathcal{R}$-normal form of M. If M has only one $\mathcal{R}$-normal form N, this is denoted by $\mathcal{R}\text{-}nf(M) = N$ or by $M \xrightarrow[\mathcal{R}\text{-}nf]{} N$. If there are no infinite $\mathcal{R}$-reduction sequences starting from M, then M is $\mathcal{R}$*-strongly normalizing*, also written as $\mathcal{R}$-SN. M is $\mathcal{R}$*-infinite*, denoted $\mathcal{R}{-}\infty(M)$, if and only if M is not $\mathcal{R}$-SN.

The standard notion of reduction, β-reduction, is of course the least relation such that:

$$((\lambda x.P)Q) \xrightarrow[\beta]{} P[x := Q]$$

It is well-known that all λ-terms are β-confluent (Church-Rosser) and that β-normal forms are unique.

3 β-Strong Normalization and ⋆-Normalization

In this section, we introduce two new notions of reduction, γ-reduction and $\star$-reduction (a combination of γ-reduction and β-reduction), which are used throughout the rest of the paper. These notions of reduction transform λ-terms in ways that are easier to analyze than β-reduction. The main result of this section is Theorem 3.11 which implies that the question of β-strong normalization can be reduced to the question of $\star$-normalization. Subsequent sections will then show for certain typed λ-calculi that all typable terms have $\star$-normal forms, implying that all typable terms are β-strongly normalizing.

3.1 γ-Reduction and γ-Normal Forms.

Definition 3.1 γ-reduction is the least reduction relation such that:

$$((\lambda x.(\lambda y.N))P) \underset{\gamma}{\longrightarrow} (\lambda y.((\lambda x.N)P))$$

We assume that $x \neq y$ and $y \notin \mathrm{FV}(P)$, using α-conversion if necessary. In a pictorial format, this reduction looks like this:

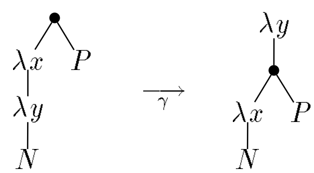

All of the standard notation for a reduction relation is used for γ-reduction.

Lemma 3.2 *For every $M \in \Lambda$:*

1. *M is γ-SN.*
2. *γ-nf(M) is unique.*

Proof: The claims are proved separately.

1. Count the number of pairs of subterm occurrences P and Q in M such that P is a β-redex, P contains Q, and there is no subterm (RS) contained within P such that Q is contained within S. It is easy to see that every γ-reduction step reduces this count. Thus, there cannot be an infinite γ-reduction sequence.

2. First, we show that γ-reduction is confluent (Church-Rosser), from which our claim follows:

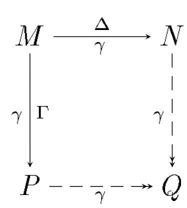

Let $\Delta \equiv ((\lambda x.(\lambda y.S))T)$ and let $M \equiv C[\Delta]$ for some context $C[\quad]$. Let $\Gamma \equiv ((\lambda w.(\lambda v.U))V)$. Consider the possible relationships between the γ-redexes Δ and Γ.

(a) $\Gamma \equiv \Delta$. The result is immediate.

(b) $\Gamma \subseteq \Delta$ or $\Delta \subseteq \Gamma$. If this is the case, then one of $\Gamma \subseteq S$ or $\Gamma \subseteq T$ or $\Delta \subseteq U$ or $\Delta \subseteq V$ holds. Since all of these case are handled identically, suppose that $\Gamma \subseteq S$. Let $S \equiv D[\Gamma]$ and let Γ' be the γ-contractum of Γ. Clearly:

$$\begin{aligned} M &\equiv C[((\lambda x.(\lambda y.D[\Gamma]))T)] \\ N &\equiv C[(\lambda y.((\lambda x.D[\Gamma])T))] \\ P &\equiv C[((\lambda x.(\lambda y.D[\Gamma']))T)] \end{aligned}$$

If we define Q as follows:

$$Q \equiv C[(\lambda y.((\lambda x.D[\Gamma'])T))]$$

then it is easy to check that $N \xrightarrow[\gamma]{\Gamma} Q$ and that $P \xrightarrow[\gamma]{} Q$.

(c) $\Gamma \not\subseteq \Delta$ and $\Delta \not\subseteq \Gamma$. Let $M \equiv C[\Delta, \Gamma]$ where $C[\ ,\ ]$ is a context with two holes. Let Γ' be the γ-contractum of Γ and Δ the γ-contractum of Δ. Clearly, $P \equiv C[\Delta, \Gamma']$ and $N \equiv C[\Delta', \Gamma]$. If we let $Q \equiv C[\Delta', \Gamma']$, then it is easy to check that $N \xrightarrow[\gamma]{\Gamma} Q$ and that $P \xrightarrow[\gamma]{\Delta} Q$.

■

It is easy to give an inductive definition of those λ-terms which happen to be in γ-normal form. The set Λ^γ of γ-normal forms is defined inductively as follows:

1. $x \in \Lambda^\gamma$ if $x \in \mathcal{V}$.
2. $(MN) \in \Lambda^\gamma$ if $M, N \in \Lambda^\gamma$ and M is not a λ-abstraction.
3. $(\lambda x.M) \in \Lambda^\gamma$ if $M \in \Lambda^\gamma$ and $x \in \mathcal{V}$.
4. $((\lambda x.M)N) \in \Lambda^\gamma$ if $M, N \in \Lambda^\gamma$, M is not a λ-abstraction, and $x \in \mathcal{V}$.

Lemma 3.3 *A λ-term M is in γ-nf if and only if $M \in \Lambda^\gamma$.*

Proof: The two directions are proved separately.

$\Leftarrow$ If $M \in \Lambda^\gamma$, then M can not contain a γ-redex, by induction on the definition of M.

$\Rightarrow$ Suppose M is in γ-*nf* and $M \notin \Lambda^\gamma$ and then derive a contradiction. Let N be a least subterm of M such that $N \notin \Lambda^\gamma$. Since N is a least such subterm, every proper subterm of N must belong to Λ^γ. For every possible shape of N, we derive a contradiction.

1. Obviously, N can not be a variable because that would contradict part 1 of the definition of Λ^γ.
2. N can not be of the form $((\lambda x.\lambda y.P)Q)$ because then N would be a γ-redex, contradicting the fact that M is in γ-*nf*.
3. It is also impossible for N to be of the form (PQ) (respectively $((\lambda x.P)Q)$) where P is not a λ-abstraction. If this were the case, then since $P \in \Lambda^\gamma$ and by part 2 (respectively part 4) of the definition of Λ^γ, N would be in Λ^γ, a contradiction.
4. N can not be of the form $(\lambda x.P)$ because the fact that $P \in \Lambda^\gamma$ would mean that $N \in \Lambda^\gamma$ by part 3 of the definition of Λ^γ.

■

A subterm occurrence N in M is *passive* if $N \equiv M$ or N occurs as $(PN) \subset M$ for some P or N occurs as $(\lambda x.N)$ where $(\lambda x.N)$ is passive, otherwise N is *active*. (Note that this definition is different from [Bar84, § 2.1.8 (iv), p. 25].) Generally, if there is a λ-abstraction $(\lambda z.Z)$ in X and $X \xrightarrow[\gamma]{} Y$ and there is a λ-abstraction $(\lambda z.Z')$ in Y, we say that $(\lambda z.Z)$ and $(\lambda z.Z')$ are the same λ-abstraction, even though the bodies Z and Z' may be different.

Lemma 3.4 *γ-reduction has these properties:*

1. *λ-abstractions are neither destroyed nor introduced.*
2. *If a λ-abstraction is the function of a β-redex before a γ-reduction step, it is still the function of a β-redex after the γ-reduction step.*
3. *If a λ-abstraction becomes the function of a fresh β-redex after a γ-reduction step, then it was active before the γ-reduction step.*
4. *If a λ-abstraction is passive before a γ-reduction step, then it is still passive after the γ-reduction step.*
5. *If a λ-term is in γ-normal form, then all of its active λ-abstractions are functions of β-redexes.*

Proof: Each property is proved separately.

1. Obvious from the definition.
2. Obvious from the definition.
3. There is only one case when a λ-abstraction which was not part of a β-redex before a γ-reduction step becomes part of a β-redex after the γ-reduction step. This can only happen when a γ-redex $\Gamma \equiv ((\lambda x.(\lambda y.M))N)$ occurs as the function of an application as in (ΓP). In this case, the λ-abstraction which becomes part of a β-redex was $(\lambda y.M)$ which was in an active position.
4. When a γ-redex $\Gamma \equiv ((\lambda x.(\lambda y.M))N)$ in the λ-term $C[\Gamma]$ is reduced, the subterm M is active both before and after the γ-reduction step, the subterm N is passive both before and after, the abstraction over x is active both before and after, the abstraction over y is active before, and all subterms not inside Γ retain their active or passive status as do all subterms inside M or N.
5. Suppose $M \in \Lambda^\gamma$ and $(\lambda y.N) \subset M$ is active but not the function of a β-redex. Since $(\lambda y.N)$ is active, it must occur either as $((\lambda y.N)P)$ or as $(\lambda x.(\lambda y.N))$ where $(\lambda x.(\lambda y.N))$ is active. The first alternative, $((\lambda y.N)P)$, is impossible since $(\lambda y.N)$ is not the function of a β-redex. By induction on the number of enclosing abstractions, we prove that the second alternative, occuring as $(\lambda x.(\lambda y.N))$ which is active, is also impossible. In the base case, there may not be any enclosing abstractions, so $(\lambda x.(\lambda y.N))$ is immediately impossible. In the induction case, there may be $n+1$ enclosing abstractions around $(\lambda y.N)$, so there may be only n enclosing abstractions around $(\lambda x.(\lambda y.N))$. By induction, it is then the case that since $(\lambda x.(\lambda y.N))$ is active it must be the function of a β-redex, so it occurs as $((\lambda x.(\lambda y.N))P)$. However, this is a γ-redex which is impossible since M is in γ-normal form.

■

3.2 γ-Reduction and β-Strong Normalization.

In this subsection, we show that if the result of γ-reduction is β-SN, then the input must also have been β-SN. (We could show that γ-reduction preserves the β-SN property, but we will not need such a general result later.) To reach this result we will need some auxiliary lemmas and a method for keeping track of the residuals of both β-redexes and γ-redexes under both β-reduction and γ-reduction.

To keep track of a β-redex $\Delta \equiv ((\lambda x.P)Q)$ and its residuals relative to β-reduction as well as γ-reduction, we mark its leading λ with a subscripted index $i \in \mathbb{N}$. For example, the marked β-redex Δ is written as $((\lambda_i x.P)Q)$. The notation we use is in the style of [Bar84, § 11.1.2, p. 279 and § 11.2.4, p. 284].

It is also necessary to keep track of γ-redexes and their residuals relative to both β-reduction and γ-reduction. For this, we also mark the leading λ of the γ-redex with an index $j \in \mathbb{N}$, but this time in superscript position. For example, for the γ-redex $((\lambda x.\lambda y.N)P)$, the marked version will be written as $((\lambda^j x.\lambda y.N)P)$. It will be possible for the same λ to be marked as part of both a β-redex and a γ-redex, in which case it will have both a subscript and a superscript.

The set $\Lambda^\sharp$ of marked terms is defined inductively as follows:

1. $x \in \Lambda^\sharp$ if $x \in \mathcal{V}$.
2. $(MN) \in \Lambda^\sharp$ if $M, N \in \Lambda^\sharp$.
3. $(\lambda x.M) \in \Lambda^\sharp$ if $M \in \Lambda^\sharp$ and $x \in \mathcal{V}$.
4. **Marked β-redex:** $((\lambda_i x.M)N) \in \Lambda^\sharp$ if $M, N \in \Lambda^\sharp$, $x \in \mathcal{V}$, and $i \in \mathbb{N}$ is a fresh index.
5. **Marked γ-redex:** $((\lambda^i x.\lambda y.M)N) \in \Lambda^\sharp$ if $M, N \in \Lambda^\sharp$, $x, y \in \mathcal{V}$, and $i \in \mathbb{N}$ is a fresh index.
6. **Simultaneously marked β-redex and γ-redex:** $((\lambda_i^j x.\lambda y.M)N) \in \Lambda^\sharp$ if $M, N \in \Lambda^\sharp$, $x, y \in \mathcal{V}$, and $i, j \in \mathbb{N}$ are fresh indices.

If $M \in \Lambda^\sharp$ then $|M|$ denotes the term in Λ resulting from erasing all indices in M.

The notions of reduction β and γ are extended to marked terms in the following manner. The notation $[i]$ means the index i may or may not be present, but if it is present in one occurrence of $[i]$ it is present in all others.

$$\begin{array}{rcl} ((\lambda_{[i]}^{[j]} x.M)N) & \xrightarrow[\beta]{} & M[x := N] \\ ((\lambda_{[i]}^{[j]} x.(\lambda y.M))N) & \xrightarrow[\gamma]{} & (\lambda y.((\lambda_{[i]} x.M)N)) \end{array}$$

It is important to notice that β-reduction of a redex that is both a marked β-redex and a marked γ-redex will erase both markings, while γ-reduction will erase only the marking of the γ-redex, preserving the marking of the β-redex.

Let M, N be terms in Λ, Δ be a β-redex (respectively γ-redex) occurrence in M, Γ be a β-redex (respectively γ-redex) occurrence in N, and σ a $\beta\gamma$-reduction from M to N:

$$\sigma : M \xrightarrow[\beta\gamma]{}\!\!\!\!\rightarrow N$$

The redex Γ is a *residual* of the redex Δ (relative to the reduction σ) if σ can be *lifted* to a reduction σ' from $M' \in \Lambda^\sharp$ to $N' \in \Lambda^\sharp$ such that:

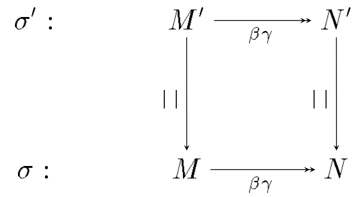

where Δ is the only marked redex in M' (with some index $i \in \mathbb{N}$) and Γ is one of the marked redexes in N' (with the same index i).

Lemma 3.5 *For every $M \in \Lambda$ it is the case that:*

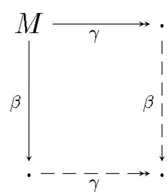

Proof: Let $\Gamma \equiv ((\lambda x.(\lambda y.N))P)$ be a γ-redex occurrence in M. Let Δ be a β-redex occurrence in M. Consider the different possible relationships between Γ and Δ.

1. If $\Delta \equiv \Gamma$, then this is the case:

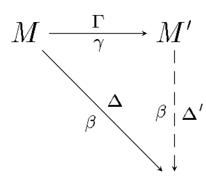

   where $\Delta' \equiv ((\lambda x.N)P)$ is the residual of Δ in M'.

2. If either $\Delta \subset \Gamma$ (in which case either $\Delta \subset N$ or $\Delta \subset P$) or $\Delta \not\subset \Gamma$ and $\Gamma \not\subset \Delta$, then it must be the case that:

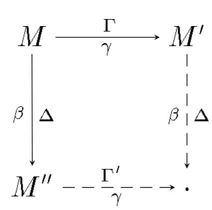

   where Γ' is the residual of Γ in M''. Notice that the the residual of Δ in M' is exactly Δ and that when $\Gamma \not\subset \Delta$ it is the case that $\Gamma \equiv \Gamma'$.

3. If $\Gamma \subset \Delta$, then this must be the case:

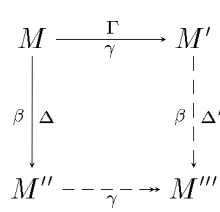

where Δ' is the residual of Δ in M' and the reduction from M'' to M''' reduces all residuals of Γ in M''.

Thus, for all cases it holds that:

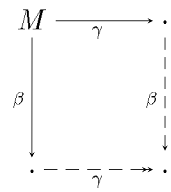

Diagram chasing then produces the desired conclusion:

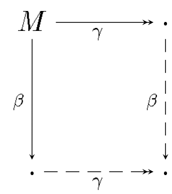

■

Lemma 3.6 *For every $M \in \Lambda$, if γ-$nf(M)$ is β-SN then M is β-SN.*

(We claim that the converse of Lemma 3.6 is also true. However, it requires a more subtle argument and it is not needed in this paper.)

Proof: Let $M' = \gamma\text{-}nf(M)$. We now prove that if $\beta\text{-}\infty(M)$ then $\beta\text{-}\infty(M')$, which is logically equivalent to the claim of the lemma. Suppose σ were an infinite β-reduction from M. Using Lemma 3.5 allows erecting an infinite β-reduction σ' from M':

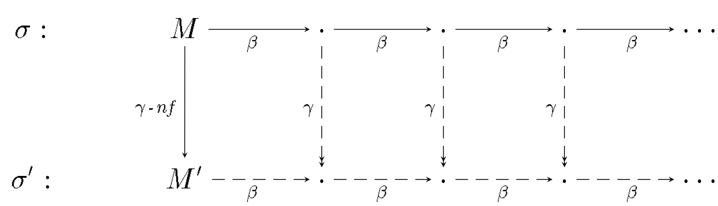

■

3.3 Preservation of β-Strong Normalization by βI-Reduction.

When β-reduction is restricted to I-redexes, we call it βI-reduction. It is already known that βI-reduction preserves the β-SN property. We present here the necessary observations to make use of this known result.

Let $\Delta \equiv ((\lambda x.P)Q)$. If $x \in \mathrm{FV}(P)$ then Δ is an *I-redex*. Otherwise, if $x \notin \mathrm{FV}(P)$ then Δ is a *K-redex*. (Following [Bar84, § 11.3.6, p. 296].)

Lemma 3.7 *Let $M \xrightarrow[\beta]{\Delta} N$ where Δ is an I-redex. Then M is β-SN if and only if N is β-SN.*

Proof: The two directions of the equivalence are proved separately.

$\Rightarrow$ Immediate.

$\Leftarrow$ This is true if and only if the contrapositive is true: β-$\infty(M)$ implies β-$\infty(N)$. This is exactly the statement of the Conservation Theorem [Bar84, § 13.4.12, p. 343].

■

3.4 $\star$-Reduction.

In this subsection, we define $\star$-reduction, a combination of γ-reduction and β-reduction. We then prove the major result of Theorem 3.11, showing that the question of β-strong normalization can be reduced to the question of $\star$-normalization. The importance of this result is the fact that it is easier to prove a normalization result than a strong normalization result, because the reduction strategy can be chosen.

Definition 3.8 For two terms $M, N \in \Lambda^\gamma$, we define $M \xrightarrow[\star]{} N$ to hold if there is an I-redex Γ in M and a term $M' \in \Lambda$ such that:

$$M \xrightarrow[\beta]{\Gamma} M' \xrightarrow[\gamma\text{-}nf]{} N$$

All of the standard notation for a reduction relation is used for $\star$-reduction.

Lemma 3.9 *For all $M, N \in \Lambda^\gamma$, if $M \twoheadrightarrow_\star N$ and N is β-SN, then M is also β-SN.*

Proof: It is sufficient to show the claim for a single $\star$-reduction step: if $M \xrightarrow[\star]{} N$ and N is β-SN, then M is also β-SN. This is a consequence of Lemma 3.7 (for the reduction $M \xrightarrow[\beta]{\Delta} M'$ where Δ is an I-redex) and Lemma 3.6 (for the reduction $M' \xrightarrow[\gamma\text{-}nf]{} N$). ■

Lemma 3.10 *Let $M \in \Lambda^\gamma$ be a term containing no I-redexes. Then M is β-SN.*

Proof: By induction on the number of β-redexes in M. For the base case where there are no redexes, the result is immediate. For the induction step, let $\{\Delta_1, \ldots, \Delta_{n+1}\}$ be the set of all β-redex occurrences in M. All of these redexes are K-redexes. Let $\Delta_{n+1} \equiv ((\lambda x.P)Q)$ where x does not occur in P. Assume for some $j \in \{1, \ldots, n\}$ it is the case that $\Delta_i \subset Q$ if and only if $j < i \leq n$. Consider the β-reduction step $M \xrightarrow[\beta]{\Delta_n} N$. We now show that the set of all β-redexes in N is exactly $\{\Delta'_1, \ldots, \Delta'_j\}$ where for $1 \leq i \leq j$ it is the case that Δ'_i is the residual of Δ_i and that $N \in \Lambda^\gamma$. Let $M \equiv C[\Delta_{n+1}]$ for some context $C[\quad]$ with exactly one hole. It is clear that $N \equiv C[P]$. Because M is in γ-normal form, P can not be a λ-abstraction, and since x does not occur in P, no new β-redex is formed by the reduction. Any β-redex that occurred inside Q was discarded and all others were kept but not duplicated, so the set of remaining β-redexes is exactly as claimed. Since N contains only $j < n+1$ β-redexes and since $N \in \Lambda^\gamma$, we can assume by induction that N is β-SN. Since the β-redex Δ_{n+1} which we reduced was picked arbitrarily and no infinite β-reduction can follow from its reduction, this proves that M is β-SN. ■

Theorem 3.11 *For any term $M \in \Lambda$, if γ-nf(M) is $\star$-normalizing (there is at least one $\star$-reduction from γ-nf(M) which terminates) then M is β-SN.*

(We claim the converse of Theorem 3.11 is also true, but do not prove it and do not need it.)

Proof: If γ-*nf*(M) is $\star$-normalizing, then by Lemma 3.9 and Lemma 3.10 (since a $\star$-normal form belongs to Λ^γ and has no I-redexes) it holds that γ-*nf*(M) is β-SN. By Lemma 3.6 we conclude that M is β-SN. ■

4 $\star$-Normalization of the Simply-Typed $\boldsymbol{\lambda}$-Calculus

In this section, we prove that every simply-typed λ-term is β-SN. This is a new proof for a well-known result and it is probably not any simpler than many of the other proofs in the literature already. The novelty of this proof is that the argument does not depend on semantic notions such as models, interpretations, proofs of soundness, etc. Instead, this proof is a decreasing-metric termination proof of the style more frequently seen in the field of term rewriting. In Section 5, we will generalize this proof to the more complicated intersection-type discipline.

4.1 The Simply-Typed $\boldsymbol{\lambda}$-Calculus.

In this paper, it is convenient to define the simply-typed λ-calculus in an explicitly-typed manner, where every subterm and bound variable of a typed λ-term is annotated with an explicit type, written in superscript position for convenience. (This can be called "Church" style.) For example, one simple typing of $((\lambda x.x)(\lambda y.y))$ might be written as:

$$((\lambda x^{\alpha\to\alpha}.x^{\alpha\to\alpha})^{(\alpha\to\alpha)\to(\alpha\to\alpha)}(\lambda y^{\alpha}.y^{\alpha})^{\alpha\to\alpha})^{\alpha\to\alpha}$$

Let $\Lambda^{\to}$ be the set of simply-typed λ-terms.

The set of simple types $\mathbb{T}$ is built from the countably infinite set of type variables $\mathbb{V}$ using the "$\to$" type constructor as specified by the grammar $\mathbb{T} ::= \mathbb{V} \mid (\mathbb{T} \to \mathbb{T})$. A type is therefore either a *type variable* or a $\to$*-type*. Small Greek letters from the beginning of the alphabet (e.g. α, β, γ, δ) are metavariables over $\mathbb{V}$ and small Greek letters towards the end of the alphabet (e.g. σ and τ) are metavariables over $\mathbb{T}$. When writing types, the arrows associate to the right so that $\sigma \to \tau \to \rho = \sigma \to (\tau \to \rho)$.

The λ-term variables are pairs of untyped variables and types, written as x^{σ}, y^{τ}, z^{ρ}, and so on. Instead of using type assignments (sometimes called contexts or environments), we require every typed λ-terms M to satisfy the property that:

(†) For all $x^{\sigma}, y^{\tau} \in \mathrm{FV}(M) \cup \mathrm{BV}(M)$, if $x = y$ then $\sigma = \tau$

The set $\Lambda^{\to}$ of simply-typed λ-terms and a type-erasing function $|\quad|$ from $\Lambda^{\to}$ to Λ are defined inductively as follows:

1. $x^{\sigma} \in \Lambda^{\to}$ and $|x^{\sigma}| = x$ if $x \in \mathcal{V}$ and $\sigma \in \mathbb{T}$.
2. $(M^{\sigma\to\tau}N^{\sigma})^{\tau} \in \Lambda^{\to}$ and $|(M^{\sigma\to\tau}N^{\sigma})^{\tau}| = (|M^{\sigma\to\tau}||N^{\sigma}|)$ if $M^{\sigma\to\tau} \in \Lambda^{\to}$ and $N^{\sigma} \in \Lambda^{\to}$.
3. $(\lambda x^{\sigma}.M^{\tau})^{\sigma\to\tau} \in \Lambda^{\to}$ and $|(\lambda x^{\sigma}.M^{\tau})^{\sigma\to\tau}| = (\lambda x.|M^{\tau}|)$ if $M^{\tau} \in \Lambda^{\to}$ and $\sigma \in \mathbb{T}$.

Provided all of the free and bound variables in a λ-term are annotated with types, the types annotating applications and λ-abstractions may be omitted with no loss of information.

We choose to present the simply-typed λ-terms in a "Church" style rather than a "Curry" style partly because this gives a natural interpretation for β-reduction and γ-reduction. In the Church style, using the natural extension of β-reduction to the simply-typed λ-calculus, a β-reduct or a γ-reduct of M automatically inherits a simple-typing from M. This will prove to be vital for our purposes. In the Curry style, if M is typable and $M \xrightarrow[\beta]{} N$, then N is also typable, but a mechanism must be defined to construct the typing for N from the typing for M. Also, if $M \xrightarrow[\beta]{} N$, the reduct N may have typings that are not necessarily derived from the typing for M.

We now define explicitly how β-reduction and γ-reduction work on simply-typed λ-terms. If $M \in \Lambda^{\rightarrow}$ and Δ is a β-redex occurrence in M such that

$$\Delta \equiv ((\lambda x^{\sigma}.P^{\tau})^{\sigma \rightarrow \tau} Q^{\sigma})^{\tau}$$

and $M \equiv C[\Delta]$ where $C[\quad]$ is a context with exactly one hole, then:

$$M \xrightarrow[\beta]{\Delta} C[P^{\tau}[x^{\sigma} := Q^{\sigma}]]$$

If Γ is a γ-redex occurrence in M such that

$$\Gamma \equiv ((\lambda x^{\sigma}.(\lambda y^{\rho}.N^{\tau})^{\rho \rightarrow \tau})^{\sigma \rightarrow \rho \rightarrow \tau} P^{\sigma})^{\rho \rightarrow \tau}$$

and $M \equiv D[\Gamma]$ where $D[\quad]$ is a context with exactly one hole, then:

$$M \xrightarrow[\gamma]{\Gamma} D[(\lambda y^{\rho}.((\lambda x^{\sigma}.N^{\tau})^{\sigma \rightarrow \tau} P^{\sigma})^{\tau})^{\rho \rightarrow \tau}]$$

4.2 A Metric on Simply-Typed $\boldsymbol{\lambda}$-Terms.

The proof for $\star$-normalization later in this section uses a metric $order^{\bullet}$ on λ-terms that decreases after each reduction step. This metric is defined on the types involved in the I-redexes in the λ-term.

For simple types, define the function *order* inductively as follows:

1. $order(\alpha) = 0$ where $\alpha \in \mathbb{V}$ is a type variable.
2. $order(\sigma \rightarrow \tau) = \max\{1 + order(\sigma), order(\tau)\}$.

Let Δ be a simply-typed β-redex so that:

$$\Delta \equiv ((\lambda x^{\sigma}.P^{\tau})^{\sigma \rightarrow \tau} Q^{\sigma})^{\tau}$$

Define $order(\Delta) = order(\sigma)$. Let M be a simply-typed λ-term. Let the set of all I-redex occurrences in M be $\{\Delta_1, \ldots, \Delta_n\}$. Define the function $order^{\bullet}$ from λ-terms to multisets over $\mathbb{N}$ so that:

$$order^{\bullet}(M) = \{order(\Delta_1), \ldots, order(\Delta_n)\}$$

Thus, for any λ-term M, $order^{\bullet}(M)$ is a finite multiset of natural numbers. Observe that K-redexes do not contribute to the value of $order^{\bullet}(M)$.

4.3 A Well-Founded Multiset Ordering.

Since the metric $order^{\bullet}$ computes multisets of natural numbers instead of just single natural numbers, we can not use the simple, numeric "$<$" ordering. However, there is a standard multiset ordering which is suitable.

For S a multiset over $\mathbb{N}$ and $n \in \mathbb{N}$, let $\#(S, n)$ denote the number of occurrences of n in S. For S and T finite multisets over $\mathbb{N}$, define $S \succ T$ if and only if both:

1. $S \neq T$.

2. For $n \in \mathbb{N}$, if $\#(S,n) < \#(T,n)$ then there is some $m > n$ such that $\#(S,m) > \#(T,m)$.

In plain English, if one starts with a finite multiset, removes some numbers from this multiset and replaces each of them with any finite number of strictly smaller numbers, then the result is defined to be smaller than the starting point.

Lemma 4.1 *The ordering "$\succ$" is well-founded, i.e. there cannot be an infinite descending chain* $S \succ S_1 \succ S_2 \succ \cdots \succ S_i \succ \cdots$.

Proof: See [DM79]. ■

4.4 A Normalizing $\star$-Reduction Strategy.

We now prove that a particular $\star$-reduction strategy terminates for all simply-typable λ-terms, implying all such terms are β-SN. For a λ-term M, the I-redex Δ is *innermost* if it does not properly contain another I-redex (but Δ may contain K-redexes). Every λ-term with one or more I-redexes contains at least one innermost I-redex.

The reduction notion $\star$ is defined so far only for untyped λ-terms in γ-normal form (members of Λ^γ). Since $\star$-reduction is stated in terms of β-reduction and γ-reduction, it is obvious how to extend it to any simply-typed λ-term M such that $|M| \in \Lambda^\gamma$ and we do so now. Denote the set of all simply-typed λ-terms M such that $|M| \in \Lambda^\gamma$ by $(\Lambda^{\rightarrow})^\gamma$. If $M, N \in (\Lambda^{\rightarrow})^\gamma$ and $\Delta \subset M$ is an I-redex, write $M \xrightarrow[\star]{\Delta} N$ to mean $M \xrightarrow[\beta]{\Delta} M' \xrightarrow[\gamma\text{-}nf]{} N$ for some $M' \in \Lambda^{\rightarrow}$.

Lemma 4.2 *If $M, N \in (\Lambda^{\rightarrow})^\gamma$ and $M \xrightarrow[\star]{\Delta} N$ where Δ is an innermost I-redex, then $order^\bullet(M) \succ order^\bullet(N)$.*

Proof: Let $\Delta \equiv ((\lambda v^\sigma.S^\tau)^{\sigma\rightarrow\tau} P^\sigma)^\tau$ be an innermost I-redex. Let $M \equiv C[\Delta]$ for some context $C[\ \ ]$. Let the set of all I-redex occurrences in M be $\{\Delta_1, \ldots, \Delta_q\}$ with Δ the same subterm occurrence as Δ_1. Let $M' \equiv C[S^\tau[v^\sigma := P^\sigma]]$. It is clear that $M \xrightarrow[\beta]{\Delta} M'$ and that (by the given part of the claim) $M' \xrightarrow[\gamma\text{-}nf]{} N$. Note that M' is in $\Lambda^{\rightarrow}$ but may not be in $(\Lambda^{\rightarrow})^\gamma$.

Consider the residuals in N of the I-redexes in M. Since Δ is an innermost I-redex, P contains no I-redexes that might be duplicated by the β-reduction step from M to M'. By the definition of γ-reduction, the γ-reduction steps from M' to N can not duplicate or remove any β-redexes. Thus, for $2 \leq i \leq q$, let Δ_i' be the single residual of Δ_i in N. For $2 \leq i \leq q$ it is clear that $order(\Delta_i) = order(\Delta_i')$, because neither the β-reduction step nor the γ-reduction steps could have changed the type assigned to the bound variable of Δ_i.

Given what we have shown so far, to prove that $order^\bullet(M) \succ order^\bullet(N)$ it is sufficient to show that for each new I-redex introduced by the β-reduction step or the subsequent γ-reduction steps, the value of *order* on this new I-redex is smaller than $order(\Delta) = order(\sigma)$. In fact, we will show the stronger claim that this is the case for all new β-redexes, both I-redexes and K-redexes.

Consider separately the simple case where P is not a λ-abstraction or every occurrence of v in S is passive. In this case, it is easy to see that both:

1. M' is in γ-normal form and belongs to $(\Lambda^{\rightarrow})^\gamma$ and thus $M' \equiv \gamma\text{-}nf(M') \equiv N$.

2. No new I-redex occurrence is "created" by the β-reduction step, i.e. the I-redex occurrences in $M' \equiv N$ are exactly $\{\Delta'_2, \ldots, \Delta'_q\}$.

Keep in mind that, because M is in γ-normal form, S is not a λ-abstraction and, thus, S can not become the function of a β-redex. Thus, it is clear that $\mathit{order}^\bullet(N)$ is exactly $\mathit{order}^\bullet(M)$ with one occurrence of $\mathit{order}(\Delta_1)$ removed, which implies that $\mathit{order}^\bullet(M) \succ \mathit{order}^\bullet(N)$.

Now consider the more complicated case where P is a λ-abstraction *and* there are active occurrences of v in S. Let P^σ be of the form:

$$(\lambda x_1^{\varphi_1}. \ldots .\lambda x_m^{\varphi_m}.Q^\psi)^\sigma$$

where Q^ψ is not a λ-abstraction. For each active occurrence of v in S, one or more β-redexes will be formed by the $\star$-reduction step. Usually, some of these β-redexes will be formed by the β-reduction step from M to M' and some will be formed by the γ-reduction steps from M' to N.

First, we show that for each β-redex Γ formed by the β-reduction step that $\mathit{order}(\Gamma) < \mathit{order}(\Delta) = \mathit{order}(\sigma)$. Wherever v occurs as

$$(v^{\varphi_1 \to \cdots \to \varphi_m \to \psi} R^{\varphi_1})$$

for some subterm R^{φ_1}, the β-reduction step will form the β-redex:

$$\Gamma \equiv ((\lambda x_1^{\varphi_1}. \ldots .\lambda x_m^{\varphi_m}.Q^\psi)^{\varphi_1 \to \cdots \to \varphi_m \to \psi} R^{\varphi_1})$$

No other kind of β-redex can be formed by the β-reduction step. By definition of *order* on a β-redex, $\mathit{order}(\Gamma) = \mathit{order}(\varphi_1)$. By the definition of *order* on a type, it is easy to see that $\mathit{order}(\varphi_1)$ is at least 1 smaller than $\mathit{order}(\sigma)$.

Now, we show that for each β-redex Γ formed by the γ-reduction steps from M' to N that $\mathit{order}(\Gamma) < \mathit{order}(\Delta) = \mathit{order}(\sigma)$. To do this, first we analyze the λ-abstractions which may become involved in fresh β-redexes. By Lemma 3.4, we know that if a β-redex is formed by one of the γ-reduction steps from M' to N, the function of the new β-redex must be a λ-abstraction that existed in M', which was not already part of a β-redex, and which was active in M'. Since M was in γ-normal form, all active λ-abstractions in M were already the functions of β-redexes. The only active abstractions in M' which are not functions of β-redexes are the outermost abstractions of the copies of P^σ wherever an active occurrence of v^σ was replaced. Thus, for each fresh β-redex Γ, the function of Γ must be a copy of one of the outermost abstractions of P^σ. Recall the definition in Subsection 4.1 of how γ-reduction works on simply-typed λ-terms. When a γ-reduction step rearranges λ-abstractions, the types assigned to the bound variables move with them. The types assigned to the bound variables of the outermost abstractions of P^σ are exactly $\varphi_1, \ldots, \varphi_m$. For $1 \leq i \leq m$ it holds that $\mathit{order}(\varphi_i) < \mathit{order}(\sigma)$ (by the definition of *order* on a type). Thus, for each fresh β-redex Γ it holds that $\mathit{order}(\Gamma) < \mathit{order}(\Delta)$. ■

Lemma 4.3 *If $M \in (\Lambda^\to)^\gamma$ then M is $\star$-normalizing.*

Proof: Suppose there is no $\star$-reduction sequence from M which reaches $\star$-normal form. Then there is an infinite $\star$-reduction of the form:

$$M \xrightarrow[\star]{\Delta_0} M_1 \xrightarrow[\star]{\Delta_1} M_2 \xrightarrow[\star]{\Delta_2} \cdots$$

where each redex Δ_0, Δ_1, Δ_2, ..., is an innermost I-redex. By Lemma 4.2, there must be this infinite sequence:

$$order^\bullet(M) \succ order^\bullet(M_1) \succ order^\bullet(M_2) \succ \cdots$$

This contradicts Lemma 4.1, so there must be a $\star$-reduction sequence from M which reaches $\star$-normal form. ■

We believe that Lemma 4.3 could be stated more strongly: If $M \in (\Lambda^{\rightarrow})^{\gamma}$ then M is $\star$-SN. However, we do not need such a strong result to prove the next theorem.

Theorem 4.4 *If $M \in \Lambda^{\rightarrow}$ then M is β-SN.*

Proof: Let $N \equiv \gamma\text{-}nf(M)$. Since N is in $(\Lambda^{\rightarrow})^{\gamma}$, N is $\star$-normalizing by Lemma 4.3. Thus, by Theorem 3.11, M is β-SN. ■

5 $\star$-Normalization of the Intersection-Type Discipline

In this section, we prove that every λ-term typable in the system of intersection types $\star$-normalizes and is therefore β-SN. As in Section 4, this is a new proof for a well-known result. The format of this section closely follows the format of Section 4.

5.1 The Intersection-Type Discipline.

The intersection-type discipline is defined as an extension of the simply-typed λ-calculus. Cardone and Coppo call the system presented here the system of *simple* intersection types, reserving the unqualified name for the system with the ω type constant that can be assigned to any λ-term [CC90]. Sometimes, a "$\leq$" rule is included with the system for type inclusion, but this is not necessary here since this rule does not change the set of typable λ-terms.

The set of intersection types $\mathbb{T}^{\wedge}$ is built from the countably infinite set of type variables $\mathbb{V}$ using the "$\rightarrow$" and "$\wedge$" type constructor as specified by the grammar $\mathbb{T}^{\wedge} ::= \mathbb{V} \mid (\mathbb{T}^{\wedge} \rightarrow \mathbb{T}^{\wedge}) \mid (\mathbb{T}^{\wedge} \wedge \mathbb{T}^{\wedge})$. A type is therefore either a *type variable,* a $\rightarrow$*-type,* or a $\wedge$*-type.* The same notation conventions are followed as with the simply-typed λ-calculus, except that "$\wedge$" is left-associative and has higher precedence than "$\rightarrow$" so that $\sigma \wedge \tau \rightarrow \rho = (\sigma \wedge \tau) \rightarrow \rho$.

The familiar Curry-style presentation of the intersection-type discipline uses the set of inference rules in Figure 1. In these inference rules, A stands for a set of types assigned to free variables and $A(x)$ is the particular type that A assigns to the free variable x. If there is a derivation in this system ending with the sequent $A \vdash M : \tau$ for some A and τ, then M is *typable* in the intersection-type discipline.

The Curry-style presentation of the intersection-type system makes it hard to define explicitly how β-reduction and γ-reduction work on typings. So we give an equivalent Church-style presentation of the intersection-type discipline. (This is done despite Barendregt's fairly accurate observation that "for ... the system of intersection types ... it is not clear how to define a Church version" [Bar92].)

As with the simply-typed λ-calculus, λ-term variables are pairs of untyped variables and types written as x^σ, y^τ, z^ρ, etc. The λ-terms will be required to satisfy the same property (†) as before, so that type assignments can be avoided. The set $\Lambda^{\wedge}$ of λ-terms in the intersection-type system and a type-erasing function $|\ \ |$ from $\Lambda^{\wedge}$ to Λ are defined inductively as follows:

VAR	$A \vdash x : \sigma$	$A(x) = \sigma$
APP	$\dfrac{A \vdash M : \sigma \rightarrow \tau, \qquad A \vdash N : \sigma}{A \vdash (M\ N) : \tau}$	
ABS	$\dfrac{A \cup \{x : \sigma\} \vdash M : \tau}{A \vdash (\lambda x.M) : \sigma \rightarrow \tau}$	
$\wedge{-}\mathcal{I}$	$\dfrac{A \vdash M : \sigma, \qquad A \vdash M : \tau}{A \vdash M : \sigma \wedge \tau}$	
$\wedge{-}\mathcal{E}$	$\dfrac{A \vdash M : \sigma \wedge \tau}{A \vdash M : \sigma, \qquad A \vdash M : \tau}$	

Figure 1: Inference Rules of Intersection-Type Discipline (Curry Style).

1. $x^\sigma \in \Lambda^\wedge$ and $|x^\sigma| = x$ if $x \in \mathcal{V}$ and $\sigma \in \mathbb{T}^\wedge$.
2. $(M^{\sigma\rightarrow\tau} N^\sigma)^\tau \in \Lambda^\wedge$ and $|(M^{\sigma\rightarrow\tau} N^\sigma)^\tau| = (|M^{\sigma\rightarrow\tau}||N^\sigma|)$ if $M^{\sigma\rightarrow\tau} \in \Lambda^\wedge$ and $N^\sigma \in \Lambda^\wedge$.
3. $(\lambda x^\sigma.M^\tau)^{\sigma\rightarrow\tau} \in \Lambda^\wedge$ and $|(\lambda x^\sigma.M^\tau)^{\sigma\rightarrow\tau}| = (\lambda x.|M^\tau|)$ if $M^\tau \in \Lambda^\wedge$ and $\sigma \in \mathbb{T}^\wedge$.
4. $(\wedge{-}\mathcal{I}\ M_1^{\sigma_1} \ldots M_n^{\sigma_n})^{\sigma_1\wedge\cdots\wedge\sigma_n} \in \Lambda^\wedge$ and $|(\wedge{-}\mathcal{I}\ M_1^{\sigma_1} \ldots M_n^{\sigma_n})^{\sigma_1\wedge\cdots\wedge\sigma_n}| = |M_1^{\sigma_1}|$ if $|M_1^{\sigma_1}| \equiv \cdots \equiv |M_n^{\sigma_n}|$ and for $1 \leq i \leq n$ it holds that σ_i is not a $\wedge$-type.
5. $(\wedge{-}\mathcal{E}\ M^{\sigma_1\wedge\cdots\wedge\sigma_n})^{\sigma_i} \in \Lambda^\wedge$ where $1 \leq i \leq n$ and $|(\wedge{-}\mathcal{E}\ M^{\sigma_1\wedge\cdots\wedge\sigma_n})^{\sigma_i}| = |M^{\sigma_1\wedge\cdots\wedge\sigma_n}|$ if $M^{\sigma_1\wedge\cdots\wedge\sigma_n} \in \Lambda^\wedge$.

This definition may seem quite unusual to some readers, but the clever reader will see that this presentation merely converts entire Curry-style derivations directly into terms. The only difference is that multiple uses of the $\wedge{-}\mathcal{I}$ and $\wedge{-}\mathcal{E}$ rules are compressed into one use and that no distinction is made between the type $\sigma \wedge (\tau \wedge \rho)$ and $(\sigma \wedge \tau) \wedge \rho$. Sometimes, we will be lazy and omit some of the type annotations from λ-terms in $\Lambda^\wedge$ when the types are not relevant.

Now it is necessary to give an interpretation for what it means to perform β-reduction and γ-reduction on members of $\Lambda^\wedge$. Before doing that, a way to eliminate redundant pairs of $\wedge{-}\mathcal{I}$ and $\wedge{-}\mathcal{E}$ must be provided. Although it can be assumed that the typing of a λ-term does not have redundant uses of $\wedge{-}\mathcal{I}$ and $\wedge{-}\mathcal{E}$, the natural way to interpret β-reduction and γ-reduction will introduce redundant pairs which must be eliminated. The auxiliary notion of reduction r is defined to handle this. Let r-reduction be the reduction relation such that:

$$(\wedge{-}\mathcal{E}\ (\wedge{-}\mathcal{I}\ M_1^{\sigma_1} \ldots M_n^{\sigma_n})^{\sigma_1\wedge\cdots\wedge\sigma_n})^{\sigma_i} \xrightarrow[r]{} M_i^{\sigma_i}$$

We assume at all times that typed λ-terms are in r-normal form, if necessary reducing them to r-normal form. It is obvious that r-reduction is strongly normalizing, since it is strictly size-reducing.

In providing an interpretation for β-reduction on members of $\Lambda^\wedge$, some basic requirements should be fulfilled. If $M^\sigma \xrightarrow[\beta]{} N^\tau$, then it should be the case that $\sigma = \tau$ and $|M^\sigma| \xrightarrow[\beta]{} |N^\tau|$. This still leaves room for many interpretations of β-reduction. For our purposes, the typing of N^σ should be essentially inherited from the typing for M^σ. To achieve this, we define the notion of the *parallel* occurrences of a subterm, define the notion of *singular* β-reduction, and then define *regular* β-reduction as the simultaneous singular reduction of all of the parallel occurrences of a β-redex.

The notion of the *parallel* occurrences of a subterm is defined as an equivalence relation on the subterm occurrences of each typed λ-term in $\Lambda^\wedge$. Let $\sim$ be the least equivalence relation such that:

1. $(\wedge{-}\mathcal{I}\, M_1^{\sigma_1} \ldots M_n^{\sigma_n})^{\sigma_1 \wedge \cdots \wedge \sigma_n} \sim M_i^{\sigma_i}$ for $1 \leq i \leq n$.
2. $(\wedge{-}\mathcal{E}\, M^{\sigma_1 \wedge \cdots \wedge \sigma_n})^{\sigma_i} \sim M^{\sigma_1 \wedge \cdots \wedge \sigma_n}$.
3. If $(M^{\sigma \to \tau} N^\sigma)^\tau \sim (M'^{\sigma' \to \tau'} N'^{\sigma'})^{\tau'}$ then $M^{\sigma \to \tau} \sim M'^{\sigma' \to \tau'}$ and $N^\sigma \sim N'^{\sigma'}$.
4. If $(\lambda x^\sigma . M^\tau)^{\sigma \to \tau} \sim (\lambda x'^{\sigma'} . M'^{\tau'})^{\sigma' \to \tau'}$ then $M^\tau \sim M'^{\tau'}$.

(Note that $\sim$ is defined on subterm *occurrences* with respect to a single, typed λ-term.) If, with respect to a typed λ-term M^τ, for the subterms $N_1^{\sigma_1}, N_2^{\sigma_2} \subset M^\tau$ it holds that $N_1^{\sigma_1} \sim N_2^{\sigma_2}$, then we say that $N_2^{\sigma_2}$ is a *parallel occurrence* of $N_1^{\sigma_1}$.

A *parallel set* X of subterms is all the members of an equivalence class under $\sim$ that are not $\wedge{-}\mathcal{I}$-terms or $\wedge{-}\mathcal{E}$-terms, for example all of the application subterms in a $\sim$-equivalence class. If $X = \{N_1^{\sigma_1}, \ldots, N_n^{\sigma_n}\}$ is a parallel set in $M^\tau \in \Lambda^\wedge$, then there is some corresponding untyped subterm $N \subseteq |M^\tau|$ such that $|N_i^{\sigma_i}| \equiv N$ for $1 \leq i \leq n$. Similarly, for every untyped subterm $N \subseteq |M^\tau|$, there is a corresponding parallel set X in M^τ. When the context is clear, N and X may be used interchangeably. The notion of the residuals of a parallel set under reduction is defined in terms of the residuals of the corresponding untyped subterm. If $C^\tau[\quad, \ldots, \quad]$ is a typed context such that all of the holes are parallel to each other, then $C^\tau[\quad, \ldots, \quad]$ is a *parallel context*.

The notion of *singular* β-reduction, denoted $\beta 1$-reduction, is exactly the ordinary notion of β-reduction from the simply-typed λ-calculus. It is the reduction relation such that:

$$((\lambda x^\sigma . M^\tau)^{\sigma \to \tau} N^\sigma)^\tau \xrightarrow[\beta 1]{} M^\tau[x^\sigma := N^\sigma]$$

Unfortunately, $\beta 1$-reduction does not stay within the set $\Lambda^\wedge$. If $M_1^{\sigma_1} \xrightarrow[\beta 1]{} N_1^{\sigma_1}$, then this implies that $(\wedge{-}\mathcal{I}\, M_1^{\sigma_1} M_2^{\sigma_2})^{\sigma_1 \wedge \sigma_2} \xrightarrow[\beta 1]{} (\wedge{-}\mathcal{I}\, N_1^{\sigma_1} M_2^{\sigma_2})^{\sigma_1 \wedge \sigma_2}$. However, since $|N_1^{\sigma_1}| \not\equiv |M_2^{\sigma_2}|$, the result of the $\beta 1$-reduction step is not even a valid member of $\Lambda^\wedge$. It is necessary to simultaneously $\beta 1$-reduce a parallel set of $\beta 1$-redexes.

The groundwork has now been laid for defining regular β-reduction. Observe that if

$$((\lambda x^\sigma . M^\tau)^{\sigma \to \tau} N^\sigma)^\tau \sim (P^{\sigma' \to \tau'} N'^{\sigma'})^{\tau'}$$

holds, then it must be the case that $P^{\sigma' \to \tau'}$ is a λ-abstraction and not a $\wedge{-}\mathcal{I}$-term or a $\wedge{-}\mathcal{E}$-term, due to our assumption that typed λ-terms are kept in r-normal form. This means that if one of the members of a parallel set is a $\beta 1$-redex, then all of them are $\beta 1$-redexes. Let $X = \{M_1^{\sigma_1}, \ldots, M_n^{\sigma_n}\}$ be a set of $\beta 1$-redexes such that $M_i^{\sigma_i} \xrightarrow[\beta 1]{} N_i^{\sigma_i}$ for $1 \leq i \leq n$. Let $M^\tau \equiv C^\tau[M_1^{\sigma_1}, \ldots, M_n^{\sigma_n}] \in \Lambda^\wedge$

where $C^\tau[\quad,\ldots,\quad]$ is a parallel context with n holes. In this case, the parallel set X in M^τ is said to be a *β-redex* and β-reduction is defined so that:

$$C^\tau[M_1^{\sigma_1},\ldots,M_n^{\sigma_n}] \xrightarrow[\beta]{X} C^\tau[N_1^{\sigma_1},\ldots,N_n^{\sigma_n}]$$

If $M^\tau \xrightarrow[\beta]{X} N^\tau$ and $|M^\tau| \xrightarrow[\beta]{\Delta} |N^\tau|$, then we may also say that M^τ reduces to N^τ by the redex Δ and we may refer to the typed β-redex X and the untyped β-redex Δ interchangeably.

It is also necessary to give a precise interpretation for γ-reduction on typed λ-terms in $\Lambda^\wedge$. It is a bit more complicated than the interpretation for β-reduction, because a use of the $\wedge{-}\mathcal{I}$ constructor can occur as part of the γ-redex. As with β-reduction, we define a notion of *singular* γ-reduction, denoted $\gamma 1$. There are two cases of $\gamma 1$-reduction, one simple and one complex. The first, simpler case is exactly as in the simply-typed λ-calculus:

$$((\lambda x^\sigma.(\lambda y^\rho.P^\tau)^{\rho\to\tau})^{\sigma\to\rho\to\tau}Q^\sigma)^{\rho\to\tau} \xrightarrow[\gamma 1]{} (\lambda y^\rho.((\lambda x^\sigma.P^\tau)^{\sigma\to\tau}Q^\sigma)^\tau)^{\rho\to\tau}$$

The more complicated case is like this, where π stands for the type $(\rho_1 \to \tau_1) \wedge \cdots \wedge (\rho_n \to \tau_n)$:

$$((\lambda x^\sigma.(\wedge{-}\mathcal{I}\,(\lambda y^{\rho_1}.P_1^{\tau_1})^{\rho_1\to\tau_1}\ldots(\lambda y^{\rho_n}.P_n^{\tau_n})^{\rho_n\to\tau_n})^\pi)^{\sigma\to\pi}Q^\sigma)^\pi$$

$$\Big\downarrow{\scriptstyle \gamma 1}$$

$$\begin{array}{l}(\wedge{-}\mathcal{I}\,(\lambda y^{\rho_1}.((\lambda x^\sigma.P_1^{\tau_1})^{\sigma\to\tau_1}Q^\sigma)^{\tau_1})^{\rho_1\to\tau_1}\\ \qquad\vdots\\ \qquad\quad(\lambda y^{\rho_n}.((\lambda x^\sigma.P_n^{\tau_n})^{\sigma\to\tau_n}Q^\sigma)^{\tau_n})^{\rho_n\to\tau_n})^\pi\end{array}$$

(Our conventions require immediate α-conversion on the result of this reduction to rename separate bound variables distinctly.) Given this definition of $\gamma 1$-reduction, regular γ-reduction is now defined in terms of $\gamma 1$-reduction in the exact same way that regular β-reduction is defined in terms of $\beta 1$-reduction.

5.2 Extending the Metric to Intersection Types.

The metric of Subsection 4.2 is now extended to handle intersection types. For intersection types, define the function *order* inductively as follows:

1. $order(\alpha) = 0$ where $\alpha \in \mathbb{V}$ is a type variable.
2. $order(\sigma \to \tau) = \max\{1 + order(\sigma), order(\tau)\}$.
3. $order(\sigma \wedge \tau) = \max\{order(\sigma), order(\tau)\}$.

For a singular $\beta 1$-redex $M^\tau \equiv ((\lambda x^\sigma.P^\tau)^{\sigma\to\tau}Q^\sigma)^\tau$, define $order(M^\tau) = order(\sigma)$, exactly as for the simply-typed λ-calculus. If $X = \{M_1^{\sigma_1},\ldots,M_n^{\sigma_n}\}$ is a set of $\beta 1$-redexes, then define $order(X) = \max\{order(M_1^{\sigma_1}),\ldots,order(M_n^{\sigma_n})\}$. For a typed λ-term M^τ, let $X_1, \ldots, X_n$ be all of I-redexes in M^τ. (Each X_i is a parallel set of $\beta 1$-redexes.) Then define $order^\bullet(M^\tau)$ to be the multiset $\{order(X_1),\ldots,order(X_n)\}$.

5.3 Normalizing $\star$-Reduction for Intersection Types.

We now prove that a particular $\star$-reduction strategy terminates for all λ-terms typable in the intersection-type discipline, implying all such terms are β-SN. The same reduction strategy is used that was presented in Subsection 4.4.

Denote the set of all typed λ-terms $M^\tau \in \Lambda^\wedge$ such that $|M^\tau| \in \Lambda^\gamma$ by $(\Lambda^\wedge)^\gamma$. Extend $\star$-reduction to terms in $\Lambda^\wedge$ in the obvious manner.

Lemma 5.1 *If $M, N \in (\Lambda^\wedge)^\gamma$ and $M \xrightarrow[\star]{\Delta} N$ where Δ is an innermost I-redex, then $order^\bullet(M) \succ order^\bullet(N)$.*

Proof: The structure of this proof is almost identical to the proof for Lemma 4.2. The main differences are the notation necessary to account for parallel sets and the handling of r-reduction in between β-reduction and γ-reduction steps.

Let $\Delta \equiv ((\lambda v.S)P)$ be the untyped, innermost I-redex and let $X = \{T_1, \ldots, T_n\}$ be the corresponding typed redex where for $1 \le i \le n$ it holds that T_i is of the form $((\lambda v_i^{\sigma_i}.S_i^{\tau_i})^{\sigma_i \to \tau_i} P_i^{\sigma_i})^{\tau_i}$. Let $M \equiv C[T_1, \ldots, T_n]$ for some parallel context $C[\ \ , \ldots, \ \ ]$. Let the set of all I-redex occurrences in M be $\{X_1, \ldots, X_q\}$ with X the same β-redex occurrence as X_1. Let $M' \equiv C[U_1, \ldots, U_n]$ where for $1 \le i \le n$ it holds that $U_i = S_i^{\tau_i}[v_i^{\sigma_i} := P_i^{\sigma_i}]$. ($r$-reduction is implicitly performed after the substitution.) It is clear that $M \xrightarrow[\beta]{\Delta} M'$ and that (by the given part of the claim) $M' \xrightarrow[\gamma\text{-}nf]{} N$. Note that M' is in $\Lambda^\wedge$ but may not be in $(\Lambda^\wedge)^\gamma$.

Consider the residuals in N of the I-redexes in M. It will be shown to be the case that there is exactly one residual for each such I-redex and that the value of *order* on the residual will be the same as for the original. At the untyped level, since Δ is an innermost I-redex, P contains no I-redexes that might be duplicated by the β-reduction step from M to M'. Thus, each I-redex except the one reduced has exactly one residual after the β-reduction step. By the definition of γ-reduction, although $\gamma 1$-reduction may duplicate $\beta 1$-redexes, it is clear that γ-reduction does not duplicate β-redexes (as parallel sets). The γ-reduction steps from M' to N do not duplicate or remove any β-redexes. After each β-reduction or γ-reduction step, r-reduction may need to be applied, but inspecting the definition of r-reduction reveals that it can not duplicate or remove β-redexes (as parallel sets). Thus, for $2 \le i \le q$, it is safe to define X_i' as the single residual of X_i in N. Now it remains to show that for $2 \le i \le q$ that $order(X_i) = order(X_i')$. The initial β-reduction step and any necessary subsequent r-reduction can not change the value of *order* on the residual of X_i, because each residual contains the same number of $\beta 1$-redexes each with the same type as before. The complicated case of $\gamma 1$-reduction can duplicate some number of $\beta 1$-redexes in the residual of X_i. However, in all copies the type of the bound variable is the same, so this does not change the value of *order*. r-reduction that occurs after a complicated $\gamma 1$-reduction may discard $\beta 1$-redexes in a residual of X_i. Since this can only happen immediately after a $\gamma 1$-reduction step which duplicated $\beta 1$-redexes, again, it can not change the value of *order*. The end result is that for $2 \le i \le q$, it must hold that $order(X_i) = order(X_i')$.

Given what we have shown so far, to prove that $order^\bullet(M) \succ order^\bullet(N)$ it is sufficient to show that for each new I-redex introduced by the β-reduction step or the subsequent γ-reduction steps, the value of *order* on this new I-redex is smaller than $order(X) = \max\{order(\sigma_1), \ldots, order(\sigma_n)\}$. In fact, we will show the stronger claim that this is the case for all new β-redexes, both I-redexes and K-redexes.

Consider separately the simple case where P is not a λ-abstraction or every occurrence of v in S is passive. In this case, it is easy to see that both:

1. M' is in γ-normal form and belongs to $(\Lambda^{\rightarrow})^{\gamma}$ and thus $M' \equiv \gamma\text{-}nf(M') \equiv N$.

2. No new I-redex occurrence is "created" by the β-reduction step, i.e. the I-redex occurrences in $M' \equiv N$ are exactly $\{X'_2, \ldots, X'_q\}$.

Keep in mind that, because M is in γ-normal form, S is not a λ-abstraction and, thus, S can not become the function of a β-redex. Thus, it is clear that $\mathit{order}^{\bullet}(N)$ is exactly $\mathit{order}^{\bullet}(M)$ with one occurrence of $\mathit{order}(X_1)$ removed, which implies that $\mathit{order}^{\bullet}(M) \succ \mathit{order}^{\bullet}(N)$.

Now consider the more complicated case where P is a λ-abstraction *and* there are active occurrences of v in S. Let P (untyped) be of the form $(\lambda x_1. \ldots .\lambda x_m.Q)$ where Q is not a λ-abstraction. For convenience, we name the various (untyped) subterms of P so that, for $1 \leq i \leq m$, $V_i \equiv (\lambda x_i.V_{i+1})$ and $P = V_1$ and $Q = V_{m+1}$. For $1 \leq i \leq n$, the (typed) subterm $P_i^{\sigma_i}$ can have a quite complicated substructure. The various subterms of P_i will be given names of the form $V_j^{\vec{k}}$ and $W_j^{\vec{k}}$ where $1 \leq j \leq m+1$ and $\vec{k}$ is a sequence of numbers beginning with i. For $1 \leq j \leq m+1$ and each $\vec{k}$, it will be the case that $|V_j^{\vec{k}}| \equiv V_j$ and also that $|W_j^{\vec{k}}| \equiv V_j$. Start by letting $P_i \equiv V_1^i$. For $1 \leq j \leq m$ and each $\vec{k}$, $V_j^{\vec{k}}$ is of one of these two forms:

$$(V_j^{\vec{k}})^{\psi_j^{\vec{k}}} \equiv \begin{cases} (W_j^{\vec{k},1})^{\theta_j^{\vec{k},1}} \\ (\wedge\text{-}\mathcal{I}\,(W_j^{\vec{k},1})^{\theta_j^{\vec{k},1}} \cdots (W_j^{\vec{k},p})^{\theta_j^{\vec{k},p}})^{\theta_j^{\vec{k},1} \wedge \cdots \wedge \theta_j^{\vec{k},p}} & \text{for some } p > 1 \end{cases}$$

and $W_j^{\vec{k}}$ is of this form:

$$(W_j^{\vec{k}})^{\theta_j^{\vec{k}}} \equiv (\lambda (x_j^{\vec{k}})^{\varphi_j^{\vec{k}}}.(V_{j+1}^{\vec{k}})^{\psi_{j+1}^{\vec{k}}})^{\varphi_j^{\vec{k}} \rightarrow \psi_{j+1}^{\vec{k}}}$$

The convention has been adopted here that the type of $V_j^{\vec{k}}$ is named $\psi_j^{\vec{k}}$, the type of $W_j^{\vec{k}}$ is named $\theta_j^{\vec{k}}$, and the type given to the bound variable $x_j^{\vec{k}}$ is named $\varphi_j^{\vec{k}}$.

It will be important later to know the maximum value of *order* for the type associated with each bound variable $x_j^{\vec{k}}$ which corresponds to the untyped variable x_j. Since either

$$\psi_j^{\vec{k}} = \theta_j^{\vec{k},1} \quad \text{or} \quad \psi_j^{\vec{k}} = (\theta_j^{\vec{k},1} \wedge \cdots \wedge \theta_j^{\vec{k},p}) \text{ for some } p > 1$$

it must be the case that either

$$\mathit{order}(\psi_j^{\vec{k}}) = \mathit{order}(\theta_j^{\vec{k},1}) \quad \text{or} \quad \mathit{order}(\psi_j^{\vec{k}}) = \max\{\mathit{order}(\theta_j^{\vec{k},1}), \ldots, \mathit{order}(\theta_j^{\vec{k},p})\} \text{ for some } p > 1$$

Thus, in either case, for any $p \geq 1$, $\mathit{order}(\psi_j^{\vec{k}}) \geq \mathit{order}(\theta_j^{\vec{k},p})$. Since $\theta_j^{\vec{k}} = (\varphi_j^{\vec{k}} \rightarrow \psi_{j+1}^{\vec{k}})$, by definition it is the case that $\mathit{order}(\theta_j^{\vec{k}}) = \max\{1 + \mathit{order}(\varphi_j^{\vec{k}}), \mathit{order}(\psi_{j+1}^{\vec{k}})\}$. Thus, both $\mathit{order}(\theta_j^{\vec{k}}) \geq \mathit{order}(\psi_{j+1}^{\vec{k}})$ and $\mathit{order}(\theta_j^{\vec{k}}) > \mathit{order}(\varphi_j^{\vec{k}})$. At this point, a simple induction establishes that for $1 \leq j \leq m$, for $1 \leq i \leq n$, and for any $\vec{k}$ that begins with i that $\mathit{order}(X) \geq \mathit{order}(\sigma_i) \geq \mathit{order}(\theta_j^{\vec{k}}) > \mathit{order}(\varphi_j^{\vec{k}})$.

For each active occurrence of v in S, one or more β-redexes will be formed by the $\star$-reduction step. Usually, some of these β-redexes will be formed by the β-reduction step from M to M' and some will be formed by the γ-reduction steps from M' to N.

First, we show that for each β-redex Y formed by the β-reduction step that $\mathit{order}(Y) < \mathit{order}(X)$. Examine the untyped version of the β-reduction step. Wherever v occurs as (vR) for some subterm R, the β-reduction step will form the β-redex $((\lambda x_1.\ldots.\lambda x_m.Q)R)$. No other kind of β-redex can be formed by the β-reduction step. Examine now the typed version of the β-reduction step. Within each S_i, occurrences of $v_i^{\sigma_i}$ are replaced by $P_i^{\sigma_i}$. Suppose the type σ_i is a $\wedge$-type. Since the rule for forming applications requires a $\rightarrow$-type for the function, there must be an occurrence of the $\wedge\!-\!\mathcal{E}$ constructor. So an instance of the substitution will look like this:

$$((\wedge\!-\!\mathcal{E}\,(\wedge\!-\!\mathcal{I}\,(W_1^{i,1})^{\theta_1^{i,1}}\cdots(W_1^{i,p})^{\theta_1^{i,p}}))R)$$

for some $p > 1$. One r-reduction step will produce this:

$$((W_1^{i,j})^{\theta_1^{i,j}}R)$$

where $1 \leq j \leq p$. Suppose instead the type σ_i is a $\rightarrow$-type. Then an instance of the substitution will look like this:

$$((W_1^{i,1})^{\theta_1^{i,1}}R)$$

In either case, the type of the function of the new $\beta 1$-redex is $\theta_1^{i,1} = (\varphi_1^{i,1} \rightarrow \psi_2^{i,1})$. Thus, the value of order on the new $\beta 1$-redex is $\mathit{order}(\varphi_1^{i,1})$ which has already been shown to be smaller than $\mathit{order}(\sigma_i)$ which is less than or equal to $\mathit{order}(X)$. The new β-redex Y is a parallel set of such new $\beta 1$-redexes. The value of $\mathit{order}(Y)$ is the maximum of the value on each of its constituent $\beta 1$-redexes. Thus, the desired result is shown.

Now, we show that for each β-redex Y formed by the γ-reduction steps from M' to N that $\mathit{order}(Y) < \mathit{order}(X)$. To do this, first we analyze the λ-abstractions which may become involved in fresh β-redexes. By Lemma 3.4, we know that (at the untyped level) if a β-redex is formed by one of the γ-reduction steps from M' to N, the function of the new β-redex must be a λ-abstraction that existed in M', which was not already part of a β-redex, and which was active in M'. Since M was in γ-normal form, all active λ-abstractions in M were already the functions of β-redexes. The only active abstractions in M' which are not functions of β-redexes are the outermost abstractions of the copies of P wherever an active occurrence of v was replaced. Thus, for each fresh (untyped) β-redex Γ, the function of Γ must be a copy of one of the outermost abstractions of P. Consider a fresh (typed) β-redex $Y = \{Z_1^{\rho_1}, \ldots, Z_r^{\rho_r}\}$. From the preceding statements we may conclude that for each $\beta 1$-redex Z_i that the bound variable of the function of Z_i is a copy of $x_j^{\vec{k}}$ for some j and $\vec{k}$. Recall the definition of how γ-reduction works on λ-terms in $\Lambda^\wedge$. When a γ-reduction step rearranges λ-abstractions, the types assigned to the bound variables move with them. When a r-reduction step rearranges subterms, any types assigned to bound variables after the step were assigned to those bound variables before the step. Thus, $\mathit{order}(Z_i^{\rho_i}) = \mathit{order}(\varphi_j^{\vec{k}}) < \mathit{order}(X)$. Since $\mathit{order}(Y)$ is the maximum of these values, it is clear that $\mathit{order}(Y) < \mathit{order}(X)$. ■

Theorem 5.2 *If $M^\tau \in \Lambda^\wedge$, then M^τ is β-SN.*

Proof: The proof is a combination of the proofs for Lemma 4.3 and Theorem 4.4 except that it depends on Lemma 5.1 instead of Lemma 4.2. ■

References

[Bar84] H. P. Barendregt. *The Lambda Calculus: Its Syntax and Semantics*. North-Holland, revised edition, 1984.

[Bar92] H. P. Barendregt. Lambda calculi with types. In S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum, eds., *Handbook of Logic in Computer Science*, vol. 2, chapter 2, pp. 117–309. Oxford University Press, 1992.

[CC90] F. Cardone and M. Coppo. Two extensions of Curry's type inference system. In Odifreddi [Odi90], chapter 1, pp. 19–75.

[CDC80] M. Coppo and M. Dezani-Ciancaglini. An extension of basic functionality theory for lambda-calculus. *Notre Dame J. Formal Log.*, 21:685–693, 1980.

[CDCV81] M. Coppo, M. Dezani-Ciancaglini, and B. Venneri. Functional characters of solvable terms. *Z. Math. Log. Grund. Math.*, 27:45–58, 1981.

[DM79] N. Dershowitz and Z. Manna. Proving termination with multiset orderings. *J. ACM*, 22:465–476, 1979.

[Gal90] J. H. Gallier. On Girard's "candidats de reductibilité". In Odifreddi [Odi90], pp. 123–203.

[Gir71] J.-Y. Girard. Une extension de l'interprétation de Gödel à l'analyse, et son application à l'elimination des coupures dans l'analyse et la théorie des types. In J. E. Fenstad, ed., *Proceedings of 2nd Scandinavian Logic Symposium*, pp. 63–92, Amsterdam, 1971. North Holland.

[Gir72] J.-Y. Girard. *Interprétation Fonctionnelle et Elimination des Coupures de l'Arithmétique d'Ordre Supérieur*. Thèse d'Etat, Université Paris VII, 1972.

[Lei86] D. Leivant. Typing and computational properties of lambda expressions. *Theoretical Comput. Sci.*, 44:51–68, 1986.

[MKO94] D. A. McAllester, J. Kucan, and D. Otth. A proof of strong normalization for $\mathbf{F_2}$, $\mathbf{F_\omega}$, and beyond. Technical report, Massachusetts Institute of Technology Laboratory for Computer Science, 1994.

[Odi90] P. Odifreddi, ed. *Logic and Computer Science*. Number 31 in the APIC Series. Academic Press, 1990.

[Pot80] G. Pottinger. A type assignment for the strongly normalizable λ-terms. In J. P. Seldin and J. R. Hindley, eds., *To H. B. Curry: Essays on Combinatory Logic, Lambda Calculus, and Formalism*, pp. 561–577. Academic Press, 1980.

[Tai67] W. W. Tait. Intensional interpretation of functionals of finite type I. *J. Symbolic Logic*, 32:198–212, 1967.