

An optimised term graph rewriting engine for \mathcal{X}

Measuring the cost of α -conversion

Steffen van Bakel and Jayshan Raghunandan

Department of Computing, Imperial College London,
180 Queen's Gate, London SW7 2BZ, UK
{svb, jr200}@doc.ic.ac.uk

Abstract

This paper studies the calculus \mathcal{X} , that has its foundation in Classical Logic; we present an implementation for \mathcal{X} using term graph rewriting techniques, and discuss improvements thereof which result in an increasingly more efficient running of the reduction engine. We show that name capture can be dealt with ‘on the fly’, by realising the avoidance of capture through adding or modifying the rewrite rules. We study two different approaches to garbage collection, and compare the various implementations by presenting benchmarks.

1 Introduction

This paper presents a term graph rewriting implementation for the (untyped) calculus \mathcal{X} as first presented in [4], and studies the efficiency of a number of approaches towards implementation issues, in particular that of α -conversion. \mathcal{X} is a sequent calculus which embodies both substitution and context call; it was first been defined in [24, 25, 18] and was later extensively studied in [2, 3].

The origin of \mathcal{X} lies within the quest for a Curry-Howard correspondence to Gentzen’s sequent calculus LK for Classical Logic, introduced in [14]. In particular, \mathcal{X} corresponds to the implicative fragment of Kleene’s G_3 [17], with implicit weakening and contraction. LK is a logical system in which the rules only introduce connectives (but on either side of a sequent), in contrast to *natural deduction* (also introduced in [14]) which uses rules that introduce or eliminate connectives in the logical formulae. Natural deduction normally derives statements with a single conclusion, whereas LK allows for multiple conclusions, deriving sequents of the form $A_1, \dots, A_n \vdash B_1, \dots, B_m$, where A_1, \dots, A_n is to be understood as $A_1 \wedge \dots \wedge A_n$ and B_1, \dots, B_m is to be understood as $B_1 \vee \dots \vee B_m$. G_3 has four rules: *axiom*, *left introduction* of the arrow, *right introduction*, and *cut*.

$$\begin{array}{ll} (Ax) : \frac{}{\Gamma, A \vdash A, \Delta} & (cut) : \frac{\Gamma \vdash A, \Delta \quad \Gamma, A \vdash \Delta}{\Gamma \vdash \Delta} \\ (\Rightarrow R) : \frac{\Gamma, A \vdash B, \Delta}{\Gamma \vdash A \Rightarrow B, \Delta} & (\Rightarrow L) : \frac{\Gamma \vdash A, \Delta \quad \Gamma, B \vdash \Delta}{\Gamma, A \Rightarrow B \vdash \Delta} \end{array}$$

Since LK has only *introduction* rules (or better, no *elimination* rules), the only way to eliminate a connective is to eliminate the whole formula in which it appears via an application of the (*cut*)-rule. Gentzen defined an informal procedure that eliminates all applications of the (*cut*)-rule from a proof of a sequent, generating a proof in *normal form* of the same sequent, that is, without a cut; this procedure is defined via local reductions of the proof tree which essentially correspond to innermost reduction.

Starting from different approaches in that area, in [18] Lengrand introduced \mathcal{X} as a calculus that enjoys the Curry-Howard isomorphism for G_3 ; similar work was presented by Urban [24, 25], who defined an expressive restriction of cut-elimination that is strongly normalisable.

\mathcal{X} can be seen as an extension of Urban’s calculus, and was studied in [2, 3] in the context of a number of calculi. A notable difference with Urban’s work is that he set out to study the structure of proofs, so the terms of his calculus carry types and correspond only to proofs. In the approach to \mathcal{X} of [2, 3] which we take also here, we study terms *without* types, and drop the condition that terms should represent proofs of the logic altogether: we study a pure calculus. This opens the research to notions as recursion, normalising strategies, confluence, etc.

The calculus \mathcal{X} achieves a Curry-Howard isomorphism for the proofs in LK by constructing *witnesses* for derivable sequents. In establishing the isomorphism, similar to calculi like Parigot’s $\lambda\mu$ [20] and $\bar{\lambda}\mu\tilde{\mu}$ [13], Roman names are attached to formulae in the left context, and Greek names for those on the right, and syntactic structure is associated to the rules. Names on the left can be seen as inputs to the term, and names to the right as outputs; since multiple formulae can appear on both sides, this implies that a term can not only have more than one input, but also more than one output. There are two kinds of names (connectors) in \mathcal{X} : *sockets* (inputs, with Roman names, that are reminiscent of values) and *plugs* (outputs, with Greek names, that are reminiscent of continuations), that correspond to *variables* and *co-variables*, respectively, in [28], or, alternatively, to Parigot’s λ and μ -variables [20].

Gentzen’s proof reductions by cut-elimination become the fundamental principle of computation in \mathcal{X} . Cuts in proofs are witnessed by $P\hat{\alpha} \dagger \hat{x}Q$ (called the *cut* of P and Q via α and x), and the reduction rules specify how to remove them: a term is in normal form if and only if it has no sub-term of this shape. The generalisation with respect to Gentzen’s proof reduction is that now all cuts can be eliminated, not just those that are innermost.

\mathcal{X} as an untyped language provides an excellent general purpose machine, very well suited to encode various calculi; [3] illustrates the expressive power of \mathcal{X} by giving consistent interpretations of the Lambda Calculus [12, 5], Bloo and Rose’s calculus of explicit substitution λx [10], $\lambda\mu$, and $\bar{\lambda}\mu\tilde{\mu}$.

Perhaps the main feature of \mathcal{X} is that it constitutes a *variable, application, and substitution-free* method of computation. Rather than having variables like x representing places where terms can be inserted, in \mathcal{X} the symbol x represents a *socket*, to which a term can be *attached* via a *plug* α ; both plugs and sockets carry *names*, and the only substitution-like operation is that of *renaming*. The definition of reduction on \mathcal{X} constitutes of rules that describe the reorganisation of a term in small steps, and shows the subtle interaction between plugs and sockets.

Reduction in \mathcal{X} can be specified as a *conditional term rewriting system*; the only non-standard aspects is the presence of binding structures. It is this observation that led us to the research we report on in this paper; i.e. the building of an interpreter for \mathcal{X} using the term graph rewriting technology. A tool was developed that allows users to input terms from \mathcal{X} , reducing cuts at will in \mathcal{X} ’s full reduction system, or to limit reduction to either call-by-name or call-by-value.

We not only set out to study \mathcal{X} and its properties, but also focus on trying to extend \mathcal{X} into a true programming language. With that in mind, we have concentrated on building an efficient interpreter, and sought different, increasingly better solutions to garbage collection, capture avoidance and name capture. In [4] we reported on the first results of our implementation efforts. In particular, to avoid problems caused by nested binding of connectors, a lazy copy mechanism was introduced, and almost all rewrite rules were defined using this.

In this paper we will improve on this result, and show that the special character of \mathcal{X} , being a conditional term rewriting system based on renaming, makes it possible to study α -conversion *on the level of the language itself*. So in \mathcal{X} , α -conversion is not just a hidden implementation issue, but a first-class citizen. This allows us to directly measure the cost of α -conversion using the same currency as for the cost of cut-elimination: with minor changes

to the rules, we can test our different solutions on the level of the language itself.

This paper does not claim to find innovative solutions to the problem of capture itself; many solutions to this problem exist, and we have, so far, chosen standard approaches. What we do achieve is the creation of a platform on which it is easy *to compare* these various solutions in terms of execution cost, thus enabling justification for the choice of a certain technology over another.

We will present improvements on the reduction engine in the tool as first presented in [4], and show that these are expressible directly in \mathcal{X} itself. We will present three solutions for the problem of capture that are both expressed as changes to the rules: the first will use the rebinding of bound connectors, the second will preserve Barendregt’s convention, whereas the third avoids the capture of free connectors by binders. Much to our surprise, the resulting rewrite rules of the latter two solutions are very similar, but for the fact that in the last freeness is an issue, rather than being bound as in the second. The reduction engine thus obtained proved to be impressively much more efficient though, especially after the addition of two different notions of garbage collection.

Outline of this paper. This paper is organised as follows. In Section 2 we will repeat the \mathcal{X} -calculus, and some of its main properties. In Section 3, we will show that more expressive rules regarding garbage collection and deactivation can be safely added to the calculus. Section 4 presents the term graph rewriting engine that is used for our implementation, and in Section 5 we discuss three ways to deal with α -conversion. In Section 6 we formally define our reduction strategy combinators, that help us implement our three solutions in comparable ways, and give our measurement results.

2 The \mathcal{X} -calculus

In this section we will give the definition of the \mathcal{X} -calculus [18] that was proven to be a fine-grained implementation model for various well-known calculi in [2]. \mathcal{X} features two separate categories of ‘connectors’, *plugs* and *sockets*, that act as input and output channels.

As mentioned in the introduction, \mathcal{X} enjoys the Curry-Howard relation with Kleene’s variant G_3 of Gentzen’s LK. This is achieved by labelling the formulae with term information, and building witnesses for proofs by associating derivation rules to syntactical constructs. In the construction of the witness of a logical statement, when in applying a rule a premise or conclusion disappears from the sequent, the corresponding name gets bound in the term that is constructed, and when a premise or conclusion gets created, a different, free (often new) name is associated to it. For example, in the creation of the term for right introduction of the arrow

$$\frac{P \cdot \Gamma, x:A \vdash_{\mathcal{X}} \alpha:B, \Delta}{\widehat{x}P\widehat{\alpha}\cdot\beta \cdot \Gamma \vdash_{\mathcal{X}} \beta:A \rightarrow B, \Delta}$$

the input x and the output α are bound, and β is free. This case is interesting in that it highlights a special feature of \mathcal{X} , not found in other calculi. In (applicative) calculi related to natural deduction, like the λ -calculus, only inputs are named, and the linking to a term that will be inserted is done via λ -abstraction and application. The output (i.e. result) on the other hand is anonymous; where a term ‘moves to’ carries a name via a *variable* that acts as a pointer to the positions where the term is to be inserted, but where it comes from is not mentioned, since it is implicit. Since in \mathcal{X} a term P can have many inputs and outputs, it is unsound to consider P a function; however, fixing *one* input x and *one* output α , we can see P as a function ‘from x to α ’. We make this limited view of P available via the output β , thereby *exporting* ‘ P is a function from x to α ’. The types given to the connectors confirm this view, which is also

supported by the logic.

Definition 2.1 (SYNTAX) The terms of the \mathcal{X} -calculus are defined by the following grammar, where x, y range over the infinite set of *sockets*, and α, β over *plugs*.

$$P, Q ::= \langle x \cdot \alpha \rangle \quad | \quad \hat{y}P\hat{\beta} \cdot \alpha \quad | \quad P\hat{\beta}[y]\hat{x}Q \quad | \quad P\hat{\alpha} \dagger \hat{x}Q$$

capsule export import cut.

Notice that \mathcal{X} has no notion of abstraction; in particular, $P\hat{\beta}$ and $\hat{x}Q$ are not terms.

The $\hat{\cdot}$ symbolises that the socket (plug) underneath is bound in the term directly to the right (left). The unconnected inputs and outputs in a term make up the collection of *free connectors*, that are *inactive* during a computational step.

Definition 2.2 (FREE AND BOUND) The sets of *free* and *bound sockets* and *free* and *bound plugs* in a term are defined by:

$$\begin{array}{ll} fs(\langle x \cdot \alpha \rangle) & = \{x\} & bs(\langle x \cdot \alpha \rangle) & = \emptyset \\ fs(\hat{y}P\hat{\beta} \cdot \alpha) & = fs(P) \setminus \{y\} & bs(\hat{y}P\hat{\beta} \cdot \alpha) & = bs(P) \cup \{y\} \\ fs(P\hat{\beta}[y]\hat{x}Q) & = fs(P) \cup (fs(Q) \setminus \{x\}) & bs(P\hat{\beta}[y]\hat{x}Q) & = bs(P) \cup bs(Q) \cup \{x\} \\ fs(P\hat{\alpha} \dagger \hat{x}Q) & = fs(P) \cup (fs(Q) \setminus \{x\}) & bs(P\hat{\alpha} \dagger \hat{x}Q) & = bs(P) \cup bs(Q) \cup \{x\} \\ \\ fp(\langle x \cdot \alpha \rangle) & = \{\alpha\} & bp(\langle x \cdot \alpha \rangle) & = \emptyset \\ fp(\hat{x}P\hat{\alpha} \cdot \beta) & = fp(P) \setminus \{\alpha\} & bp(\hat{x}P\hat{\alpha} \cdot \beta) & = bp(P) \cup \{\alpha\} \\ fp(P\hat{\alpha}[y]\hat{x}Q) & = (fp(P) \setminus \{\alpha\}) \cup fp(Q) & bp(P\hat{\alpha}[y]\hat{x}Q) & = bp(P) \cup \{\alpha\} \cup bp(Q) \\ fp(P\hat{\alpha} \dagger \hat{x}Q) & = (fp(P) \setminus \{\alpha\}) \cup fp(Q) & bp(P\hat{\alpha} \dagger \hat{x}Q) & = bp(P) \cup \{\alpha\} \cup bp(Q) \end{array}$$

The set of *bound connectors* is defined as $bc(P) = bs(P) \cup bp(P)$, and we sometimes write, for example, $bs(P, Q)$ as shorthand for $bs(P) \cup bs(Q)$; we use the notation $fc(P) (= fs(P) \cup fp(P))$ for the free connectors.

Note that, in $\hat{x}\langle x \cdot \alpha \rangle\hat{\alpha} \cdot \alpha$, α is bound *and* free. The standard way to avoid this anomaly is to assume Barendregt's convention, which states that *free* and *bound* names should be distinct. As illustrated in Section 5, this then becomes an implementation issue, which we will address here for \mathcal{X} . As usual, we will identify terms that only differ in the names of bound connectors (modulo α -conversion, as usual). To maintain this convention during reduction, implicit α -conversion needs to take place, which is the main issue we tackle in this paper (see Section 5).

2.1 Reduction on \mathcal{X}

The calculus, defined by the reduction rules below, explains in detail how cuts are propagated through terms to be eventually evaluated at the level of capsules.

The intuition behind reduction is: the cut $P\hat{\alpha} \dagger \hat{x}Q$ expresses the intention to connect all α s in P and x s in Q , and reduction will realise this by either connecting all α s to all x s (if x does not exist in Q , P will disappear), or all x s to all α s (if α does not exist in P , Q will disappear).

Reduction is defined by specifying both the interaction between well-connected syntactic structures, and how to deal with propagating active nodes to points in the term where they can interact. This strongly depends on the following notion.

Definition 2.3 (INTRODUCTION OF CONNECTORS) P introduces x : Either $P = Q\hat{\beta}[x]\hat{y}R$ and $x \notin fs(Q, R)$, or $P = \langle x \cdot \alpha \rangle$.

P introduces α : Either $P = \hat{x}Q\hat{\beta} \cdot \alpha$ and $\alpha \notin fp(Q)$, or $P = \langle x \cdot \alpha \rangle$.

We consider P and Q in $P\hat{\alpha} \dagger \hat{x}Q$ well connected when both α and x are introduced.

The main reduction rules are:

Definition 2.4 (LOGICAL RULES) The logical rules are presented by (where both α and x are introduced in the cuts):

$$\begin{aligned}
(\text{cap}) : & \quad \langle y \cdot \alpha \rangle \hat{\alpha} \dagger \hat{x} \langle x \cdot \beta \rangle \rightarrow \langle y \cdot \beta \rangle \\
(\text{exp}) : & \quad (\hat{y}P\hat{\beta} \cdot \alpha) \hat{\alpha} \dagger \hat{x} \langle x \cdot \gamma \rangle \rightarrow \hat{y}P\hat{\beta} \cdot \gamma \\
(\text{imp}) : & \quad \langle y \cdot \alpha \rangle \hat{\alpha} \dagger \hat{x} (P\hat{\beta}[x]\hat{z}Q) \rightarrow P\hat{\beta}[y]\hat{z}Q \\
(\text{exp-imp}) : & \quad (\hat{y}P\hat{\beta} \cdot \alpha) \hat{\alpha} \dagger \hat{x} (Q\hat{\gamma}[x]\hat{z}R) \rightarrow \begin{cases} Q\hat{\gamma} \dagger \hat{y}(P\hat{\beta} \dagger \hat{z}R) \\ (Q\hat{\gamma} \dagger \hat{y}P)\hat{\beta} \dagger \hat{z}R \end{cases}
\end{aligned}$$

The first three logical rules above specify a renaming (reconnecting) procedure, whereas the last rule specifies the basic computational step: it links the export of a function, available on the plug α , to an adjacent import via the socket x .

In case the connectors in a cut are *not* introduced, the logical rules cannot be directly applied, and the propagation rules specified below come into play. Applying these will either try to connect the socket to the plug, or the plug to the socket. This choice is expressed via the *activation* of a cut, represented by the tilting of a dagger.

Definition 2.5 (ACTIVE CUTS) The syntax is extended with two *flagged* or *active* cuts:

$$P, Q ::= \dots \mid P\hat{\alpha} \not\wedge \hat{x}Q \mid P\hat{\alpha} \backslash \hat{x}Q$$

Terms constructed without these flagged cuts are called *pure*.

A right-activated cut $P\hat{\alpha} \backslash \hat{x}Q$ corresponds to the substitution of P for x via α , and will attempt to connect all α s to all x s, whereas the left-activated cut $P\hat{\alpha} \not\wedge \hat{x}Q$ is its dual, the substitution of Q for α via x , and will attempt to connect all x s to all α s.

Definition 2.6 (ACTIVATING) We define two cut-activation rules.

$$\begin{aligned}
(a \not\wedge) : & \quad P\hat{\alpha} \dagger \hat{x}Q \rightarrow P\hat{\alpha} \not\wedge \hat{x}Q \quad \text{if } P \text{ does not introduce } \alpha \\
(\backslash a) : & \quad P\hat{\alpha} \dagger \hat{x}Q \rightarrow P\hat{\alpha} \backslash \hat{x}Q \quad \text{if } Q \text{ does not introduce } x
\end{aligned}$$

An activated cut is processed by ‘pushing’ it systematically through the syntactic structure of the term in the direction indicated by the tilting of the dagger, as described by the propagation rules below. The pushing of the active cut continues until the level of capsules is reached, where it is either deactivated or destroyed.

Definition 2.7 (PROPAGATION RULES) **Left propagation:**

$$\begin{aligned}
(d \not\wedge) : & \quad \langle y \cdot \alpha \rangle \hat{\alpha} \not\wedge \hat{x}P \rightarrow \langle y \cdot \alpha \rangle \hat{\alpha} \dagger \hat{x}P \\
(\text{cap} \not\wedge) : & \quad \langle y \cdot \beta \rangle \hat{\alpha} \not\wedge \hat{x}P \rightarrow \langle y \cdot \beta \rangle, \quad (\beta \neq \alpha) \\
(\text{exp-outs} \not\wedge) : & \quad (\hat{y}Q\hat{\beta} \cdot \alpha) \hat{\alpha} \not\wedge \hat{x}P \rightarrow (\hat{y}(Q\hat{\alpha} \not\wedge \hat{x}P)\hat{\beta} \cdot \gamma) \hat{\gamma} \dagger \hat{x}P, \quad (\gamma \text{ fresh}) \\
(\text{exp-ins} \not\wedge) : & \quad (\hat{y}Q\hat{\beta} \cdot \gamma) \hat{\alpha} \not\wedge \hat{x}P \rightarrow \hat{y}(Q\hat{\alpha} \not\wedge \hat{x}P)\hat{\beta} \cdot \gamma, \quad (\gamma \neq \alpha) \\
(\text{imp} \not\wedge) : & \quad (Q\hat{\beta}[z]\hat{y}R) \hat{\alpha} \not\wedge \hat{x}P \rightarrow (Q\hat{\alpha} \not\wedge \hat{x}P)\hat{\beta}[z]\hat{y}(R\hat{\alpha} \not\wedge \hat{x}P) \\
(\text{cut} \not\wedge) : & \quad (Q\hat{\beta} \dagger \hat{y}R) \hat{\alpha} \not\wedge \hat{x}P \rightarrow (Q\hat{\alpha} \not\wedge \hat{x}P)\hat{\beta} \dagger \hat{y}(R\hat{\alpha} \not\wedge \hat{x}P)
\end{aligned}$$

ii) Likewise, we define *call by name* reduction $P \rightarrow_N Q$ as the reduction system obtained by replacing rules $(a \not\lambda)$ and $(exp-imp)$ by:

$$(a \not\lambda_N): \quad P\hat{\alpha} \dagger \hat{x}Q \rightarrow P\hat{\alpha} \not\lambda \hat{x}Q, \text{ (if } Q \text{ introduces } x \text{ and } P \text{ does not introduce } \alpha)$$

$$(exp-imp_N): (\hat{y}P\hat{\beta} \cdot \alpha)\hat{\alpha} \dagger \hat{x}(Q\hat{\gamma}[x]\hat{z}R) \rightarrow (Q\hat{\gamma} \dagger \hat{y}P)\hat{\beta} \dagger \hat{z}R$$

If a cut can be activated in two ways, the CBV strategy only allows to activate it via $(a \not\lambda)$; the CBN strategy can only activate such a cut via $(\not\lambda a)$.

This way, we obtain two notions of reduction that are clearly confluent: all rules are left-linear and non-overlapping.

2.2 Typing for \mathcal{X}

The notion of type assignment on \mathcal{X} corresponds to the basic implicative system G_3 for Classical Logic.

Definition 2.11 (TYPES AND CONTEXTS [2]) *i)* The set of types², ranged over by A, B , is defined over a set of *type-variables* ranged over by φ , by the grammar:

$$A, B ::= \varphi \mid A \rightarrow B$$

ii) A *context of sockets* Γ is a mapping from sockets to types, denoted as a finite set of *statements* $x:A$, such that the *subjects* of the statements (the sockets) are distinct. We write $\Gamma, x:A$ for the context $\Gamma \cup \{x:A\}$.

Contexts of plugs Δ and the notation $\alpha:A, \Delta$ are defined similarly.

Definition 2.12 (TYPING FOR \mathcal{X} [2]) *i)* *Type judgements* are expressed via the ternary relation $P : \cdot \Gamma \vdash \Delta$, where Γ is a context of sockets and Δ is a context of plugs, and P is a term. We say that P is the *witness* of this judgement.

ii) *Type assignment for \mathcal{X}* is defined by the following sequent rules:

$$(ax): \frac{}{\langle x \cdot \alpha \rangle : \cdot \Gamma, x:A \vdash \alpha:A, \Delta} \quad (cut): \frac{P : \cdot \Gamma \vdash \alpha:A, \Delta \quad Q : \cdot \Gamma, x:A \vdash \Delta}{P\hat{\alpha} \dagger \hat{x}Q : \cdot \Gamma \vdash \Delta}$$

$$(\Rightarrow R): \frac{P : \cdot \Gamma, x:A \vdash \alpha:B, \Delta}{\hat{x}P\hat{\alpha} \cdot \beta : \cdot \Gamma \vdash \beta:A \rightarrow B, \Delta} \quad (\Rightarrow L): \frac{P : \cdot \Gamma \vdash \alpha:A, \Delta \quad Q : \cdot \Gamma, x:B \vdash \Delta}{P\hat{\alpha}[y]\hat{x}Q : \cdot \Gamma, y:A \rightarrow B \vdash \Delta}$$

We write $P : \cdot \Gamma \vdash \Delta$ if there exists a derivation built out of these rules that has this judgement in the bottom line.

Notice that, by the special meaning we associate to the comma in contexts, contraction is implicit. Also, since Γ and Δ are non-specific, so is weakening.

In $P : \cdot \Gamma \vdash \Delta$, the term P acts as a *witness* of the judgement; Γ and Δ carry the types of the free connectors in P , as unordered sets.

The following result was shown in [2]:

Theorem 2.13 (WITNESS REDUCTION [2]) *If $P : \cdot \Gamma \vdash \Delta$, and $P \rightarrow Q$, then $Q : \cdot \Gamma \vdash \Delta$.*

²The types considered in this paper are normally known as *Curry* types.

2.3 \mathcal{X} as a reduction machine for the λ -calculus

In [2] the relation between \mathcal{X} and many other calculi is studied; as an illustration, in this section, we will briefly highlight the relation between the $\lambda\mathbf{x}$ and \mathcal{X} . We assume the reader to be familiar with the λ -calculus [5].

Definition 2.14 ($\lambda\mathbf{x}$ [10]) The set of $\lambda\mathbf{x}$ -terms is defined by the grammar:

$$M, N ::= x \mid \lambda x.M \mid MN \mid M\langle x := N \rangle$$

A term of the form $M\langle x := N \rangle$ is called a *closure*. A term which contains no closure is called a *pure term*.

The set of *free* and *bound* variables in a term M , $fv(M)$ and $bv(M)$, are defined by:

$$\begin{aligned} fv(x) &= \{x\} & bv(x) &= \emptyset \\ fv(\lambda x.M) &= fv(M) \setminus \{x\} & bv(\lambda x.M) &= bv(M) \cup \{x\} \\ fv(MN) &= fv(M) \cup fv(N) & bv(MN) &= bv(M) \cup bv(N) \\ fv(M\langle x := N \rangle) &= fv(M) \setminus \{x\} \cup fv(N) & bv(M\langle x := N \rangle) &= bv(M) \cup bv(N) \cup \{x\} \end{aligned}$$

We accept Barendregt's convention on free and bound variables, so assume that no variable appears both free and bound in a term.

Definition 2.15 (REDUCTION IN $\lambda\mathbf{x}$ [10]) Reduction in $\lambda\mathbf{x}$ is defined through the rules:

$$\begin{aligned} (\mathbf{B}) : & \quad (\lambda x.M)P \rightarrow M\langle x := P \rangle \\ (\mathbf{Var}) : & \quad x\langle x := P \rangle \rightarrow P \\ (\mathbf{VarK}) : & \quad y\langle x := P \rangle \rightarrow y, \quad (y \neq x) \\ (\mathbf{App}) : & \quad (MN)\langle x := P \rangle \rightarrow M\langle x := P \rangle N\langle x := P \rangle \\ (\mathbf{Abs}) : & \quad (\lambda y.M)\langle x := P \rangle \rightarrow \lambda y.M\langle x := P \rangle \end{aligned}$$

The $\lambda\mathbf{x}_{gc}$ calculus is defined by adding a rule for garbage collection. We point out that this rule requires the use of a more complicated side-condition.

$$(\mathbf{gc}) \quad M\langle x := P \rangle \rightarrow M, \quad (x \notin fv(M))$$

In fact, the rules (\mathbf{VarK}) and (\mathbf{gc}) overlap.

To illustrate the expressive power of \mathcal{X} , we now define the direct encoding of $\lambda\mathbf{x}$ -terms into \mathcal{X} , inspired by Gentzen's encoding of natural deduction into the sequent calculus [14].

Definition 2.16 (INTERPRETING THE λ -CALCULUS [2]) The interpretation of λ -terms into terms of \mathcal{X} in the context α , $\llbracket M \rrbracket_\alpha^G$, is defined by:

$$\begin{aligned} \llbracket x \rrbracket_\alpha^G &= \langle x \cdot \alpha \rangle \\ \llbracket \lambda x.M \rrbracket_\alpha^G &= \widehat{x} \llbracket M \rrbracket_\beta^G \widehat{\beta} \cdot \alpha, & (\beta \text{ fresh}) \\ \llbracket MN \rrbracket_\alpha^G &= \llbracket M \rrbracket_\gamma^G \widehat{\gamma} \dagger \widehat{x} (\llbracket N \rrbracket_\beta^G \widehat{\beta} [x] \widehat{y} \langle y \cdot \alpha \rangle), & (x, y, \beta, \gamma \text{ fresh}) \\ \llbracket M\langle x := N \rangle \rrbracket_\alpha^G &= \llbracket M \rrbracket_\gamma^G \widehat{\gamma} \backslash \widehat{x} \llbracket M \rrbracket_\alpha^G, & (\gamma \text{ fresh}) \end{aligned}$$

For this encoding, [2] shows:

Theorem 2.17 ([2]) *If $\Gamma \vdash_\lambda M : A$, then $\llbracket M \rrbracket_\alpha^G : \Gamma \vdash \alpha : A$.*

In [2], also the following relation is shown between reduction in $\lambda\mathbf{x}$ and \mathcal{X} :

Theorem 2.18 ([2]) *i) If $M \rightarrow_x N$, then $\llbracket M \rrbracket_\alpha^G \rightarrow^* \llbracket N \rrbracket_\alpha^G$.*

ii) If $M \rightarrow_{\beta} N$, then $\llbracket M \rrbracket_{\alpha}^G \rightarrow^* \llbracket N \rrbracket_{\alpha}^G$.

Alternatively, we can consider Prawitz' encoding of natural deduction into classical logic [21, 23]. This encoding preserves normal forms when interpreting a particular test case (i.e. if a λ -term is in normal form, then the interpretation in the \mathcal{X} -calculus is also in normal form), and will also allow us to indirectly obtain a cost of using the \mathcal{X} -calculus to reduce arbitrary λ -terms in Section 6.

Definition 2.19 (PRAWITZ' INTERPRETATION) There are three parts to the interpretation. We use the symbol ' \cdot ' to represent the list concatenation operator and ' $[]$ ' for the empty list. In addition, we use the symbols L to represent lists of λ -terms, and the symbols M, N, P to represent arbitrary λ -terms.

$$\begin{aligned}
\llbracket x \rrbracket_{\alpha}^P &= \langle x \cdot \alpha \rangle \\
\llbracket \lambda x. M \rrbracket_{\alpha}^P &= \widehat{x} \llbracket M \rrbracket_{\beta}^P \widehat{\beta} \cdot \alpha \\
\llbracket MN \rrbracket_{\alpha}^P &= \llbracket MN, [] \rrbracket_{\alpha}^P \\
\llbracket M \langle x := N \rangle \rrbracket_{\alpha}^P &= \llbracket N \rrbracket_{\gamma}^P \widehat{\gamma} \backslash \widehat{x} \llbracket M \rrbracket_{\alpha}^P \\
\llbracket x N, L \rrbracket_{\alpha}^P &= \llbracket N \rrbracket_{\beta}^P \widehat{\beta} [x] \widehat{y} \llbracket L \rrbracket_{\alpha}^P y \\
\llbracket (\lambda x. M) N, L \rrbracket_{\alpha}^P &= \llbracket \lambda x. M \rrbracket_{\beta}^P \widehat{\beta} \dagger \widehat{y} (\llbracket N \rrbracket_{\gamma}^P \widehat{\gamma} [y] \widehat{z} \llbracket L \rrbracket_{\alpha}^P z) \\
\llbracket MNP, L \rrbracket_{\alpha}^P &= \llbracket MN, P : L \rrbracket_{\alpha}^P \\
\llbracket [] \rrbracket_{\alpha}^P x &= \langle x \cdot \alpha \rangle \\
\llbracket M : L \rrbracket_{\alpha}^P x &= \llbracket M \rrbracket_{\beta}^P \widehat{\beta} [x] \widehat{y} \llbracket L \rrbracket_{\alpha}^P y
\end{aligned}$$

For this encoding, we can show the expected results.

Theorem 2.20 i) If $M \rightarrow_{\mathcal{X}} N$, then $\llbracket M \rrbracket_{\alpha}^P \rightarrow^* \llbracket N \rrbracket_{\alpha}^P$.

ii) If $M \rightarrow_{\beta}^* N$, then $\llbracket M \rrbracket_{\alpha}^P \rightarrow^* \llbracket N \rrbracket_{\alpha}^P$.

Proof: By induction on the definition of one-step explicit reduction, of which we show the more interesting cases.

$$\begin{aligned}
(\lambda x. M) N \rightarrow_{\beta} M \langle x := N \rangle : \llbracket (\lambda x. M) N \rrbracket_{\alpha}^P &\triangleq \\
\llbracket (\lambda x. M) N, [] \rrbracket_{\alpha}^P &\triangleq \\
\llbracket \lambda x. M \rrbracket_{\beta}^P \widehat{\beta} \dagger \widehat{y} (\llbracket N \rrbracket_{\gamma}^P \widehat{\gamma} [y] \widehat{z} \llbracket [] \rrbracket_{\alpha}^P z) &\triangleq \\
(\widehat{x} \llbracket M \rrbracket_{\delta}^P \widehat{\delta} \cdot \beta) \widehat{\beta} \dagger \widehat{y} (\llbracket N \rrbracket_{\gamma}^P \widehat{\gamma} [y] \widehat{z} \langle z \cdot \alpha \rangle) &\rightarrow (\text{exp-imp}) \\
\llbracket N \rrbracket_{\gamma}^P \widehat{\gamma} \dagger \widehat{x} (\llbracket M \rrbracket_{\delta}^P \widehat{\delta} \dagger \widehat{z} \langle z \cdot \alpha \rangle) &\rightarrow (2.9) \\
\llbracket N \rrbracket_{\gamma}^P \widehat{\gamma} \dagger \widehat{x} \llbracket M \rrbracket_{\alpha}^P &\rightarrow (\backslash a) \\
\llbracket N \rrbracket_{\gamma}^P \widehat{\gamma} \backslash \widehat{x} \llbracket M \rrbracket_{\alpha}^P &\triangleq \llbracket M \langle x := N \rangle \rrbracket_{\alpha}^P
\end{aligned}$$

$$x \langle x := P \rangle \rightarrow P : \llbracket x \langle x := P \rangle \rrbracket_{\alpha}^P \triangleq \llbracket P \rrbracket_{\gamma}^P \widehat{\gamma} \backslash \widehat{x} \langle x \cdot \alpha \rangle \rightarrow^* (\backslash d, 2.9) \llbracket P \rrbracket_{\alpha}^P$$

$$y \langle x := P \rangle \rightarrow y, \text{ if } y \neq x : \llbracket y \langle x := P \rangle \rrbracket_{\alpha}^P \triangleq \llbracket P \rrbracket_{\gamma}^P \widehat{\gamma} \backslash \widehat{x} \langle y \cdot \alpha \rangle \rightarrow^* (2.9) \langle y \cdot \alpha \rangle \triangleq \llbracket y \rrbracket_{\alpha}^P$$

$$\begin{aligned}
(MN) \langle x := P \rangle \rightarrow M \langle x := P \rangle N \langle x := P \rangle : \llbracket (MN) \langle x := P \rangle \rrbracket_{\alpha}^P &\triangleq \\
\llbracket P \rrbracket_{\gamma}^P \widehat{\gamma} \backslash \widehat{x} \llbracket MN \rrbracket_{\alpha}^P &\triangleq \llbracket P \rrbracket_{\gamma}^P \widehat{\gamma} \backslash \widehat{x} \llbracket MN, [] \rrbracket_{\alpha}^P
\end{aligned}$$

We now distinguish the cases:

$$\begin{aligned}
M = x : \llbracket P \rrbracket_{\gamma}^P \widehat{\gamma} \backslash \widehat{x} \llbracket x N, [] \rrbracket_{\alpha}^P &\triangleq \\
\llbracket P \rrbracket_{\gamma}^P \widehat{\gamma} \backslash \widehat{x} (\llbracket N \rrbracket_{\beta}^P \widehat{\beta} [x] \widehat{z} \langle z \cdot \alpha \rangle) &\rightarrow (\backslash \text{imp-outs}), (2.9) \\
\llbracket P \rrbracket_{\gamma}^P \widehat{\gamma} \backslash \widehat{v} ((\llbracket P \rrbracket_{\gamma}^P \widehat{\gamma} \backslash \widehat{x} \llbracket N \rrbracket_{\beta}^P) \widehat{\beta} [v] \widehat{z} \langle z \cdot \alpha \rangle) &\triangleq \\
\llbracket P \rrbracket_{\gamma}^P \widehat{\gamma} \backslash \widehat{v} \llbracket v N \langle x := P, [] \rangle \rrbracket_{\alpha}^P &\triangleq \\
\llbracket v N \langle x := P \rangle \langle v := P \rangle, [] \rrbracket_{\alpha}^P &\triangleq \llbracket x N \langle x := P \rangle, [] \rrbracket_{\alpha}^P
\end{aligned}$$

$$\begin{aligned}
M = y, y \neq x : \llbracket P \rrbracket_{\gamma}^p \hat{\gamma} \hat{x} \llbracket yN, [] \rrbracket_{\alpha}^p &\triangleq \\
\llbracket P \rrbracket_{\gamma}^p \hat{\gamma} \hat{x} (\llbracket N \rrbracket_{\beta}^p \hat{\beta} [y] \hat{z} \langle z \cdot \alpha \rangle) &\rightarrow (\backslash \text{imp-ins}), (2.9) \\
(\llbracket P \rrbracket_{\gamma}^p \hat{\gamma} \hat{x} \llbracket N \rrbracket_{\beta}^p \hat{\beta} [y] \hat{z} \langle z \cdot \alpha \rangle) &\triangleq \\
\llbracket N \langle x := P \rangle \rrbracket_{\beta}^p \hat{\beta} [y] \hat{z} \langle z \cdot \alpha \rangle &\triangleq \llbracket yN \langle x := P \rangle, [] \rrbracket_{\alpha}^p
\end{aligned}$$

The other cases follow by induction.

The second part follows from the fact that \rightarrow_x implements \rightarrow_{β} . □

3 Optimising reduction

We should point out that, using the rules above, not all typeable terms are strongly normalisable. For example, to allow for the propagation of cuts over cuts immediately leads to non-termination, since we can always choose the outermost cut as the one to contract. Although the notion of cut-elimination as proposed here has no rule that would allow this behaviour, it can be mimicked, which can lead to non-termination for typeable terms, as already observed by Urban [24].

Take $P\hat{\alpha} \dagger \hat{x} (\langle x \cdot \beta \rangle \hat{\beta} \dagger \hat{z}Q)$, such that $x \notin fs(Q), \beta \notin fp(P)$, and P, Q pure, then:

$$\begin{aligned}
P\hat{\alpha} \dagger \hat{x} (\langle x \cdot \beta \rangle \hat{\beta} \dagger \hat{z}Q) &\rightarrow^* (\backslash a), (\backslash \text{cut}) \\
(P\hat{\alpha} \backslash \hat{x} \langle x \cdot \beta \rangle) \hat{\beta} \dagger \hat{z} (P\hat{\alpha} \backslash \hat{x} Q) &\rightarrow^* (\backslash d), (2.9) \\
(P\hat{\alpha} \dagger \hat{x} \langle x \cdot \beta \rangle) \hat{\beta} \dagger \hat{z} Q &\rightarrow^* (a \cancel{\prime}), (\text{cut} \cancel{\prime}) \\
(P\hat{\beta} \cancel{\prime} \hat{z}Q) \hat{\alpha} \dagger \hat{x} (\langle x \cdot \beta \rangle \hat{\beta} \cancel{\prime} \hat{z}Q) &\rightarrow^* (d \cancel{\prime}), (2.9) \\
P\hat{\alpha} \dagger \hat{x} (\langle x \cdot \beta \rangle \hat{\beta} \dagger \hat{z}Q) &
\end{aligned}$$

(example communicated by Alexander J. Summers)

Urban gives a solution for this unwanted reduction behaviour, and shows it sufficient to obtain strong-normalisation of typeable terms. He adds the rules

$$\begin{aligned}
(P\hat{\alpha} \dagger \hat{x} \langle x \cdot \beta \rangle) \hat{\beta} \cancel{\prime} \hat{y} Q &\rightarrow (P\hat{\beta} \cancel{\prime} \hat{y} Q) \hat{\alpha} \dagger \hat{y} Q \\
P\hat{\alpha} \backslash \hat{x} (\langle x \cdot \beta \rangle \hat{\beta} \dagger \hat{y} Q) &\rightarrow P\hat{\alpha} \dagger \hat{y} (P\hat{\alpha} \backslash \hat{x} Q)
\end{aligned}$$

and gives them priority over the rules $(\text{cut} \cancel{\prime})$ and $(\backslash \text{cut})$ by changing those to

$$\begin{aligned}
(P\hat{\alpha} \dagger \hat{x} Q) \hat{\beta} \cancel{\prime} \hat{y} R &\rightarrow (P\hat{\beta} \cancel{\prime} \hat{y} R) \hat{\alpha} \dagger \hat{x} (Q\hat{\beta} \cancel{\prime} \hat{u} R), \quad Q \neq \langle x \cdot \beta \rangle \\
P\hat{\alpha} \backslash \hat{x} (Q\hat{\beta} \dagger \hat{y} R) &\rightarrow (P\hat{\alpha} \backslash \hat{x} Q) \hat{\beta} \dagger \hat{y} (P\hat{\alpha} \backslash \hat{x} R), \quad Q \neq \langle x \cdot \beta \rangle
\end{aligned}$$

Notice that the side-condition $Q \neq \langle x \cdot \beta \rangle$ is quite different in character from the rules for \mathcal{X} we presented above, in that now *syntactic equality* between terms is tested, rather than just a syntactic property of a term.

We can remedy this by avoiding the deactivation of cuts altogether.

Definition 3.1 We can remove the rules $(d \cancel{\prime})$ and $(\backslash d)$, and add the following rules:

$$\begin{aligned}
(\text{flip} \backslash) : \quad \langle z \cdot \alpha \rangle \hat{\alpha} \cancel{\prime} \hat{x} P &\rightarrow \langle z \cdot \alpha \rangle \hat{\alpha} \backslash \hat{x} P, & (P \text{ does not introduce } x) \\
(\cancel{\prime} \text{imp}) : \quad \langle y \cdot \alpha \rangle \hat{\alpha} \cancel{\prime} \hat{x} (P\hat{\delta} [x] \hat{z}Q) &\rightarrow P\hat{\delta} [y] \hat{z}Q, & (x \text{ introduced}) \\
(\cancel{\prime} \text{cap}) : \quad \langle z \cdot \alpha \rangle \hat{\alpha} \cancel{\prime} \hat{x} \langle x \cdot \beta \rangle &\rightarrow \langle z \cdot \beta \rangle \\
(\cancel{\prime} \text{flip}) : \quad P\hat{\alpha} \backslash \hat{x} \langle x \cdot \beta \rangle &\rightarrow P\hat{\alpha} \cancel{\prime} \hat{x} \langle x \cdot \beta \rangle, & (P \text{ does not introduce } \alpha) \\
(\text{exp} \backslash) : \quad (\hat{y}P\hat{\delta} \cdot \alpha) \hat{\alpha} \backslash \hat{x} \langle x \cdot \gamma \rangle &\rightarrow \hat{y}P\hat{\delta} \cdot \gamma, & (\alpha \text{ introduced}) \\
(\text{cap} \backslash) : \quad \langle x \cdot \alpha \rangle \hat{\alpha} \backslash \hat{y} \langle y \cdot \beta \rangle &\rightarrow \langle x \cdot \beta \rangle
\end{aligned}$$

We do not need to check if a term matches another, nor need to give priority to rules.

This new set of rules does not change the end result of reduction:

Theorem 3.2 *Let \rightarrow_o denote the notion of reduction obtained by the changes as suggested above, and assume P and Q are pure.*

- i) If $P \rightarrow_o^* Q$, then $P \rightarrow_{\mathcal{X}}^* Q$.*
- ii) If $P \rightarrow_{\mathcal{X}}^* Q$ and Q is in normal form, then $P \rightarrow_o^* Q$.*

Proof: Easy. □

In Section 6, we will present an evaluation strategy that will choose to run the inactive cut that is created by the deactivation rule immediately after its creation, thus side-stepping the problem.

The set of reduction rules can be optimised further than just those expressed in Lemma 2.9. For example, the applicability of the garbage collection rules is limited, since they both involve pure terms. However, we are able to generalise these results to include terms with active cuts.

We aim to add more generic garbage collection rules; in fact, we will show their admissibility below (Theorem 3.5), for which we first need to show a number of results.

Lemma 3.3 *i) For all P, Q pure, there exists an R pure such that $P\hat{\alpha} \not\prec \hat{\alpha}Q \rightarrow^* R$.*

ii) For all P, Q pure, there exists an R pure such that $P\hat{\alpha} \setminus \hat{\alpha}Q \rightarrow^ R$.*

Proof: *i)* By induction on structure of terms. We highlight one case:

$$\begin{aligned} P = \hat{y}P'\hat{\beta}\cdot\alpha : (\hat{y}P'\hat{\beta}\cdot\alpha)\hat{\alpha} \not\prec \hat{\alpha}Q &\rightarrow \\ (\hat{y}(P'\hat{\alpha} \not\prec \hat{\alpha}Q)\hat{\beta}\cdot\alpha)\hat{\alpha} \dagger \hat{\alpha}Q &\rightarrow^* \text{(IH, } R \text{ pure)} \\ (\hat{y}R\hat{\beta}\cdot\alpha)\hat{\alpha} \dagger \hat{\alpha}Q & \end{aligned}$$

Notice that this last term is pure.

ii) By induction on structure of terms. We highlight again one case:

$$\begin{aligned} Q = Q_1\hat{\beta}[x]\hat{y}Q_2 : P\hat{\alpha} \setminus \hat{\alpha}(Q_1\hat{\beta}[x]\hat{y}Q_2) &\rightarrow (Q_1, Q_2 \text{ pure}) \\ P\hat{\alpha} \dagger \hat{v}((P\hat{\alpha} \setminus \hat{\alpha}Q_1)\hat{\beta}[v]\hat{y}(P\hat{\alpha} \setminus \hat{\alpha}Q_2)) &\rightarrow^* \text{(IH, } R_1, R_2 \text{ pure)} \\ P\hat{\alpha} \dagger \hat{v}(R_1\hat{\beta}[v]\hat{y}R_2) & \end{aligned}$$

Notice again that this last term is pure. □

We can now use this lemma to give a stronger result.

Lemma 3.4 *For all \mathcal{X} -terms P , there exists a reduction path $P \rightarrow^* P'$, with P' pure.*

Proof: For each active cut in a term, we define its *depth* d as the distance (calculated in nodes) from the root of the tree that represents the term. For any particular depth of the term, we define its *group size*, g , as the number of active cuts at that depth of the tree. We define the *class*, c , of a term as the pair of the depth of the innermost cut and its group size: $c = \langle d, g \rangle$. We finish the proof by lexicographic induction on the class of a term.

If P is pure, then $P \equiv P'$. Otherwise: take the set (of size g) of innermost active cuts at depth d . There are two cases to consider:

$g > 1$: take any term T in this set, $T = R\hat{\alpha} \not\prec \hat{\alpha}Q$ (or $T = R\hat{\alpha} \setminus \hat{\alpha}Q$). Then R and Q are pure (being involved in an innermost flagged cut). By Lemma 3.3, we know $T \rightarrow^* S$ (pure).

This eliminates the active cut from the proof, so the group-size (g) of the term reduces by one, and the class decreases.

$g = 1$: the active cut is eliminated from the term by Lemma 3.3. Since there are no more

active cuts at this level, the depth of the innermost cut decreases (to the next lowest), and the class reduces. \square

We are now ready to justify some more general garbage collection rules, that essentially equates the nets $P\hat{\alpha} \not\wedge \hat{x}Q$ and P , provided $\alpha \notin fp(P)$.

Lemma 3.5 (GENERALISED GARBAGE COLLECTION)

i) If $\alpha \notin fp(P)$, then $P\hat{\alpha} \not\wedge \hat{x}Q \stackrel{\mathcal{X}}{=} P$.

ii) If $x \notin fs(Q)$, then $P\hat{\alpha} \not\wedge \hat{x}Q \stackrel{\mathcal{X}}{=} Q$.

Proof: i) For the second part ($nf(P) \subseteq nf(P\hat{\alpha} \not\wedge \hat{x}Q)$), take $T \in nfP$, reduce $P\hat{\alpha} \not\wedge \hat{x}Q$ to $T\hat{\alpha} \not\wedge \hat{x}Q$, remark that T is pure, and apply Lemma 2.9.

For the first, we show that if $P\hat{\alpha} \not\wedge \hat{x}Q \rightarrow^* T$, then $P \rightarrow^* T$, where T is a normal form. We achieve this by showing that we can run a reduction on P , mimicking a reduction taking place on $P\hat{\alpha} \not\wedge \hat{x}Q$ by essentially ignoring all reductions inside Q ; the only problem might be when the presence of $[\]\hat{\alpha} \not\wedge \hat{x}Q$ disturbs the reduction behaviour.

The proof completes by co-induction (we only show some interesting cases):

- $P = (P_1\hat{\beta} [v] \hat{y}P_2)\hat{\gamma} \dagger \hat{z}P_3$. We can run $((P_1\hat{\beta} [v] \hat{y}P_2)\hat{\gamma} \dagger \hat{z}P_3)\hat{\alpha} \not\wedge \hat{x}Q$ in a number of ways. Any reduction inside P_1, P_2 or P_3 is dealt with by induction, so we can focus on the cuts involved. Assume we apply rule (*cut*) to propagate the outermost cut and obtain

$$((P_1\hat{\beta} [v] \hat{y}P_2)\hat{\alpha} \not\wedge \hat{x}Q)\hat{\gamma} \dagger \hat{z}(P_3\hat{\alpha} \not\wedge \hat{x}Q)$$

Now the top-most (inactive) cut can be activated in two directions; let's go left:

$$((P_1\hat{\beta} [v] \hat{y}P_2)\hat{\alpha} \not\wedge \hat{x}Q)\hat{\gamma} \not\wedge \hat{z}(P_3\hat{\alpha} \not\wedge \hat{x}Q)$$

This (outermost) activated cut cannot propagate, since the cut directly underneath it is active; propagating that first gives

$$((P_1\hat{\alpha} \not\wedge \hat{x}Q)\hat{\beta} [v] \hat{y}(P_2\hat{\alpha} \not\wedge \hat{x}Q))\hat{\gamma} \not\wedge \hat{z}(P_3\hat{\alpha} \not\wedge \hat{x}Q)$$

Now the top-cut can propagate, to give

$$((P_1\hat{\alpha} \not\wedge \hat{x}Q)\hat{\gamma} \not\wedge \hat{z}(P_3\hat{\alpha} \not\wedge \hat{x}Q))\hat{\beta} [v] \hat{y}((P_2\hat{\alpha} \not\wedge \hat{x}Q)\hat{\gamma} \not\wedge \hat{z}(P_3\hat{\alpha} \not\wedge \hat{x}Q))$$

By induction we can mimic $P_i\hat{\alpha} \not\wedge \hat{x}Q$ by P_i , for $i \in \{1, 2, 3\}$. We can simulate this particular reduction on P as follows:

$$(P_1\hat{\beta} [v] \hat{y}P_2)\hat{\gamma} \dagger \hat{z}P_3 \rightarrow^* (a \not\wedge), (imp \not\wedge) (P_1\hat{\gamma} \not\wedge \hat{z}P_3)\hat{\beta} [v] \hat{y}(P_2\hat{\gamma} \not\wedge \hat{z}P_3)$$

- $P = (\hat{y}P_1\hat{\beta} \cdot \gamma)\hat{\gamma} \dagger \hat{z}P_2$. As above, a reduction inside P_1 or P_2 creates no problems, so we can focus on the cuts involved. When we propagate the top cut, we get

$$((\hat{y}P_1\hat{\beta} \cdot \gamma)\hat{\alpha} \not\wedge \hat{x}Q)\hat{\gamma} \dagger \hat{z}(P_2\hat{\alpha} \not\wedge \hat{x}Q)$$

If we now left-activate the top cut, similar to above, we can only propagate the innermost cut, and obtain:

$$(\hat{y}(P_1\hat{\alpha} \not\wedge \hat{x}Q)\hat{\beta} \cdot \gamma)\hat{\gamma} \not\wedge \hat{z}(P_2\hat{\alpha} \not\wedge \hat{x}Q)$$

Now applying rule (*exp-outs*) will give:

$$(\hat{y}((P_1\hat{\alpha} \not\wedge \hat{x}Q)\hat{\gamma} \not\wedge \hat{z}(P_2\hat{\alpha} \not\wedge \hat{x}Q))\hat{\beta} \cdot \delta)\hat{\delta} \dagger \hat{z}(P_2\hat{\alpha} \not\wedge \hat{x}Q)$$

Now, if γ is introduced in $\hat{y}P_1\hat{\beta}\cdot\gamma$, then γ does not appear free inside P_1 nor in Q , and, by induction, we can assume that $(P_1\hat{\alpha}\not\sim\hat{x}Q)\hat{\gamma}\not\sim\hat{z}(P_2\hat{\alpha}\not\sim\hat{x}Q)$ can be simulated by P_1 , and that $P_2\hat{\alpha}\not\sim\hat{x}Q$ can be simulated by P_2 . We can mimic this reduction with a reduction path of length zero, performing an α -conversion:

$$(\hat{y}P_1\hat{\beta}\cdot\gamma)\hat{\gamma}\dagger\hat{z}P_2 \rightarrow (\hat{y}P_1\hat{\beta}\cdot\delta)\hat{\delta}\dagger\hat{z}P_2$$

If γ is not introduced, we still can assume that $P_i\hat{\alpha}\not\sim\hat{x}Q$ can be simulated by P_i , for $i \in \{1,2\}$, and we can simulate this reduction on P via a series of steps:

$$(\hat{y}P_1\hat{\beta}\cdot\gamma)\hat{\gamma}\dagger\hat{z}P_2 \rightarrow (\hat{y}P_1\hat{\beta}\cdot\gamma)\hat{\gamma}\not\sim\hat{z}P_2 \rightarrow (\hat{y}(P_1\hat{\gamma}\not\sim\hat{z}P_2)\hat{\beta}\cdot\delta)\hat{\delta}\dagger\hat{z}P_2$$

- All other cases are shown in a similar fashion.

So every reduction to a normal form, starting from $P\hat{\alpha}\not\sim\hat{x}Q$, can be mimicked by reducing P , so every normal form of $P\hat{\alpha}\not\sim\hat{x}Q$, can be reached from P .

ii) Similar. □

This result now helps to justify more general deactivation rules.

Theorem 3.6 (GENERALISED DEACTIVATION) i) $P\hat{\alpha}\not\sim\hat{x}Q \stackrel{x}{\underline{=}} P\hat{\alpha}\dagger\hat{x}Q$, if P introduces α .
ii) $P\hat{\alpha}\not\sim\hat{x}Q \stackrel{x}{\underline{=}} P\hat{\alpha}\dagger\hat{x}Q$, if Q introduces x .

Proof: i) If P introduces α , we have two cases:

$$P = \langle x\cdot\alpha \rangle : \text{By rule } (d\not\sim).$$

$$P = \hat{x}P'\hat{\beta}\cdot\alpha : \text{Then } \alpha \notin fp(P), \text{ and}$$

$$\begin{aligned} (\hat{x}P'\hat{\beta}\cdot\alpha)\hat{\alpha}\not\sim\hat{x}Q &\rightarrow (\text{exp-outs}\not\sim) \\ (\hat{x}(P'\hat{\alpha}\not\sim\hat{x}Q)\hat{\beta}\cdot\gamma)\hat{\gamma}\dagger\hat{x}Q &\stackrel{x}{\underline{=}} (3.5) \\ (\hat{x}P'\hat{\beta}\cdot\gamma)\hat{\gamma}\dagger\hat{x}Q &=_{\alpha} (\hat{x}P'\hat{\beta}\cdot\alpha)\hat{\alpha}\dagger\hat{x}Q \end{aligned}$$

ii) If Q introduces x , we have two cases:

$$Q = \langle x\cdot\beta \rangle : \text{By rule } (\not\sim d).$$

$$Q = Q_1\hat{\beta}[x]\hat{y}Q_2 : \text{Then } x \notin fs(Q_1, Q_2), \text{ and}$$

$$\begin{aligned} P'\hat{\alpha}\not\sim\hat{x}(Q_1\hat{\beta}[x]\hat{y}Q_2) &\rightarrow (\not\sim \text{imp-outs}) \\ P'\hat{\alpha}\dagger\hat{v}((P'\hat{\alpha}\not\sim\hat{x}Q_1)\hat{\beta}[v]\hat{y}(P'\hat{\alpha}\not\sim\hat{x}Q_2)) &\stackrel{x}{\underline{=}} (3.5) \\ P'\hat{\alpha}\dagger\hat{v}(Q_1\hat{\beta}[v]\hat{y}Q_2) &=_{\alpha} P'\hat{\alpha}\dagger\hat{x}(Q_1\hat{\beta}[x]\hat{y}Q_2) \quad \square \end{aligned}$$

Using coinduction, similar to the results above, we can also show:

Lemma 3.7 (GENERALISED RENAMING) i) $P\hat{\alpha}\not\sim\hat{x}\langle x\cdot\beta \rangle \stackrel{x}{\underline{=}} P[\beta/\alpha]$.

ii) $\langle y\cdot\alpha \rangle\hat{\alpha}\not\sim\hat{x}P \stackrel{x}{\underline{=}} P[y/x]$.

These results give that we can safely extend the notion of reduction on terms by adding:

$$\begin{aligned} (d\not\sim) : P\hat{\alpha}\not\sim\hat{x}Q &\rightarrow P\hat{\alpha}\dagger\hat{x}Q && (P \text{ introduces } \alpha) \\ (d\not\sim) : P\hat{\alpha}\not\sim\hat{x}Q &\rightarrow P\hat{\alpha}\dagger\hat{x}Q && (Q \text{ introduces } x) \\ (\not\sim gc) : P\hat{\alpha}\not\sim\hat{x}Q &\rightarrow P && (\alpha \notin fp(P)) \\ (gc\not\sim) : P\hat{\alpha}\not\sim\hat{x}Q &\rightarrow Q && (x \notin fs(Q)) \end{aligned}$$

thereby replacing the original deactivation rules, and $(cap\not\sim)$ and $(\not\sim cap)$.

Adding these more generic reduction rules gave another significant improvement on the efficiency of reduction, on which we report in Section 6.

4 A Term Graph Rewriting System for \mathcal{X}

As mentioned in the introduction, \mathcal{X} corresponds to sequent calculus rather than natural deduction, leading to a reduction mechanism that is quite different from the applicative style of the λ -calculus. In fact, it is not specified through (implicit) substitution, but as a *conditional term rewriting system*; the non-standard aspect with respect to rewriting is a notion of *binding*.


The choice to use term graph rewriting for the implementation of an interpreter of \mathcal{X} was therefore the natural one: it provides the best platform for an efficient, elegant, and transparent implementation of a term rewriting system. The technique applied is the conventional one of [22, 6, 7], where terms and rewrite rules are *lifted* to graphs, and we use the standard *match, build, link, re-direct, and garbage collection* approach. By the process of lifting, the connectors appear only once in the generated graph, which immediately introduces sharing. Rewrite rules also become graphs with *two* sub-graphs that each possess a root, and are joined via shared leaves.

Term graphs are defined by:

Definition 4.1 (TERM GRAPHS) Using the signature $\{\text{Cap}, \text{Exp}, \text{Imp}, \text{Cut}, \text{CutL}, \text{CutR}\}$, an infinite set of *graph variables* ranged over by P, Q, R, \dots , an infinite set of *labels* ranged over by k, l, m, n, \dots and the set of *connector variables*, ranged over by x, y, z, \dots and $\alpha, \beta, \gamma, \dots$, we define (ordered) graphs by the following grammar:

$$\begin{aligned} G ::= & k : P \mid k : \text{Cap}(l : x, m : \alpha) \mid k : \text{Exp}(l : x, G, m : \alpha, n : \beta) \mid \\ & k : \text{Imp}(G_1, l : \alpha, m : y, n : x, G_2) \mid k : \text{Cut}(G_1, l : \alpha, m : x, G_2) \mid \\ & k : \text{CutL}(G_1, l : \alpha, m : x, G_2) \mid k : \text{CutR}(G_1, l : \alpha, m : x, G_2) \end{aligned}$$

which in graphs get represented by:



Since the graphs for left and right-activated cuts are much like unactivated cuts, we will only treat them separately when necessary.

In the second graph, for example, the connectors x, α and β could appear in G ; then edges would point out from G to these connectors. We now give a formal definition of interpreting terms in \mathcal{X} as graphs (using the notation of [9]).

Definition 4.2 (GRAPH INTERPRETATION) *i*) For each term, its *graph interpretation*, $\llbracket \cdot \rrbracket_G$, expressed as a pair consisting of a label for the *root* of the graph and its *edges* and *sub-graphs* is inductively defined by in Figure 1.


ii) We define the set of *initial* \mathcal{X} -graphs as the image of \mathcal{X} terms under $\llbracket \cdot \rrbracket_G$.

Notice that, by α -conversion, we can always assume that all bound connectors in a term have different names, and therefore will keep bound connectors separate when building the graph interpretation.

$$\begin{aligned}
\llbracket x \rrbracket_G &= \langle x \mid \emptyset \rangle \\
\llbracket \alpha \rrbracket_G &= \langle \alpha \mid \emptyset \rangle \\
\llbracket \langle x \cdot \alpha \rangle \rrbracket_G &= \langle r \mid \{r : \mathbf{Cap}(r_1, r_2)\} \cup G_1 \cup G_2 \rangle, \\
\text{where } \langle r_1 \mid G_1 \rangle &= \llbracket x \rrbracket_G, \\
\langle r_2 \mid G_2 \rangle &= \llbracket \alpha \rrbracket_G \\
\llbracket \hat{y} P \hat{\alpha} \cdot \beta \rrbracket_G &= \langle r \mid \{r : \mathbf{Exp}(r_1, r_2, r_3, r_4)\} \cup G_1 \cup G_2 \cup G_3 \cup G_4 \rangle \\
\text{where } \langle r_1 \mid G_1 \rangle &= \llbracket \beta \rrbracket_G \\
\langle r_2 \mid G_2 \rangle &= \llbracket y \rrbracket_G \\
\langle r_3 \mid G_3 \rangle &= \llbracket P \rrbracket_G \\
\langle r_4 \mid G_4 \rangle &= \llbracket \alpha \rrbracket_G \\
\llbracket P \hat{\alpha} [y] \hat{x} Q \rrbracket_G &= \langle r \mid \{r : \mathbf{Imp}(r_1, r_2, r_3, r_4, r_5)\} \\
&\quad \cup G_1 \cup G_2 \cup G_3 \cup G_4 \cup G_5 \rangle \\
\text{where } \langle r_1 \mid G_1 \rangle &= \llbracket P \rrbracket_G \\
\langle r_2 \mid G_2 \rangle &= \llbracket \alpha \rrbracket_G \\
\langle r_3 \mid G_3 \rangle &= \llbracket y \rrbracket_G \\
\langle r_4 \mid G_4 \rangle &= \llbracket x \rrbracket_G \\
\langle r_5 \mid G_5 \rangle &= \llbracket Q \rrbracket_G \\
\llbracket P \hat{\alpha} \dagger \hat{x} Q \rrbracket_G &= \langle r \mid \{r : \mathbf{Cut}(r_1, r_2, r_3, r_4)\} \cup G_1 \cup G_2 \cup G_3 \cup G_4 \rangle \\
\text{where } \langle r_1 \mid G_1 \rangle &= \llbracket P \rrbracket_G \\
\langle r_2 \mid G_2 \rangle &= \llbracket \alpha \rrbracket_G \\
\langle r_3 \mid G_3 \rangle &= \llbracket x \rrbracket_G \\
\langle r_4 \mid G_4 \rangle &= \llbracket Q \rrbracket_G
\end{aligned}$$

Figure 1: Term graph interpretation of terms

Example 4.1 $\llbracket \langle x \cdot \alpha \rangle \hat{\alpha} \dagger \hat{y} (\langle x \cdot \beta \rangle \hat{\beta} [y] \hat{z} \langle z \cdot \gamma \rangle) \rrbracket_G$ becomes



Definition 4.3 (\mathcal{X} -GRAPH REWRITING) *i*) The *lifting* of the reduction rules to term graph rewriting rules is expressed by first extending the interpretation of terms to graphs with the case for the (term) variables that occur in the rewrite rules:

$$\llbracket P \rrbracket_G = \langle P \mid \emptyset \rangle$$

and then to define:

$$\begin{aligned}
\llbracket \mathit{left} \rightarrow \mathit{right} \rrbracket_G &= \langle r_l \mid G_l \cup G_r \rangle \\
\text{where } \langle r_l \mid G_l \rangle &= \llbracket \mathit{left} \rrbracket_G \\
\langle r_r \mid G_r \rangle &= \llbracket \mathit{right} \rrbracket_G
\end{aligned}$$

Following [6, 7], these rules induce a notion $G \rightarrow_G G'$ of term graph rewriting.


ii) We define the set of \mathcal{X} -graphs by closure under rewriting of initial \mathcal{X} -graphs.

Notice that since $\llbracket x \rrbracket_G = \langle x \mid \emptyset \rangle$, and we can assume this to be unique, the left and right-hand side graphs for a rule are joined on the connectors.


Example 4.2 We give some term graph rules³:

³Notice that it is perhaps more common to use \perp for the nodes that can be matched against a graph or a connector node; we write the original names for readability and ease of definition of the interpretation function.


- $(exp) : (\hat{y}P\hat{\beta}\cdot\alpha)\hat{a}\dagger\hat{x}\langle x\cdot\gamma \rangle \rightarrow \hat{y}P\hat{\beta}\cdot\gamma$




- $(exp-imp) : (\hat{y}P\hat{\beta}\cdot\alpha)\hat{a}\dagger\hat{x}(Q\hat{\gamma}\dagger[x]\hat{z}R) \rightarrow Q\hat{\gamma}\dagger\hat{y}(P\hat{\beta}\dagger\hat{z}R)$



- $(exp-outs') : (\hat{y}Q\hat{\beta}\cdot\alpha)\hat{a}\dagger\hat{x}P \rightarrow (\hat{y}(Q\hat{\alpha}\dagger\hat{x}P)\hat{\beta}\cdot\gamma)\hat{\gamma}\dagger\hat{x}P, (\gamma \text{ fresh})$



- $(\lambda imp-outs) : P\hat{\alpha}\dagger\hat{x}(Q\hat{\beta}\dagger[x]\hat{y}R) \rightarrow P\hat{\alpha}\dagger\hat{z}((P\hat{\alpha}\dagger\hat{x}Q)\hat{\beta}\dagger[z]\hat{y}(P\hat{\alpha}\dagger\hat{x}R)),$
with z fresh (with apologies for the spaghetti):



As can be seen, the amount of nodes added to the graph is small in comparison to the complexity of the graph generated by the rewriting; notice, for example, that an application of the third rule ($exp-outs'$) would only add the nodes containing Cut , Exp , γ , and $CutL$ (that are accessible from the right-hand root r_r). Also, all edges coming into the node in the graph that is matched against the left-hand root r_l would be redirected into the new node Cut . The nodes containing $CutL$ and Exp accessible from the left-hand root would become potential garbage.

Example 4.3 Take the reduction:

$$\begin{aligned}
& (\hat{y}\langle y\cdot\beta \rangle\hat{\beta}\cdot\gamma)\hat{\gamma}\dagger\hat{x}\langle (x\cdot\delta)\hat{\delta}\dagger[x]\hat{v}\langle v\cdot\alpha \rangle \rangle && \rightarrow^* (\lambda a), (\lambda imp-outs), (\lambda d), (exp) \\
& (\hat{y}\langle y\cdot\beta \rangle\hat{\beta}\cdot\gamma)\hat{\gamma}\dagger\hat{z}((\hat{y}\langle y\cdot\beta \rangle\hat{\beta}\cdot\delta)\hat{\delta}\dagger[z]\hat{v}\langle v\cdot\alpha \rangle) && \rightarrow^* (exp-imp) \\
& (\hat{y}\langle y\cdot\beta \rangle\hat{\beta}\cdot\sigma)\hat{\sigma}\dagger\hat{y}\langle (y\cdot\beta)\hat{\beta}\dagger\hat{v}\langle v\cdot\alpha \rangle \rangle && \rightarrow^* (cap) \\
& (\hat{y}\langle y\cdot\beta \rangle\hat{\beta}\cdot\sigma)\hat{\sigma}\dagger\hat{y}\langle y\cdot\alpha \rangle && \rightarrow^* (exp) \\
& \hat{y}\langle y\cdot\beta \rangle\hat{\beta}\cdot\alpha &&
\end{aligned}$$

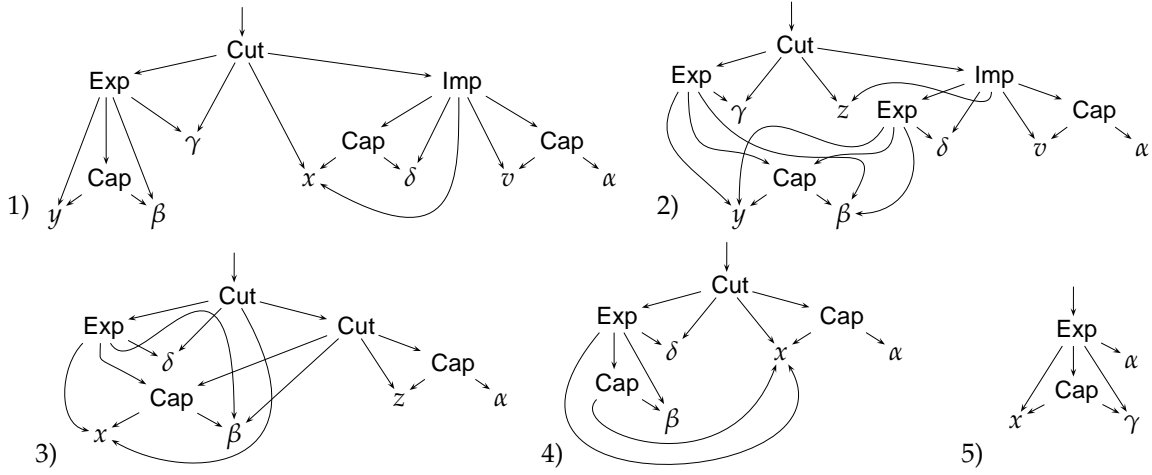


Figure 2: Graphs for Example 4.3

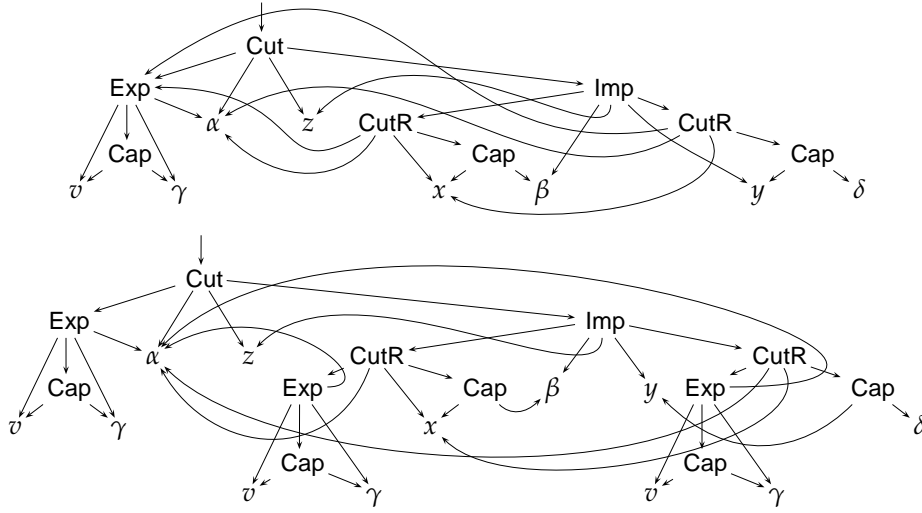


Figure 3: Unravelling a graph

The graph interpretation of this first term gives the first graph in Figure 2; reducing that graph using the corresponding term-graph rewriting rules generates the corresponding consecutive graphs.

In addition to the interpretation of terms to graphs, we would like an operation that transforms an \mathcal{X} -graph with sharing into one whose structure more closely resembles an \mathcal{X} -term. This is achieved by ‘unravelling’ the graph; copying out the shared nodes as far down as the connectors (which only appear once in a graph).

Definition 4.4 (CF [16]) *Unrv* G , the *unravelling* of a \mathcal{X} -graph G is obtained by traversing the graph top-down (notice that we have no cyclic structures), and copying, for all shared graphs, all nodes in that graph that are not free connectors.

Notice that both the set of *initial* \mathcal{X} -graphs and the image of the set of \mathcal{X} -graphs under *Unrv* are graphs containing sharing only at the level of connectors. This setup gives us a method of comparing an \mathcal{X} -term P with an \mathcal{X} -graph G , by comparing $\llbracket P \rrbracket_G$ with *Unrv* G . This will be useful for formulating results later in the paper.

Example 4.4 Take G to be the first graph in Figure 3, then $\text{Unrv } G$ is the second. Notice that the bound connectors v and γ within the shared graph

$$\text{Exp}(1:v, \text{Cap}(1,2), 2:\gamma, \alpha)$$

are copied out, but α is not, and that the in-degree of α increases from four to six.

We now have the following results.

Lemma 4.5 If $G_1 \rightarrow_G G_2$, then there exists G_3 such that $\text{Unrv } G_1 \rightarrow_G G_3$, as well as $\text{Unrv } G_2 = \text{Unrv } G_3$.

Proof: In each step of $G_1 \rightarrow_G G_2$ a cut K is contracted. Using colouring, we can build a reduction sequence $\text{Unrv } G_1 \rightarrow_G G_3$, for some G_3 , by contracting, for each step in $G_1 \rightarrow_G G_2$, always only all copies of K (using the same rule repeatedly). Notice that this reduction might have introduced sharing, and that G_2 and G_3 differ only in that G_3 contains less sharing than G_2 , i.e. G_3 is a partially unravelled version of G_2 . Since no other manipulation has been performed, we get $\text{Unrv } G_2 = \text{Unrv } G_3$. \square

We also have the following adequacy result:

Theorem 4.6 (ADEQUACY) Let G_1, G_2 be \mathcal{X} -graphs, and P_1, P_2 be \mathcal{X} -terms such that $\text{Unrv } G_i = \llbracket P_i \rrbracket_G$, for $i = 1, 2$. If $G_1 \rightarrow_G G_2$, then $P_1 \rightarrow P_2$. Moreover, if G_2 is in normal form, then so is P_2 .

Proof: By Lemma 4.5, we get that there exists a G_3 such that $\llbracket P_1 \rrbracket_G \rightarrow_G G_3$, as well as $\llbracket P_2 \rrbracket_G = \text{Unrv } G_3$. This reduction induces, similar to Lemma 4.5, a reduction from P_1 to P_2 . If G_2 contains no cuts, then neither does $\llbracket P_2 \rrbracket_G$, nor P_2 . \square

We can now prove the following result:

Theorem 4.7 If $P \rightarrow Q$ in one step, then there exists a \mathcal{X} -graph G such that: $\llbracket P \rrbracket_G \rightarrow_G G$, and $\text{Unrv } G = \llbracket Q \rrbracket_G$.

Proof: Easy: in $\llbracket P \rrbracket_G$, redexes are not shared; the only sharing in G is introduced by the reduction, which gets erased by unravelling. \square

Notice that, by the non-confluent character for \mathcal{X} , we cannot prove a similar result for many-steps reduction paths, as illustrated by the following example.


Example 4.5 Let P and Q be such that $\alpha \notin \text{fp}(P)$ and $x \notin \text{fp}(Q)$, so $P \leftarrow P\hat{\alpha} \dagger \hat{x}Q \rightarrow Q$. Now (assume $z \neq v$):

$$\begin{aligned} (P\hat{\alpha} \dagger \hat{x}Q) \hat{\gamma} \hat{\lambda} \hat{z}(\langle z \cdot \beta \rangle \hat{\beta} [v] \hat{w}(\hat{y} \langle z \cdot \delta \rangle \hat{\delta} \cdot \epsilon)) &\rightarrow (\lambda \text{imp-ins}), (\lambda \text{exp}), (\lambda d) (2 \times) \\ ((P\hat{\alpha} \dagger \hat{x}Q) \hat{\gamma} \dagger \hat{z} \langle z \cdot \beta \rangle) \hat{\beta} [v] \hat{w}(\hat{y}((P\hat{\alpha} \dagger \hat{x}Q) \hat{\gamma} \dagger \hat{z} \langle z \cdot \delta \rangle) \hat{\delta} \cdot \epsilon) & \\ \rightarrow^* (a \neq), (2.9), (\lambda a), (2.9) & \\ (P\hat{\gamma} \dagger \hat{z} \langle z \cdot \beta \rangle) \hat{\beta} [v] \hat{w}(\hat{y}(Q\hat{\gamma} \dagger \hat{z} \langle z \cdot \delta \rangle) \hat{\delta} \cdot \epsilon) & \end{aligned}$$


Notice that we have explicitly used the non-confluence of the cut $P\hat{\alpha} \dagger \hat{x}Q$, and reduced it once to P , and once to Q .

We cannot simulate this in the setting of Term Graph Rewriting. Instead, we get the follow-

ing graph for the first term



which, by $(\lambda imp-ins)$, (λexp) , and $(\lambda d)(2\times)$ reduces to the graph



Of course, although shared, the cut $Cut(P, \alpha, x, Q)$ can be reduced only *once*, resulting in either P or Q , implying that the above result cannot be achieved.

This is not unexpected, however, since all implementations of reduction systems will use a *reduction strategy*, preferring certain redexes over others, and thereby excluding other reduction paths; unfortunately, in a non-confluent setting, this means that certain reachable normal forms are no longer achievable.

5 Dealing with α -conversion in \mathcal{X}

In the rest of this paper we will discuss a number of solutions to the problem of capture avoidance / α -conversion in the context of \mathcal{X} . Capture avoidance is a well-known implementation issue; the most familiar context in which this problem occurs is perhaps the λ -calculus, where, when reducing a term like $(\lambda xy.xy)(\lambda xy.xy)$, α -conversion is essential. Without it, one would get

$$\begin{aligned} (\lambda xy.xy)(\lambda xy.xy) &\rightarrow (\lambda y.xy)[\lambda xy.xy/x] \\ &= \lambda y.(\lambda xy.xy)y \\ &\rightarrow \lambda y.(\lambda y.xy)[y/x] \rightarrow \lambda yy.yy \end{aligned}$$

The conflict is caused by the fact that during the second reduction step, the free occurrence of y in $(\lambda xy.xy)y$ is brought *under* the binding λy , *i.e.* is *captured*.

A particular problem in dealing with α -conversion here is that the only reduction rule is $(\lambda x.M)N \rightarrow M[N/x]$, where the substitution is *implicit*, and supposed to be performed *immediately*. For example, when reducing

$$(\lambda xy.xy)(\lambda xy.xy) \rightarrow \lambda y.(\lambda xy.xy)y,$$

the latter term is *identical* to $\lambda y.xy[(\lambda xy.xy)/x]$. The actual performance of the substitution, which brings the right-most binder under the left-most is not part of the reduction system itself, but specified in the auxiliary definition of substitution.

A way to avoid this problem is, normally, to assume Barendregt's convention. In the reduction above, the first term satisfies this criterion, but the second does not: y is free as well as

bound in $(\lambda xy.xy)y$. Since this is not true for $(\lambda y.xy)[\lambda xy.xy/x]$, α -conversion must take place *during* the execution of substitution, and replace the bound name; hence, the result should be:

$$\begin{aligned} (\lambda xy.xy)(\lambda xy.xy) &\rightarrow (\lambda y.xy)[\lambda xy.xy/x] \\ &= \lambda y.(\lambda xz.xz)y \\ &\rightarrow \lambda y.(\lambda z.xz)[y/x] = \lambda yz.yz \end{aligned}$$

In the pure λ -calculus, it is impossible to deal with α -conversion since capturing cannot be expressed. The historical definition contained a notion of α -reduction in the language itself, but now normally the problem is avoided altogether by considering terms ‘modulo α -conversion’, and Barendregt’s convention is cited, stating that bound and free names should be kept distinct by, when necessary, performing a silent α -conversion, thereby in fact leaving it to the implementer to treat.

For λx the situation is slightly better, in that now substitution is an explicit part of term manipulation, so the creation of a naming conflict while reducing can be expressed. In fact, variable capturing can be avoided radically by changing the reduction rule

$$(\lambda y.M)\langle x:=P \rangle \rightarrow \lambda y.(M\langle x:=P \rangle)$$

into

$$(\lambda y.M)\langle x:=P \rangle \rightarrow \lambda z.(M\langle y:=z \rangle\langle x:=P \rangle), \quad (z \text{ fresh})$$

thereby preventing a capture on a possibly bound y in P . This is expensive though, as it is performed on *all* substitutions on abstractions, and does not actually detect the capture, but just prevents it.

If one could use side-conditions on the rules as in λx_{gc} we could express that a free occurrence in a term will be captured, and define:

$$\begin{aligned} (\lambda y.M)\langle x:=P \rangle &\rightarrow \lambda z.(M\langle y:=z \rangle\langle x:=P \rangle), & (z \text{ fresh \& } y \in fs(P)) \\ (\lambda y.M)\langle x:=P \rangle &\rightarrow \lambda z.(M\langle x:=P \rangle), & (y \notin fs(P)) \end{aligned}$$

This observation lies at the basis of how we deal with α -conversion in Section 5.3.

Because it is an implementation issue, and we aim to have an efficient implementation of \mathcal{X} , we deal with the problem more formally. First we show that it is an issue to begin with.

Example 5.1 Take the term⁴

$$(\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \gamma) \hat{\gamma} \dagger \hat{x}(\langle x \cdot \delta \rangle \hat{\delta} [x] \hat{w}\langle w \cdot \alpha \rangle)$$

which is also used in Example 4.3. Reducing this term differently, we get:

$$\begin{aligned} (\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \gamma) \hat{\gamma} \dagger \hat{x}(\langle x \cdot \delta \rangle \hat{\delta} [x] \hat{w}\langle w \cdot \alpha \rangle) &\rightarrow^* (\hat{x}a), (\hat{x} \text{ imp-outs}) \\ (\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \gamma) \hat{\gamma} \dagger \hat{z}(((\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \gamma) \hat{\gamma} \hat{x}\langle y \cdot \mu \rangle) \hat{\delta} [z] \hat{w}((\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \gamma) \hat{\gamma} \hat{x}\langle w \cdot \alpha \rangle)) &\rightarrow^* (\hat{x}d), (\text{exp}), (\hat{x} \text{ cap}) \\ (\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \gamma) \hat{\gamma} \dagger \hat{z}((\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \delta) \hat{\delta} [z] \hat{w}\langle w \cdot \alpha \rangle) &\rightarrow (\text{exp-imp}) \\ (\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \delta) \hat{\delta} \dagger \hat{y}(\langle y \cdot \mu \rangle \hat{\mu} \dagger \hat{w}\langle w \cdot \alpha \rangle) &\rightarrow^* (\hat{x}a), (\hat{x} \text{ cut}), (\hat{x}d), (\hat{x} \text{ cap}) \\ ((\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \delta) \hat{\delta} \dagger \hat{y}\langle y \cdot \mu \rangle) \hat{\mu} \dagger \hat{w}\langle w \cdot \alpha \rangle &\rightarrow (\text{exp}) \\ (\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \mu) \hat{\mu} \dagger \hat{w}\langle w \cdot \alpha \rangle &\end{aligned}$$

⁴it is easy to check this term results from reducing the λ -term $xx(x:=\lambda y.y)$; reducing this λ -term poses no α -conversion problem.

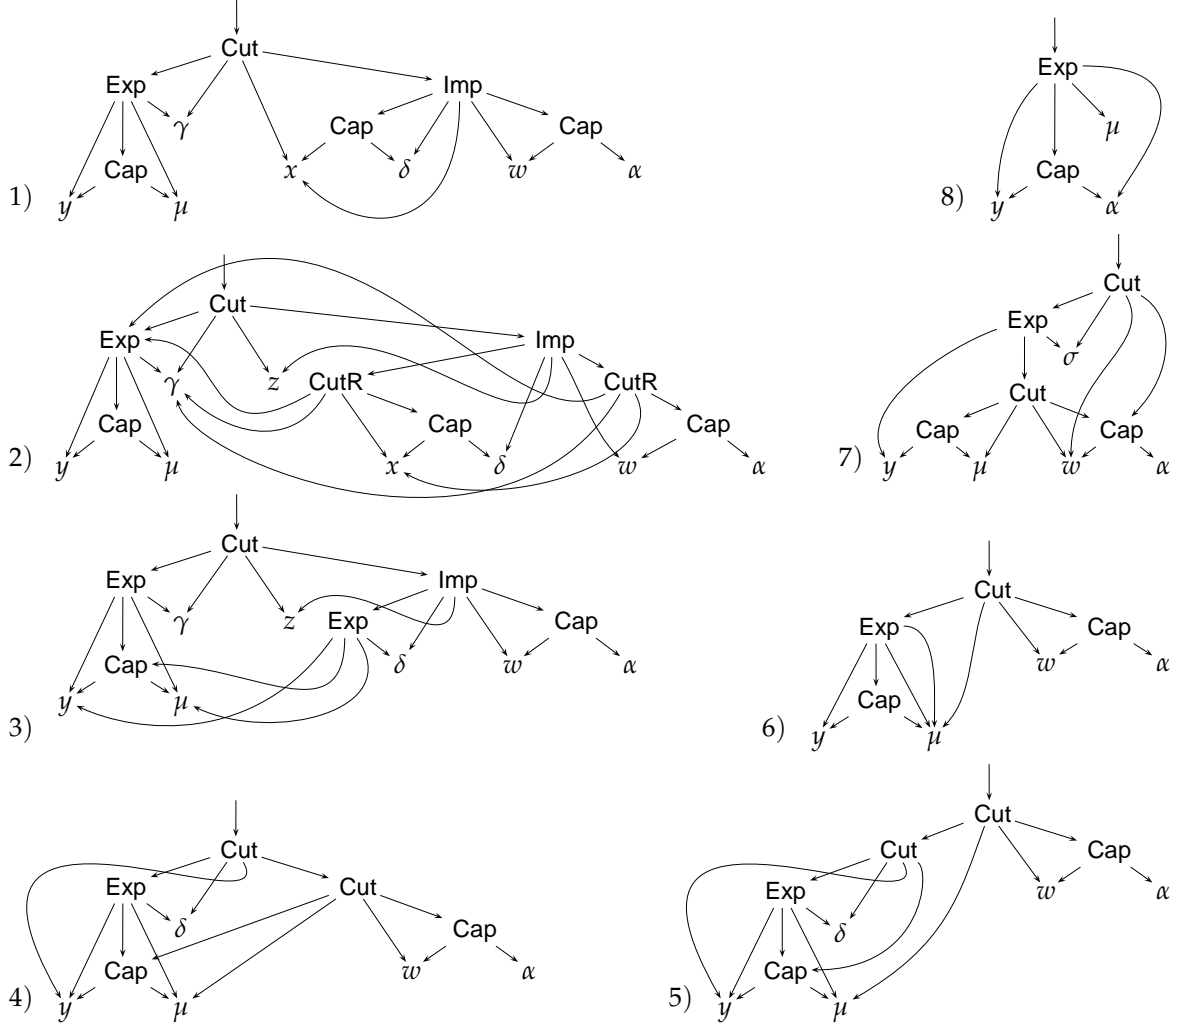


Figure 4: Graphs for Example 5.1

Notice that, in this term, μ is bound *twice*. To reduce the cut, we first check if μ is introduced; this is not so, since $\mu \in fp(\langle y \cdot \mu \rangle)$, thus denying the intended application of the rule (*exp*). Instead, we propagate the cut and obtain, using rules (*a'*), (*exp-outs'*) and (*d'*),

$$\begin{aligned}
 (\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \mu) \hat{\mu} \dagger \hat{w}\langle w \cdot \alpha \rangle &\rightarrow^* (a'), (exp-outs'), (d') \\
 (\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \dagger \hat{w}\langle w \cdot \alpha \rangle) \hat{\mu} \cdot \sigma \hat{\sigma} \dagger \hat{w}\langle w \cdot \alpha \rangle &\rightarrow^* (cap), (exp) \\
 \hat{y}\langle y \cdot \alpha \rangle \hat{\mu} \cdot \alpha &
 \end{aligned}$$


which is not the intended $\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \alpha$. The problem is of course that we should have renamed one of the bound μ : this is an α -conversion problem, which is not solved by merely running the graph interpretation, as can be seen in Figure 4, where we show the corresponding graph reduction.

5.1 Lazy copying of shared graphs

The solution to the problem of capture we propose in this section is the one we first presented in [4]; it aims to avoid, as for λ -graphs, the sharing of graphs that are involved in more than one cut. Similarly to the case for the λ -calculus [29, 15], binding of connectors can be considered problematic in the context of sharing. Sharing an abstraction $\lambda x.G$ in λ -graphs is problematic

since the substitution is implemented via a redirection on G . This can be done only once, blocking a re-use of a shared abstraction, that therefore has to be copied first. To tackle this problem within the context of \mathcal{X} , a notion of *rebinding of sockets* and of *plugs* was introduced.

The basic idea is the following: suppose we are dealing with the graph



The fact that α is bound *twice* might cause the binders to become *nested* during computation. We avoid that by copying that part P that depends on α : we will ‘peel off a copy’ of the graphs which might get affected by the double binding of connectors.

This method is similar to that of [29], and differs in that we must copy several constructors and considering two classes of variables. An eager copy method will be too restrictive, in that it destroys a large amount of sharing. We will specify a lazy strategy that avoids much of the unnecessary copying in Section 6.3.

We extend the syntax of \mathcal{X} -graphs with two further constructors

$$P ::= \dots \mid \text{rp}(P, \alpha, \beta) \mid \text{rs}(P, x, y)$$

that will represent the renaming of a bound (rebinding) plug or socket, respectively. This results in the (term graph) definition of rebinding a socket (rp) as given in Definition 5.1. These prevent a connector from being doubly bound by, essentially, copying that structure of a graph which contains that binder whilst introducing the new connector, thereby destroying the sharing of the connector via binding edges.

The function $\text{rp}(P, \alpha, \beta)$ as given in Definition 5.1 is defined to build a new graph G' where the free occurrences of α in G are replaced with β and any binders encountered *in* G are made fresh. Since this is, essentially, a copying function, when we move the rebinding mechanism *under* binders, as in the third case below, we would create double binders for those bound connectors we have just passed. Therefore, we need to rebind those as well.

Definition 5.1 (REBINDING REWRITE RULES) The function rp is defined by the term graph rewriting rules in Figure 5.

$$\begin{aligned} (\text{rpGC}) \quad & \text{rp}(P, \beta, \gamma) \rightarrow P, \quad \beta \text{ not free in } P \\ (\text{rpCapRen}) \quad & \text{rp}(\langle x \cdot \beta \rangle, \beta, \gamma) \rightarrow \langle x \cdot \gamma \rangle \\ (\text{rpExp}) \quad & \text{rp}(\widehat{y}P\widehat{\alpha} \cdot \eta, \beta, \gamma) \rightarrow \widehat{k} \text{rs}(\text{rp}(\text{rp}(P, \beta, \gamma), \alpha, \lambda), y, k) \widehat{\lambda} \cdot \eta, \quad (\eta \neq \beta) \\ (\text{rpExpRen}) \quad & \text{rp}(\widehat{y}P\widehat{\alpha} \cdot \beta, \beta, \gamma) \rightarrow \widehat{k} \text{rs}(\text{rp}(\text{rp}(P, \beta, \gamma), \alpha, \lambda), y, k) \widehat{\lambda} \cdot \gamma \\ (\text{rpMed}) \quad & \text{rp}(P\widehat{\alpha}[x]\widehat{y}Q, \beta, \gamma) \rightarrow \text{rp}(\text{rp}(P, \beta, \gamma), \alpha, \eta) \widehat{\eta}[x]\widehat{z} \text{rs}(\text{rp}(Q, \beta, \gamma), y, z) \\ (\text{rpCut}) \quad & \text{rp}(P\widehat{\alpha} \dagger \widehat{y}Q, \beta, \gamma) \rightarrow \text{rp}(\text{rp}(P, \beta, \gamma), \alpha, \eta) \widehat{\eta} \dagger \widehat{z} \text{rs}(\text{rp}(Q, \beta, \gamma), y, z) \end{aligned}$$

(The function rs is defined similarly.) Notice that the call to the function rp builds a version of P that uses a fresh socket γ to connect rather than β . Also, all bound connectors are renamed: evaluating the rebinding rules builds a version of P with fresh binder names. This ensures there is only ever one pointer to nodes that bind over P or the local binders in P . This comes, however, at the price of a large rebinding overhead.

The functions rs and rp are expressed as higher-order term-graph rewriting rules. Because these higher-order functions may not necessarily be evaluated eagerly, they may interfere

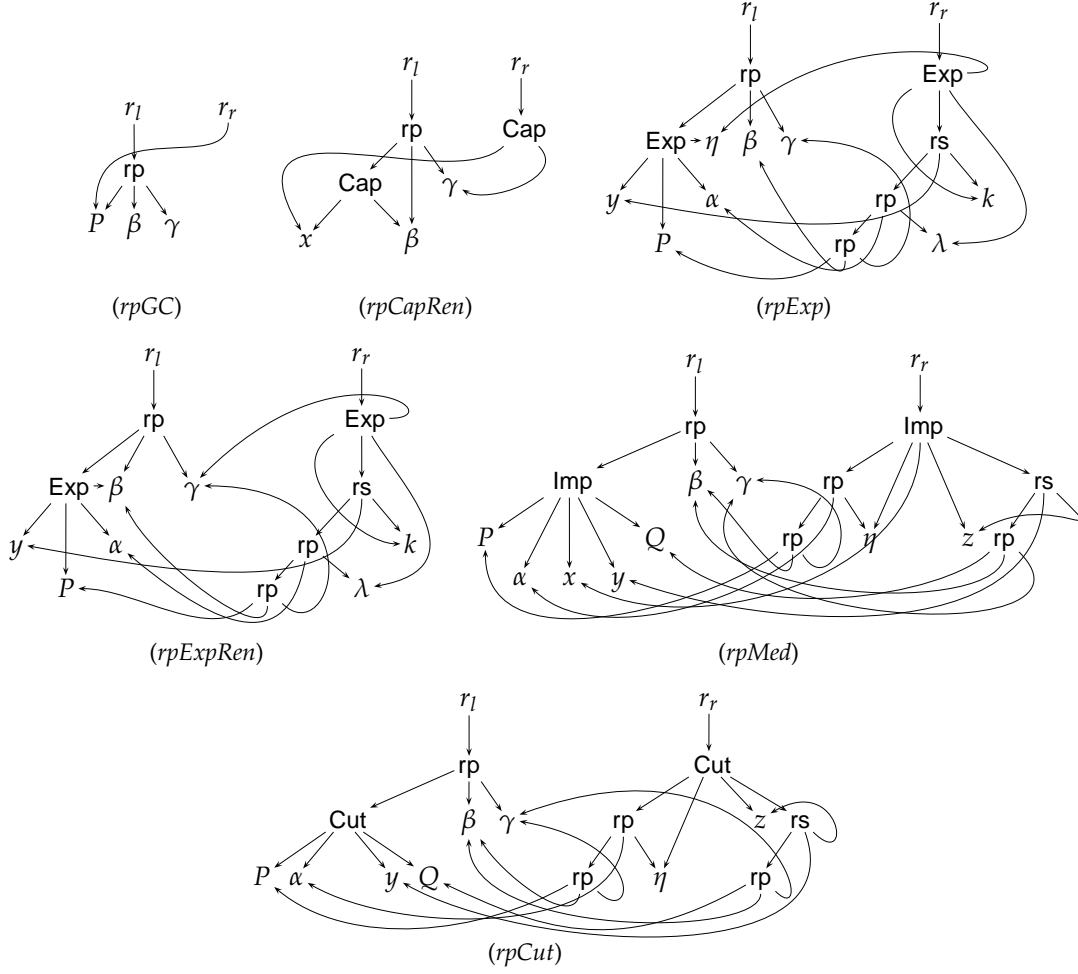


Figure 5: Rebinding plugs


with the reductions of the λ -calculus: if the sub-circuit of an inactive cut is a rebinding term, an activation will be forced even though the sub-circuit of the rebinding term introduces the appropriate connector of the cut (and a logical rule should therefore have been applied).

Rather than forcing the evaluation of these rebinding constructs to completion via an ‘eager’ reduction strategy, we will define a lazier evaluation strategy that avoids this mis-activation in Section 6.3.

Using the functions rs and rp gives a different formal definition of interpreting rewrite rules in \mathcal{X} as graphs. As suggested by the example above (Example 5.1), term rewrite rules which introduce sharing of binders need to copy these in order to avoid capture.

Definition 5.2 (COPYING TGRS REWRITE RULES) Left propagation

$$(exp-outs^{\nearrow}) : (\hat{y}P\hat{\beta}\cdot\alpha)\hat{\alpha}\nearrow\hat{x}Q \rightarrow (\hat{y}(P\hat{\alpha}\nearrow\hat{x}Q)\hat{\beta}\cdot\gamma)\hat{\gamma}\dagger\hat{z}rs(Q,x,z)$$



$$\begin{aligned}
& (\hat{y}\langle y\cdot\mu\rangle\hat{\mu}\cdot\gamma)\hat{\gamma}\dagger\hat{x}\langle(x\cdot\delta)\hat{\delta}[x]\hat{w}\langle w\cdot\alpha\rangle\rangle \quad \rightarrow^* (\lambda a), (\lambda \text{imp-outs}) \\
& (\hat{y}\langle y\cdot\mu\rangle\hat{\mu}\cdot\gamma)\hat{\gamma}\dagger\hat{z}((\text{rp}((\hat{y}\langle y\cdot\mu\rangle\hat{\mu}\cdot\gamma),\gamma,\tau)\hat{\tau}\lambda\hat{x}\langle x\cdot\delta\rangle))\hat{\beta}[z]\hat{y}(\text{rp}(\hat{y}\langle y\cdot\mu\rangle\hat{\mu}\cdot\gamma,\gamma,\sigma)\hat{\sigma}\lambda\hat{x}\langle w\cdot\alpha\rangle)) \quad \rightarrow^* (\text{rp}) \\
& (\hat{y}\langle y\cdot\mu\rangle\hat{\mu}\cdot\gamma)\hat{\gamma}\dagger\hat{z}((\hat{v}\langle v\cdot\nu\rangle\hat{v}\cdot\tau)\hat{\tau}\lambda\hat{x}\langle x\cdot\delta\rangle)\hat{\delta}[z]\hat{w}((\hat{u}\langle u\cdot\eta\rangle\hat{\eta}\cdot\sigma)\hat{\sigma}\lambda\hat{x}\langle w\cdot\alpha\rangle)) \quad \rightarrow^* (\lambda d), (\text{exp}), (\lambda \text{cap}) \\
& (\hat{y}\langle y\cdot\mu\rangle\hat{\mu}\cdot\gamma)\hat{\gamma}\dagger\hat{z}((\hat{v}\langle v\cdot\nu\rangle\hat{v}\cdot\delta)\hat{\delta}[z]\hat{w}\langle w\cdot\alpha\rangle) \quad \rightarrow (\text{exp-imp}) \\
& (\hat{v}\langle v\cdot\nu\rangle\hat{v}\cdot\delta)\hat{\delta}\dagger\hat{y}\langle y\cdot\mu\rangle\hat{\mu}\dagger\hat{w}\langle w\cdot\alpha\rangle \quad \rightarrow^* (\lambda a), (\lambda \text{cut}), (\lambda \text{exp}), (\lambda d) \\
& ((\hat{v}\langle v\cdot\nu\rangle\hat{v}\cdot\delta)\hat{\delta}\dagger\hat{y}\langle y\cdot\mu\rangle)\hat{\mu}\dagger\hat{w}\langle w\cdot\alpha\rangle
\end{aligned}$$


Figure 6: Correction using rebinding of Example 5.1.

$$(\text{imp}\lambda): (Q\hat{\beta}[z]\hat{y}R)\hat{\alpha}\lambda\hat{x}P \rightarrow (Q\hat{\alpha}\lambda\hat{x}P)\hat{\beta}[z]\hat{y}(R\hat{\alpha}\lambda\hat{k}\text{rs}(P,x,k))$$

$$(\text{cut}\lambda): (Q\hat{\beta}\dagger\hat{y}R)\hat{\alpha}\lambda\hat{x}P \rightarrow (Q\hat{\alpha}\lambda\hat{x}P)\hat{\beta}\dagger\hat{y}(R\hat{\alpha}\lambda\hat{k}\text{rs}(P,x,k))$$

Right propagation

$$\begin{aligned}
(\lambda \text{imp-outs}): P\hat{\alpha}\lambda\hat{x}(Q\hat{\beta}[x]\hat{y}R) & \rightarrow \\
P\hat{\alpha}\dagger\hat{z}((\text{rp}(P,\alpha,\mu)\hat{\mu}\lambda\hat{x}Q)\hat{\beta}[z]\hat{y}(\text{rp}(P,\alpha,\eta)\hat{\eta}\lambda\hat{x}R)) &
\end{aligned}$$



$$(\lambda \text{imp-ins}): P\hat{\alpha}\lambda\hat{x}(Q\hat{\beta}[z]\hat{y}R) \rightarrow (P\hat{\alpha}\lambda\hat{x}Q)\hat{\beta}[z]\hat{y}(\text{rp}(P,\alpha,\gamma)\hat{\gamma}\lambda\hat{x}R)$$

$$(\lambda \text{cut}): P\hat{\alpha}\lambda\hat{x}(Q\hat{\beta}\dagger\hat{y}R) \rightarrow (P\hat{\alpha}\lambda\hat{x}Q)\hat{\beta}\dagger\hat{y}(\text{rp}(P,\alpha,\gamma)\hat{\gamma}\lambda\hat{x}R)$$

That this (partial) copy action gives a solution to the name capture problem is perhaps not obvious. In fact, running the modified system on terms that have a bound connector occur more than once would not be guaranteed to be safe. However, if the initial term has different names for all bound connectors, which in turn are different from all free names, then this property is preserved during reduction. So it is impossible for names to be captured.

Example 5.2 A corrected reduction, using rebinding, for Example 5.1 can be found in Figure 6. Notice that this time there is no possibility of variable clash, since there are no shared binders (the other copy of the μ binder together with the μ in the capsule is renamed to ν). The highlighted cut $\hat{\mu}\dagger\hat{w}$ can be activated and safely propagated through the left sub-circuit.

The solution using rebinding is surprisingly easy to formulate, and only the rules that use explicit replication need to be changed, but comes at the price of having to extend the signature of the calculus, as well as the set of rewrite rules. Moreover, it turns out to be highly inefficient; this is of course mainly due to the loss of sharing. The main objection to rebinding is that it creates unnecessary overhead in that it invokes rebinding for all double bindings of connectors, regardless of whether or not they created a conflict; as we will see in the final section, the cost of rebinding is high.

5.2 Preserving Barendregt's convention

Although the solution of the previous section is correct in that it avoids the creation of nested bindings, the implementation was slow, mainly due to the large amount of extra reduction that were introduced as a result of the copying process, even when done lazily.

More importantly, we observed that, using term graph rewriting techniques, the particular λ -graph problem disappears; substitution on a graph G is now specified *explicitly* by reduction rules, and application of these rules will have the effect that a term is generated *of the same shape as G* , to which the substitution has been applied. In fact, a large part of the necessary copying action is done during the *build* phase of the rewrite mechanism, and does not need to be treated at the level of the reduction rules. The conclusion of this was that, although the copying solved the α -conversion problem, here the problem is of a different nature, i.e. that of *capturing*.

In this section, we will propose a solution for the α -conversion problem in \mathcal{X} , by preserving Barendregt's convention on names, i.e. make sure that names never occur both free and bound. We will achieve this by detecting and avoiding name clashes, without having to extend the syntax of the calculus; we just modify the (side-conditions of) the rules.

To tackle it in a formal way, we first introduce the notion of α -safety.

Definition 5.3 (α -SAFETY) We call a *term* (\mathcal{X} -graph) α -safe if it adheres to Barendregt's convention, i.e.: (1) no connector occurs both free and bound, and (2) no nesting of binders to the same connector occurs. We call a *rewrite rule* α -safe if it respects α -safety, that is, it rewrites an α -safe term (graph) to an α -safe term (graph). We call a rewrite system α -safe if all its rules are α -safe.

Example 5.4 The term $(\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \mu) \hat{\mu} \dagger \hat{w}\langle w \cdot \alpha \rangle$ is not α -safe (it fails both criteria); neither is $(\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \delta) \hat{\mu} \dagger \hat{w}\langle w \cdot \alpha \rangle$, by the first criterion.

In order to obtain an α -safe implementation of \mathcal{X} , we need to identify the rewrite rules that are *not* α -safe.

Example 5.5 In Example 5.1, the application of (*exp-imp*) in Step 3 violates our α -safety property since μ is both bound and free in $(\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \delta) \hat{\delta} \dagger \hat{y}\langle y \cdot \mu \rangle$. So rule (*exp-imp*) is *not* α -safe; for the left-hand side $(\hat{y}P \hat{\beta} \cdot \alpha) \hat{\alpha} \dagger \hat{x}(Q \hat{\gamma}[x] \hat{z}R)$ to be α -safe, the connectors y and z , and β and γ are allowed to be the same: they are not nested. However, this is no longer true for the right-hand sides: y and z occur nested in the first alternative $Q \hat{\gamma} \dagger \hat{y}(P \hat{\beta} \dagger \hat{z}R)$, and β and γ in the second $(Q \hat{\gamma} \dagger \hat{y}P) \hat{\beta} \dagger \hat{z}R$, which might force an α -conversion to be necessary.

Since we do not necessarily need to avoid connectors being bound twice (as long as they are not nested), in dealing with the necessary renaming of bound connectors we can take advantage of the explicit renaming feature of \mathcal{X} , using new cuts such as $\langle v \cdot \delta \rangle \hat{\delta} \chi \hat{y}P$ or $P \hat{\beta} \chi \hat{v} \langle v \cdot \delta \rangle$ to rename y by v , or β by δ respectively in P , where v, δ are fresh (see Lemma 2.9). By activating the cuts, the renaming is forced to take place and other cuts are prohibited from propagating over the renaming.

On the down-side, unlike for the rebinding solution where we added extra nodes to our graphs to express the copying, it is no longer possible to force eager or lazy evaluation of α -conversion without doing the same with the general propagation rules.

We discussed above in Example 5.5 that we should perform the α -conversion in rule (*exp-imp*). Let us consider the first choice (assume α and x are introduced); remember that y might be z or appear bound in P without violation of Barendregt's convention:

$$(\hat{y}R \hat{\beta} \cdot \alpha) \hat{\alpha} \dagger \hat{x}(Q \hat{\gamma}[x] \hat{z}P) \rightarrow Q \hat{\gamma} \dagger \hat{y}(R \hat{\beta} \dagger \hat{z}P)$$

If we assume that the α -safety criterion holds on an instance of the left-hand side of the rule (so $\gamma \notin bp(\hat{y}P\hat{\beta}\cdot\alpha)$ and $y \notin bs(Q\hat{\gamma}[x]\hat{z}R)$), then the application of the rule should not break the α -safety criterion. We require the right-hand side of the rule to also be α -safe. Notice that in its current form, the rule does indeed break the criterion, since β and γ are nested on the right-hand side and the left-hand side places no constraints on the relation between β and γ . A term graph on which this rule is applied may have $\beta=\gamma$, in which case the application of the rule will have created a nested binding.

Since now we do not necessarily need to avoid connectors being bound twice (as long as they are not nested), we do not need to completely copy terms. Instead, in dealing with the necessary renaming of bound connectors we can take advantage of the explicit renaming feature of \mathcal{X} , introducing to the rules new cuts such as $\langle v \cdot \delta \rangle \hat{\delta} \backslash \hat{y}P$ or $P\hat{\beta} \nearrow \hat{v} \langle v \cdot \delta \rangle$ to rename y by v , or β by δ respectively in P , where v, δ are fresh (see Lemma 2.9). By activating the cuts, the intention is to force the renaming to take place first. We will also need to adopt our proposed strongly normalising rules, which prevent a cut from deactivating, thereby enforcing priority to the renaming cuts.

Returning to the violation in the ($exp\text{-}imp_N$) rule, in order to ensure the rewrite will be executed correctly (with respect to α -safety), we need to introduce an extra constraint to the applicability of the rule ($exp\text{-}imp_N$), namely $\gamma \notin bs(\hat{y}P\hat{\beta}\cdot\alpha)$. (This can be equivalently formulated as $\gamma \notin bs(P) \wedge \beta \neq \gamma$.) If the side-condition does not hold, then applying the rule will create a nested binding of (the image of) γ in the term graph. To remedy the situation, we must pay the cost of a renaming. This implies that there are now two alternatives for this rule:

$$(\hat{y}R\hat{\beta}\cdot\alpha)\hat{\alpha} \dagger \hat{x}(Q\hat{\gamma}[x]\hat{z}P) \rightarrow \begin{cases} Q\hat{\gamma} \dagger \hat{y}(R\hat{\beta} \dagger \hat{z}P) & (y \neq z, y \notin bs(P)) \\ Q\hat{\gamma} \dagger \hat{v}(\langle v \cdot \delta \rangle \hat{\delta} \backslash \hat{y}R)\hat{\beta} \dagger \hat{z}P & (y = z \vee y \in bs(P), v, \delta \text{ fresh}) \end{cases}$$

Likewise, there are two alternatives for the second choice:

$$(\hat{y}R\hat{\beta}\cdot\alpha)\hat{\alpha} \dagger \hat{x}(Q\hat{\gamma}[x]\hat{z}P) \rightarrow \begin{cases} (Q\hat{\gamma} \dagger \hat{y}R)\hat{\beta} \dagger \hat{z}P & (\gamma \neq \beta \notin bp(Q)) \\ (Q\hat{\gamma} \dagger \hat{y}(R\hat{\beta} \nearrow \hat{v} \langle v \cdot \delta \rangle))\hat{\delta} \dagger \hat{z}P & (\beta = \gamma \vee \beta \in bp(Q), v, \delta \text{ fresh}) \end{cases}$$

Example 5.6 Applying this solution to Example 5.1, we have, instead of the problematic step

$$\begin{aligned} & (\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \gamma) \hat{\gamma} \dagger \hat{k}((\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \delta) \hat{\delta} [k] \hat{w}\langle w \cdot \alpha \rangle) \rightarrow (exp\text{-}imp) \\ & ((\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \delta) \hat{\delta} \dagger \hat{y}\langle y \cdot \mu \rangle) \hat{\mu} \dagger \hat{w}\langle w \cdot \alpha \rangle \end{aligned}$$

the correction

$$\begin{aligned} & (\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \gamma) \hat{\gamma} \dagger \hat{k}((\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \delta) \hat{\delta} [k] \hat{w}\langle w \cdot \alpha \rangle) \rightarrow (exp\text{-}imp) \\ & ((\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \delta) \hat{\delta} \dagger \hat{y}(\langle y \cdot \mu \rangle \hat{\mu} \nearrow \hat{v} \langle v \cdot \beta \rangle)) \hat{\beta} \dagger \hat{w}\langle w \cdot \alpha \rangle \rightarrow (d \nearrow, \cap) \\ & ((\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \delta) \hat{\delta} \dagger \hat{y}\langle y \cdot \beta \rangle) \hat{\beta} \dagger \hat{w}\langle w \cdot \alpha \rangle \rightarrow (exp) \\ & (\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \beta) \hat{\beta} \dagger \hat{w}\langle w \cdot \alpha \rangle \rightarrow (exp) \\ & \hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \alpha \end{aligned}$$

To guarantee α -safety, we need to make a similar change to each rule where a possible α -conflict is introduced, like ($\backslash cut$):

$$P\hat{\alpha} \backslash \hat{x}(Q\hat{\beta} \dagger \hat{y}R) \rightarrow (P\hat{\alpha} \backslash \hat{x}Q)\hat{\beta} \dagger \hat{y}(P\hat{\alpha} \backslash \hat{x}R)$$

There are two points of concern here: $\alpha = \beta$, and β or y occurs bound in P (notice that $x \neq y$ as, by assumption, the left-hand side is an α -safe term). With this in mind, the rule ($\backslash cut$)

is amended with extra side conditions and replaced by the following variants (where v, δ are fresh):

$$P\hat{\alpha} \times \hat{x}(Q\hat{\beta} \dagger \hat{y}R) \rightarrow \begin{cases} (P\hat{\alpha} \times \hat{x}Q)\hat{\beta} \dagger \hat{y}(P\hat{\alpha} \times \hat{x}R) & (\beta \notin bp(P) \ \& \ y \notin bs(P)) \\ (P\hat{\alpha} \times \hat{x}Q)\hat{\beta} \dagger \hat{v}(P\hat{\alpha} \times \hat{x}(\langle v \cdot \delta \rangle \hat{\delta} \times \hat{y}R)) & (\beta \notin bp(P) \ \& \ y \in bs(P)) \\ (P\hat{\alpha} \times \hat{x}(Q\hat{\beta} \dagger \hat{v}\langle v \cdot \delta \rangle))\hat{\delta} \dagger \hat{y}(P\hat{\alpha} \times \hat{x}R) & (\beta \in bp(P) \ \& \ y \notin bs(P)) \\ (P\hat{\alpha} \times \hat{x}(Q\hat{\beta} \dagger \hat{v}\langle v \cdot \delta \rangle))\hat{\delta} \dagger \hat{v}(P\hat{\alpha} \times \hat{x}(\langle v \cdot \delta \rangle \hat{\delta} \times \hat{y}R)) & (\beta \in bp(P) \ \& \ y \in bs(P)) \end{cases}$$

Almost all propagation rules (except for $(d \dagger)$, $(cap \dagger)$, $(\times d)$, and $(\times cap)$) should be treated like this; of the logical rules, only rule $(exp-imp)$ needs dealing with as specified above. In all, this gives a much more complicated rewriting system, with a great many rewrite rules. The advantage of this approach is that α -conversion itself is detected and dealt with:

Theorem 5.7 *Let \rightarrow_α stand for the notion of rewriting on \mathcal{X} obtained by changing the rules as above. Then: if P, Q are pure terms such that P is α -safe, and $P \rightarrow_\alpha Q$, then Q is α -safe.*

Proof: Straightforward, by verifying that each individual rewriting step preserves α -safeness. \square

The computational cost of this approach is low compared to the rebinding approach; the price to pay is an increase in the number of rules. However, the detection of a possible ' α -danger' in a rule is straightforward, since it only depends on the information present in the rule, so it is even possible to, at the user level, allow for the definition of the normal rules, and to automatically generate the α -safe variants.

5.3 Avoiding capture

Barendregt's convention is a perfectly adequate solution to the α -conversion problem, since it forbids a term with nested binders to the same variable to be created, thereby totally avoiding any ambiguity to the system. However, one can justifiably argue that the convention is restrictive, and expensive to uphold at run-time. After all, allowing nesting of bound variables as in $\lambda y.(\lambda xy.xy)y$ (and therefore also of variables occurring free and bound variables as in $(\lambda xy.xy)y$) need not be ambiguous at all when considering the *innermost* of nested binders the strongest; in this paradigm, the only thing that needs to be avoided during reduction is that of *capture* of free connectors, bringing a connector under a binder.

Example 5.8 Referring back to Example 5.1, recognising that μ is bound twice, we would like to apply rule (exp) to the term

$$(\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \mu) \hat{\mu} \dagger \hat{w}\langle w \cdot \alpha \rangle$$

But since that rule checks if μ occurs free in $\langle y \cdot \mu \rangle$, which it does, we are forced to left-activate the term, which eventually gives the wrong result. However, once we accept that we only need to check that a connector does not get captured - notice that this is not the case here - and that we allow connectors to appear both free and bound, we only need to check that those free occurrences of α are indeed those that we would associate with the free α on the outside of the export term. So, instead of applying rule (exp) ,

$$(\hat{x}P\hat{\beta} \cdot \alpha) \hat{\alpha} \dagger \hat{y}\langle y \cdot \gamma \rangle \rightarrow \hat{x}P\hat{\beta} \cdot \gamma \quad (\alpha \notin fp(P))$$

we should apply the variant

$$(\hat{x}P\hat{\beta} \cdot \alpha) \hat{\alpha} \dagger \hat{y}\langle y \cdot \gamma \rangle \rightarrow \hat{x}P\hat{\alpha} \cdot \gamma \quad (\alpha = \beta \vee \alpha \notin fp(P))$$

Then the reduction of Example 5.1 should contain:

$$(\hat{y}\langle y \cdot \mu \rangle \hat{\mu} \cdot \mu) \hat{\mu} \dagger \hat{w} \langle w \cdot \alpha \rangle \rightarrow \hat{y} \langle y \cdot \mu \rangle \hat{\mu} \cdot \alpha$$

We will show that we can identify the situations in which the application of a rule results in capturing, and will modify the rules in such a way that the necessary α -conversions are automatically preformed. The solution is, in appearance, strikingly similar to that of Section 5.2 but for the fact that *freeness* is used rather than *boundness*. In Section 6.6, this approach will be shown to be much more efficient; this is mainly because the solution of Section 5.2, many terms which are not ' α -safe' (Definition 5.3) are left untouched here.

Example 5.9 Let us consider the rule ($exp\text{-}imp_N$):

$$(\hat{y}R\hat{\beta} \cdot \alpha) \hat{\alpha} \dagger \hat{x}(Q\hat{\gamma}[x]\hat{z}P) \rightarrow (Q\hat{\gamma} \dagger \hat{y}R)\hat{\beta} \dagger \hat{z}P\alpha, \quad (x \text{ introduced})$$

In order to allow the rewrite to be executed like this, the side condition should express an extra criterion to avoid the capture of a free β in Q ; if $\beta \in fs(Q)$, then the rule would bring that β under the binder \hat{y} on the right-hand side, and renaming should take place. Also, notice that if $\beta = \gamma$, there would be no capture, since the order of nested binders are preserved. This implies that there are now two alternatives for the rule ($exp\text{-}imp_N$). We define these respectively as:

$$(\hat{y}R\hat{\beta} \cdot \alpha) \hat{\alpha} \dagger \hat{x}(Q\hat{\gamma}[x]\hat{z}P) \rightarrow \begin{cases} (Q\hat{\gamma} \dagger \hat{y}R)\hat{\beta} \dagger \hat{z}P & (\beta \notin fp(Q) \vee \beta = \gamma) \\ (Q\hat{\gamma} \dagger \hat{y}(R\hat{\beta} \not\wedge \hat{v} \cdot \delta)) \hat{\delta} \dagger \hat{z}P & (\beta \in fp(Q) \wedge \beta \neq \gamma) \end{cases}$$

where v, δ are fresh. Likewise, the respective rules for the CBV variant are:

$$(\hat{y}R\hat{\beta} \cdot \alpha) \hat{\alpha} \dagger \hat{x}(Q\hat{\gamma}[x]\hat{z}P) \rightarrow \begin{cases} Q\hat{\gamma} \dagger \hat{y}(R\hat{\beta} \dagger \hat{z}P) & (y \notin fs(P) \vee y = z) \\ Q\hat{\gamma} \dagger \hat{v}((\langle v \cdot \delta \rangle \hat{\delta} \not\wedge \hat{y}R)\hat{\beta} \dagger \hat{z}P) & (y \in fs(P) \wedge y \neq z) \end{cases}$$

where v, δ are fresh.

Also, since now we explicitly allow for connectors to occur both free and bound in a term, the rules need to check if the connector we try to connect to in a cut is actually really free.

For example, rule ($exp\text{-}outs \not\wedge$) now becomes:

$$(\hat{y}Q\hat{\beta} \cdot \alpha) \hat{\alpha} \not\wedge \hat{x}P \rightarrow \begin{cases} (\hat{y}(Q\hat{\alpha} \not\wedge \hat{x}P)\hat{\beta} \cdot \gamma) \hat{\gamma} \dagger \hat{x}P, & (\alpha \in fp(Q) \ \& \ \alpha \neq \beta) \\ (\hat{y}Q\hat{\beta} \cdot \alpha) \hat{\alpha} \dagger \hat{x}P, & (\alpha \notin fp(Q) \vee \alpha = \beta) \end{cases}$$

Rule ($\not\wedge cut$) becomes:

$$P\hat{\alpha} \not\wedge \hat{x}(Q\hat{\beta} \dagger \hat{y}R) \rightarrow \begin{cases} (P\hat{\alpha} \not\wedge \hat{x}Q)\hat{\beta} \dagger \hat{y}(P\hat{\alpha} \not\wedge \hat{x}R) & (y \notin fs(P) \wedge y \neq x \wedge \beta \notin fp(P) \wedge \beta \neq \alpha) \\ (P\hat{\alpha} \not\wedge \hat{x}Q)\hat{\beta} \dagger \hat{v}(P\hat{\alpha} \not\wedge \hat{x}(\langle v \cdot \delta \rangle \hat{\delta} \not\wedge \hat{y}R)) & (y \notin fs(P) \wedge y \neq x \wedge (\beta \in fp(P) \vee \beta = \alpha)) \\ (P\hat{\alpha} \not\wedge \hat{x}(Q\hat{\beta} \not\wedge \hat{v} \cdot \delta)) \hat{\delta} \dagger \hat{y}(P\hat{\alpha} \not\wedge \hat{x}R) & ((y \in fs(P) \vee y = x) \wedge \beta \notin fp(P) \wedge \beta \neq \alpha) \\ (P\hat{\alpha} \not\wedge \hat{x}(Q\hat{\beta} \not\wedge \hat{v} \cdot \delta)) \hat{\delta} \dagger \hat{v}(P\hat{\alpha} \not\wedge \hat{x}(\langle v \cdot \delta \rangle \hat{\delta} \not\wedge \hat{y}R)) & (y \in fs(P) \vee y = x \vee \beta \in fp(P) \vee \beta = \alpha) \end{cases}$$

All the rules need to be modified to check for possible capture of connectors. Although the structure of these new rules is similar to those in Section 5.2, the improvement in execution speed is impressive, as can be seen in the last section.

6 Reduction Strategies for \mathcal{X}

In Section 5, we studied different schemes for avoiding capture in the context of the \mathcal{X} -calculus. In this section, we compare the cost of each scheme when used to reduce expressions, and obtain a cost of the renaming operations required to uphold a particular α -safety criteria; e.g. in the ‘preserving Barendregt’s convention’ scheme, the cost refers to the rewrite operations that must be carried out to ensure no nested binding is created during a reduction.

To allow for a fair and accurate comparison across the schemes we proposed, it is important that, aside from the α -conversion steps, the same reduction paths are chosen when evaluating a term. This implies the need for a *deterministic* reduction strategy. Furthermore, this strategy should not be affected by any renaming cuts performing α -conversions. The following example shows how a naive reduction strategy (in the call-by-name subsystem) could be affected, and motivates the need for an extension to the \mathcal{X} implementation to allow for complex strategies to be defined.

Example 6.1 Consider an instance of a graph where the rule (insRrenP) from the capture avoidance scheme (Section 5.3) is applicable. Let us assume the subgraphs Q and R are pure and R introduces y . We have the following reduction.

1. $(\hat{y}R\hat{\beta}\cdot\alpha)\hat{\alpha}\dagger\hat{x}(\langle v\cdot\beta\rangle\hat{\delta}[x]\hat{z}Q) \rightarrow$
2. $(\langle v\cdot\beta\rangle\hat{\delta}\dagger\hat{y}(R\hat{\beta}\hat{\gamma}\hat{f}(f\cdot\sigma)))\hat{\sigma}\dagger\hat{z}Q$

Note that if the α -conversion steps had been instantaneous, we would have obtained:

$$2'. (\langle v\cdot\beta\rangle\hat{\delta}\dagger\hat{y}R')\hat{v}\dagger\hat{z}Q$$

(where $R' = R[\sigma/\beta]$). In Step 2, although R introduces y , we cannot evaluate the cut $\hat{\delta}\dagger\hat{y}$ without activating it first, which will cause it to (wastefully) propagate through R' (once the renaming cut $\hat{\beta}\hat{\gamma}\hat{f}$ has been evaluated) searching for sockets named y . Each time this effect is observed, the cost of the α -conversion scheme will be skewed by a factor proportional to the size of the subgraph R when no optimization is used, and by a constant factor when the (optimized) garbage collection rules are used.

One way of avoiding this is to ensure the reduction strategy prioritises reduction of the cuts which perform α -conversions. How these extra cuts are identified from renaming cuts that belong to the original circuit is also a point that needs to be addressed. In Example 6.1, this would involve evaluating the renaming cut $\hat{\beta}\hat{\gamma}\hat{f}$ in step 2, before the other cuts in the expression. This action would rightly prevent the cut $\hat{\delta}\dagger\hat{y}$ activating, then propagating, to the right.

Although the \mathcal{X} -calculus specifies three kinds of cut (inactive, left-activated and right-activated), an input term should only contain instances of inactive cuts⁵. When an inactive cut is chosen as the next redex, there is an implicit understanding that the cut should be run to ‘completion’. That is, when we choose to execute a redex $P\hat{\alpha}\dagger\hat{x}Q$, we are making the choice to connect *all* α 's in P with *all* x 's in Q ; a good reduction strategy will evaluate all cuts in this way. The following example enforces this view and illustrates why a naive outermost reduction strategy is not a favourable strategy in the setting of the \mathcal{X} -calculus.

⁵The system of active cuts was originally designed by Urban [24] to obtain a strongly normalising cut-elimination transformation on sequent proofs. They were specified as annotations on cut formulae used to direct cut-rule permutations to the leaves of a sequent derivation.

Example 6.2 We begin with a nesting of cuts between arbitrary terms P, Q, R , and S . A call-by-value left-most outermost reduction strategy is applied to the graph root node, which traverses the term attempting to match each rule with the current graph node. After a successful rewrite, the redex searching process restarts from the root node of the term graph. Recall that active cuts cannot propagate over each other.

1. $(P\hat{\alpha} \dagger \hat{x}Q)\hat{\beta} \dagger \hat{y}R \rightarrow (a \dagger)$
2. $(P\hat{\alpha} \dagger \hat{x}Q)\hat{\beta} \dagger \hat{y}R \rightarrow (\text{cut} \dagger)$
3. $(P\hat{\beta} \dagger \hat{y}R)\hat{\alpha} \dagger \hat{x}(Q\hat{\beta} \dagger \hat{y}R) \rightarrow (a \dagger)$
4. $(P\hat{\beta} \dagger \hat{y}R)\hat{\alpha} \dagger \hat{x}(Q\hat{\beta} \dagger \hat{y}R) \rightarrow \dots$

In this relatively simple example, we are left in a situation resembling a traffic-jam. By step (4), the propagation of the outermost cut is blocked by the innermost active cuts. When the innermost cut propagates down the graph one level, the second innermost cut is permitted to propagate down one level. This pattern expands to more complicated examples, where each outer cut follows in the wake of an innermost cut (as would be seen in a traffic-jam, one *cut* moves along a place, and each following *cut* shifts along, filling the empty space).

The overall effect of this is an undesired increase in the cost of searching for the next redex (which involves graph traversal, structural matching and checking side-conditions).

When a term graph is rewritten using the standard notion of graph rewriting, new nodes may be added to the graph. In step 1 of Example 6.2, an inactive cut is activated. Although the graph nodes *Cut* and *CutL* (for the cuts $\hat{\delta} \dagger \hat{z}$ and $\hat{\delta} \dagger \hat{z}$) are represented by two distinct node objects, our strategy must recognise that they are related in order to evaluate the cut to completion. The strategy will need to sequentially apply a number of rules to the term graph while ‘following’ the propagating cut through the expression.

Visser [26] proposes a generic language for specifying reduction strategies on term graphs. This language is rich enough to describe the kind of strategies we seek. The following section introduces the idea of *strategy combinators* and explains how to implement a reduction strategy for \mathcal{X} that can reduce a cut to completion.

6.1 Strategy Combinators

Strategy combinators as defined by Visser [26, 11, 27] allow for the controlled traversal of a data structure. These combinators are then extended with rewrite rules, which allow on-site modification of the graph. Examples of some strategies that may be constructed with these combinators are: (1) normalize the graph using a set of rules, R , according to an innermost strategy; (2) repeatedly apply a rewrite rule to a node until failure; or (3) visit all nodes at level three of the graph, and so on.

For our purposes, it is sufficient to restrict ourselves to a subset of the language made up of the following combinators.

Definition 6.3 (STRATEGY COMBINATOR LANGUAGE [26])

$s ::=$	id	Identity
	fail	Fail
	$L \rightarrow R$	Rewrite rule
	$[L \rightarrow R]$	List of Rewrite rules
	$\text{seq}(s, s)$	Sequential Composition
	$\text{choice}(s, s)$	Left-biased choice
	$\text{all}(s)$	All immediate children
	$\text{one}(s)$	One immediate child




Figure 7: Applications of Basic Strategy Combinators to Arbitrary Graphs

Definition 6.4 (APPLICATION OF A STRATEGY) The application of a strategy to an \mathcal{X} -graph is a pair consisting of a strategy combinator, s , and a rooted subgraph $(g|n)$. The strategy combinator system has a global *fail flag* which, when raised, indicates a *fail state*; this fail state affects the operational behaviour of some of the combinators.

We will write $s@n$ for the application of the strategy combinator s to the \mathcal{X} graph rooted at n .

The main results for the application of strategy combinators are either: (1) the strategy results in another set of strategies being applied to some node(s) of the graph; (2) the graph is modified by a rewrite rule; or (3) the state of the global fail flag is altered.

Figure 7 accompanies the following description of strategy combinators.

$id@n$: the identity strategy which simply leaves the supplied node unmodified.

$fail@n$: raise the fail flag indicating a state of failure.

$L \rightarrow R@n$: attempt to match the left-hand side L with the subgraph rooted at n . If the match is successful, n will be rewritten to some subgraph rooted at n' as dictated by the rewrite rule; any further strategies to be applied to n are updated to refer to n' . If the match is unsuccessful, the fail flag is raised.

$[L \rightarrow R]@n$: sequentially traverse a list of rewrite rules while attempting to apply each rewrite rule to n . The strategy terminates the traversal of the list upon the successful application of a rewrite rule. If the list is exhausted and no rule was applicable, the fail flag is raised.

$seq(s1,s2)@n$: sequentially apply the argument strategies to n ; if at any time a state of failure is reported the entire sequence strategy fails.

$choice(s1,s2)@n$: applies $s1@n$, then applies $s2@n$ if and only if $s1@n$ reported a failure.

$all(s)@n$: attempt to apply s to each immediate successor (left-to-right) of the node n . Any

successive applications are aborted if at any point the fail flag is raised by the application of s to the immediate successors of n .

$\text{one}(s)@(n)$: attempt to apply s left-to-right to a single immediate successor of n ; if no successful application is found, the entire strategy fails.

In addition to these basic combinators, we will allow users to define their own complex combinators via the following construction.

$$C(x_1, \dots, x_k) = s$$

C 's arguments x_1, \dots, x_n may occur, and are bound, in the definition body, s . Definition 6.3 is then extended with a case for a user-defined combinator:

$$s ::= \dots \\ | C(s_1, \dots, s_k) \quad \text{User-defined combinator}$$

An application of a user-defined combinator to a node n , $C(s_1, \dots, s_k)@(n)$, denotes the instantiation $s[x_1 = s_1, \dots, x_k = s_k]$ of the body of s in the definition of C . Because this extension allows recursive strategies to be defined, we will dismiss nonsense definitions such as $C(x) = C(x)$, by forbidding left-recursion.

We will also make use of some helper strategies, which can each be defined using the strategy language described above.

Definition 6.5 (HELPER STRATEGY COMBINATORS [27]) Below, we define some helper combinators, followed by an informal description of the effect of applying the combinator to a node of some term graph. We assume that, before the application of each strategy, the fail flag is not raised.

$\text{try}(s) = \text{choice}(s, \text{id})$: attempt to apply the argument strategy s to the node. If s fails, clear the fail flag.

$\text{repeat}(s) = \text{try}(\text{seq}(s, \text{repeat}(s)))$: repeatedly apply the strategy s to the node, until no more applications are possible, leaving the system in an unfail state.

$\text{oncetd}(s) = \text{choice}(s, \text{one}(\text{oncetd}(s)))$: search once top-down from the node and terminate after the first successful application of s . Raise the fail flag if no application was successful.

$\text{outermost}(s) = \text{repeat}(\text{oncetd}(s))$: search depth-first from the node and attempt to apply s to each node of the term graph; after a successful application restart the search from the node on which the strategy was first called (modulo rewriting of the subgraph rooted at that node).

6.2 Reduction Strategies for \mathcal{X}

In this section we will define a strategy combinator that when applied to \mathcal{X} -graphs will evaluate an inactive cut to completion. We will then extend this strategy to work with our proposed solutions to the problems of name clash and capture.

First, however, we will give a detailed example of the steps involved in applying a rewrite rule strategy to an \mathcal{X} -graph following a simple traversal scheme. The example is intended to mimic the steps taken by our implementation.

Example 6.6 (A REDUCTION USING oncetd) In Figure 8, we illustrate the steps taken by the strategy language to apply $\text{oncetd}(\text{cap})$ to the term graph

$$\llbracket \langle y \cdot \gamma \rangle \hat{\gamma} [z] \hat{k}(\langle k \cdot \alpha \rangle \hat{\alpha} \dagger \hat{x} \langle x \cdot \mu \rangle) \rrbracket_G$$

During the application of a strategy, we maintain:




Figure 8: Application of oncetd (*cap*) to $\langle y \cdot \gamma \rangle \hat{\gamma} [z] \hat{k} (\langle k \cdot \alpha \rangle \hat{\alpha} \dagger \hat{x} \langle x \cdot \mu \rangle)$.

a node stack : a stack of node ids that records the path of the strategy through the term;
a combinator stack : a stack of combinator states which records the progress of the strategy;
 and
a failure flag : to record whether a strategy has resulted in failure.

The stack trace in Figure 8(a) begins with the root node of the term graph (Figure 8(b)) and the root node of the strategy graph (Figure 9) on the node stack and combinator stack, respectively; the failure flag is cleared. We write ch/i to indicate the choice strategy is currently evaluating its i th argument strategy, and one/j to indicate the one strategy currently is applying its argument strategy to the j th child of the node at which it was first applied. We summarise the interesting steps of the strategy below.

- 1-2 : The zeroth argument of choice is pushed onto the combinator stack.
- 3-4 : The failed match of the rule (*cap*) with node 1, results in a failure state. The choice combinator recovers from the failure, expanding its second argument.
- 5 : The one strategy pushes the zeroth child of node 1 (i.e. node 2) onto the node stack, and pushes its argument strategy (one) onto the combinator stack. Notice that one's argument strategy results in a recursive call being made to choice.
- 6-9 : The application of the rule (*cap*) to node 2 fails. The choice combinator recovers from the failure, pushing the one strategy onto the combinator stack. Since one is a traversal combinator, it pushes node 2's zeroth child onto the node stack.
- 10-14 : The application of the rule (*cap*) to node 6 fails, and the choice strategy once again recovers from the failure. However, notice that this time the one strategy also fails, since node 6 has no children.
- 15 : The combinators are popped off the stack until a combinator is found that can reset the failure flag. The combinator happens to be the one combinator which was evaluating the zeroth child of node 2. one clears the failure flag and proceeds to apply its argument strategy to the first child of node 2 (i.e. node 7).
- 16-19 : A repeat of the steps 10-14 occurs, except with node 7 on the node stack.
- 20-21 : Upon returning to the one combinator which was visiting the first child of node 2, it finds node 2 has no more children. Therefore, the one strategy fails and propagates the failure state.
- 22 : The one strategy failed to apply its argument strategy to the zeroth child of node 1 (i.e. node 2). It recovers from the failure state and attempts to apply the same argument strategy to the first child of node 1 (i.e. node 3).
- 23-34 : This application fails, and one applies the argument strategy to the second child of node 1 (i.e. node 4), then to the third child (node 5) when this fails.
- 35 : The application of the rule (*cap*) is successful and node 5 is rewritten to node 8 (see Figure

	Node Stack	Combinator Stack	Fail		Node Stack	Combinator Stack	Fail
1	1	-		21	1,2,7	ch/1, one/0, ch/1, one/1	×
2	1	ch/0, (cap)		22	1,3	ch/1, one/1	
3	1	ch/0	×	23	1,3	ch/1, one/1, ch/0, (cap)	
4	1	ch/1		24	1,3	ch/1, one/1, ch/0	×
5	1,2	ch/1, one/0		25	1,3	ch/1, one/1, ch/1	
6	1,2	ch/1, one/0, ch/0, (cap)		26	1,3	ch/1, one/1, ch/1, one/0	
7	1,2	ch/1, one/0, ch/0	×	27	1,3	ch/1, one/1, ch/1	×
8	1,2	ch/1, one/0, ch/1		28	1,4	ch/1, one/2	
9	1,2,6	ch/1, one/0, ch/1, one/0		29	1,4	ch/1, one/2, ch/0, (cap)	
10	1,2,6	ch/1, one/0, ch/1, one/0, ch/0, (cap)		30	1,4	ch/1, one/2, ch/0	×
11	1,2,6	ch/1, one/0, ch/1, one/0, ch/0	×	31	1,4	ch/1, one/2, ch/1	
12	1,2,6	ch/1, one/0, ch/1, one/0, ch/1		32	1,4	ch/1, one/2, ch/1, one/0	
13	1,2,6	ch/1, one/0, ch/1, one/0, ch/1, one/0		33	1,4	ch/1, one/2, ch/1	×
14	1,2,6	ch/1, one/0, ch/1, one/0, ch/1	×	34	1,5	ch/1, one/3	
15	1,2,7	ch/1, one/0, ch/1, one/1		35	1,5	ch/1, one/3, ch/0, (cap)	
16	1,2,7	ch/1, one/0, ch/1, one/1, ch/0, (cap)		36	1,8	ch/1, one/3, ch/0	
17	1,2,7	ch/1, one/0, ch/1, one/1, ch/0	×	37	1	-	
18	1,2,7	ch/1, one/0, ch/1, one/1, ch/1					
19	1,2,7	ch/1, one/0, ch/1, one/1, ch/1, one/0					
20	1,2,7	ch/1, one/0, ch/1, one/1, ch/1	×				

Figure 9: Stack Traces

8(c)), and the node stack is updated.

36-37 : choice does not evaluate its second argument strategy since the failure state is clear when it is the head of the combinator stack. The remaining combinators are popped off the stack and the reduction terminates.

Many optimisations have been made to this reduction engine that can bypass a significant number of stack operations. For example, noticing that the *Cap*-nodes have only child nodes which are names (i.e. no variable nodes), steps 9-21 may be skipped.

In the previous section, we discussed some features we would require in an λ -calculus reduction strategy—these are summarised below.

Definition 6.7 (CRITERIA FOR EVALUATING A CUT) Given a pure λ -term, for any single *inactive* cut, a ‘good’ λ -calculus reduction strategy will:

- i) Evaluate the inactive cut, $P\hat{\alpha} \dagger \hat{x}Q$, to completion so that all the plugs α in P are connected with all the sockets x in Q . The resultant term should have eliminated the cut $\hat{\alpha} \dagger \hat{x}$ from the term, as well as the connectors α and x .
- ii) Transform a pure term to a pure term. This will involve ‘following’ related active cuts through the term until they are all destroyed.
- iii) Give priority to α -conversion/renaming cuts over other cuts so that they do not interfere with parts ((a) and ((b)).

We will consider the (call-by-name) λ -calculus rules as described in Section 2 to explain how the criteria is met. Realising that the λ -calculus rewrite rules are fixed, we can use our knowledge of the rules to guide the propagation of active cuts through λ -terms. We first define the following strategies as strict partitions over the set of rewrite rules (as ordered lists).

Definition 6.8

```

rename   = [(cap), (exp), (med)]
logical  = [(exp-impN), (exp-impV), renaming]
activate = [(λa), (aλ)]
prop_a1  = [(exp-insλ), (impλ), (cutλ), (λexp), (λimp-ins), (λcut)]
prop_a2i0 = [(exp-outsλ), (λimp-outs)]
gc       = [(capλ), (λcap)]
deact    = [(dλ), (λd)]

```

Notice that, in fact, rule $(exp-imp_V)$ is obsolete here, since it will never match if $(exp-imp_N)$ does not.

Using the above combinators, we will describe a user-defined strategy combinator for reducing an inactive cut so that the criteria outlined in Definition 6.7 are obeyed.

The first two rule definitions (rename and logical) deal directly with inactive cuts. Their application is straightforward since they destroy the inactive cut being reduced. The next rule definition (activate) replaces an inactive cut with an active cut. To uphold part (c) of our criteria, we must continue evaluation of this active cut until it is destroyed. To that purpose, we will build a user defined strategy, `propagate()`, which propagates an active cut through a term until it is destroyed. Clearly `propagate()` will make use of the remaining definitions since they cover the cases that deal with activated cuts. We can now define a strategy `evalcut()` which evaluates a cut to completion.

```

evalcut() = choice(logical,
                   seq(activate, propagate())
                   )

```

This strategy considers the *only* two cases of how to evaluate an inactive cut: it can either be reduced by a logical rule, or activated and then propagated through the term. The remainder of this section looks at how to define the `propagate()` combinator.

We can break down the work that needs to be done by the `propagate()` combinator into four cases: (1) propagate an active cut through a circuit that *does not* mention any connectors involved in the cut; (2) propagate an active cut through a circuit that *does* mention the connectors involved in the cut; (3) garbage collect the active cut, since it has reached the level of the capsules; and (4) deactivate the cut, and attempt to reduce the deactivated cut.

The first two cases are covered by the following combinators.

```

(1) seq(prop_a1,
        all(try(propagate())))
    )
(2) seq(prop_a2i0,
        seq(all(all(try(propagate()))),
            evalcut())
        ) )

```

We explain these terms in the following paragraphs. First we remark that by nesting several ‘all’ combinators, we can visit all the nodes at a particular depth of a term graph (relative to the node at which the strategy was first applied). For instance, `all(s)@(n)` will apply `s` to the nodes at depth 1 relative to the node `n`, while `all(all(s))@(n)` will apply `s` to the nodes of `n` at

depth 2, and so on.

For (1), all the rules in the list `prop_a1` have active cuts (that must be further propagated) at depth 1. Therefore, after a rule in `prop_a1` has been applied, the `propagate()` strategy is applied to the nodes at depth 1 to propagate these newly introduced cuts.

For (2), all the rules of `prop_a2i0` have active cuts at depth 2 *in addition to* an inactive cut at depth 0. The newly introduced active cuts at depth 2 must be further propagated, so there is a double nesting of the all combinator. The inactive cut at depth 0 also needs to be evaluated, so the `evalcut()` strategy is recursively applied; notice that to avoid the situation shown in Example 6.2, the active cuts are propagated *before* the inactive cut is evaluated.

The final two cases for the `propagate` strategy are described by the following combinators:

- (3) `gc`
- (4) `seq(deact, evalcut())`

For (3), the strategy `gc` is applicable if the active cut is with a capsule that does not introduce the connector bound by the active cut. The rewrite rules state that in this case the active cut should be destroyed. Therefore, following a successful application of a `gc` strategy, no further work needs to be done: there is no more active cut to propagate.

For (4), if the strategy `deact` is applicable, the active cut is deactivated, creating a new inactive cut at that level which must be evaluated; therefore the `evalcut` strategy is recursively applied to that inactive cut.

Combining these four parts, we obtain a definition for the `propagate()` strategy.

```
propagate() = choice(gc,
                    choice(seq(deact, evalcut()),
                          choice(seq(prop_a1,
                                     all(try(propagate())))),
                          ),
                    seq(prop_a2i0,
                       seq(all(all(try(propagate()))),
                           evalcut())
                    ) ) ) )
```

The nesting of ‘choice’ combinators ensures each propagation case is considered at the current node. If no case is successful, the strategy combinator leaves the system in the fail state.

The above discussion shows how to evaluate a cut to completion in the CBN λ -calculus - notice that in `activate`, rule (λa) comes first, as does $(exp\text{-}imp_N)$ (see also Definition 6.9). We would like to extend this strategy to the sets of rules that solve the name clash and name capture problems described in Section 5. This involves working with a larger set of rules, but the same idea of evaluating a cut exists. We have two sets of rules to consider: the renaming using activated-cuts schemes and the lazy-copying α -conversion schemes with special rebinding symbols.

6.3 α -conversion with Rebinding Nodes

In this section, we look to define a ‘good’ strategy for evaluating an inactive cut for the solution of lazy copying as defined in Section 5.1. We can follow a similar strategy to that of the previous section, except some extra care must be taken with regards to the rebinding nodes. The nodes which perform the rebinding (`rs` and `rp`) are not part of the λ -calculus and would ideally be transparent.

Barendsen and Smetsers remark in [8]:

“ *mixing copy rules with the reduction rules [in the set of rewrite rules] may destroy properties [of the set of rewrite rules] such as confluence, or at least make it very difficult to check whether known properties [of the term graph rewriting system] extend to [the term graph rewriting system with copy rules added].* ”

This turns out to be true in our case; the rebinding nodes interfere with structural matching and the ‘introduces’ side conditions of the rewrite rules. For example, the term $rs(\langle x \cdot \alpha \rangle, \alpha, \beta)$ does not introduce β , although it evaluates to $\langle x \cdot \beta \rangle$.⁶ A naive solution could force the copying operation to completion, though this may turn out to be unnecessary.

We propose a lazier solution that *hides* the existence of rebinding nodes from the structural matching step of the rewriting procedure. Observing the lazy copying rewrite rules (Definition 5.2), this can be achieved by ensuring all rebinding nodes are at least two levels from any redex; in other words, a rebinding node should never be an immediate successor of a cut.

We specify a strategy `pushRebind()`, which when applied to a rebinding node, pushes that rebinding node down through the term-graph by one level. There are two cases to consider—either a rebinding rule (from Definition 5.1) is directly applicable to the current rebinding node, or not — i.e. there is a chain of one or more successive rebinding nodes prohibiting propagation. When such a chain exists, the strategy traverses to the lowest rebind node of that chain where a rule *will* be applicable. The lowest node is propagated one level further, followed in turn by each blocked ancestor.

The skeleton definition of this strategy is given below and makes use a group of all the rebinding evaluation rules (`rebind_rules`) as given in Definition 5.1, and a strategy `repeat'(s)` which repeatedly applies its argument strategy to the current node or fails.

```
repeat'()      = seq(s, repeat'(s))
pushRebind() = choice(repeat'(rebind_rules),
                      choice(seq(all(try(pushRebind()))),
                              rebind_rules
                              ),
                      fail
                      )
)
```

The `propagate()` strategy can now be extended to make use of `pushRebind()`. For each rebinding node introduced by a rewrite rule of the λ -calculus, the `pushRebind()` strategy is applied to that node, guaranteeing it is never the successor of a cut.

6.4 α -conversion with Renaming Cuts

The strategy `evalcut()` is first extended to cater for the case of ($exp\text{-}imp_N$) when a plug needs renaming. In this case, the active cut (at depth 2) is propagated through the term, using

```
seq(exp-impN,
    all(all(try(propagate()))))
)
```

⁶We could alter our definition of *introduces* so that $rs(\langle x \cdot \alpha \rangle, \alpha, \beta)$ introduces β , but this will still not side-step the problem of the rebinding node structure interfering with the graph matching process.

The `propagate()` strategy should then be modified to cater for any additional α -conversion structure. This involves partitioning this larger set of rules into lists of rules which have cuts (to be further propagated) at common depths. The renaming cuts, which will be the innermost active cuts, must then be given priority over the other cuts in the rule. We will consider the variant of the rule (λexp):

$$Q\hat{\alpha}\lambda\hat{x}(\hat{y}P\hat{\beta}\cdot\delta) \rightarrow \hat{w}(Q\hat{\alpha}\lambda\hat{x}((w\cdot\mu)\hat{\mu}\lambda\hat{y}P))\hat{\beta}\cdot\delta \quad (y \in fs(Q), \beta \notin fp(Q))$$

(the side-conditions of the rule do not come into play in the following discussion).

This rule would have been placed in a ‘list of rewrite rules’ combinator named `prop_a1a2`, indicating the rule has active cuts at a depth of 1 and also at a depth of 2; in the rule above these are respectively the cuts $\hat{\alpha}\lambda\hat{x}$ and $\hat{\mu}\lambda\hat{y}$ on right-hand side of the rule. The propagation of active cuts for this combinator is described by the combinator:

```
seq (prop_a1a2,
      seq (seq (all (all (try (propagate())))),
            all (try (propagate()))))
) ) )
```

In other words, if a rule from the combinator `prop_a1a2` is applied, propagate the inner active cuts at depth 2, and then propagate the remaining inner active cuts at depth 1. Lists of rules grouped together by the common depth of active and inactive cuts also need to be modified appropriately.

Now a ‘good’ outermost reduction strategy for the λ -calculus can be defined as:

```
outermost(evalcut())
```

This strategy will search for the outermost inactive cut, which, when found, will apply the `evalcut()` strategy to that cut, evaluating that cut to completion.

6.5 Optimisations

We highlight a simple optimisation to the outermost strategy that will greatly decrease the search time for the next redex. Currently, after a successful reduction of an inactive cut, the search for the next inactive cut restarts from the point at which the original call to the strategy was made, i.e. the root node of the term graph. For a general outermost strategy, this is a safe course of action, since an outermost redex may have been skipped while an inner redex is evaluated c.f. call-by-value reduction in the λ -calculus. The λ -terms we evaluate are pure terms, and according to our `evalcut()` reduction strategy inactive cuts are evaluated to completion. Since inactive cuts cannot block propagating active cuts, the depth the subsequent outermost cuts are therefore guaranteed to be at a depth lower or equal to the current node. Using this fact, we can define an outermost reduction strategy to continue redex-searching from the current node pointed to.

```
outermost'(s) = seq (repeat(s),
                    all (try (outermost'(s))))
) )
```

By parametising `evalcut` with an ordered list of activation rules and a variant of the ($exp-imp$) rule, we can define two outermost strategies for call-by-name and call-by-value as follows.

Definition 6.9 (CBN AND CBV REDUCTION STRATEGIES)

$$\begin{aligned}
\text{cbnact} &= [(\lambda a), (a\lambda)] \\
\text{cbvact} &= [(a\lambda), (\lambda a)] \\
\text{outermostCBN} &= \text{outermost}'(\text{evalcut}(\text{cbnact}, (\text{exp-imp}_N))) \\
\text{outermostCBV} &= \text{outermost}'(\text{evalcut}(\text{cbvact}, (\text{exp-imp}_V)))
\end{aligned}$$

6.6 Benchmarks

In this section we present our benchmarks comparing the costs of the solutions to name clash and capture we proposed above. In the previous section we described how to extend our implementation with ‘strategy combinators’ in order to define a fair reduction strategy that could be used to compare the proposed solutions. Incorporating the strategy combinator language into our tool presented us with some problems, which we summarise below.

Visser has provided the community with a Java implementation of the strategy combinator framework, called JJTraveler. The framework allows modular extensions to the combinator language, allowing one to add *user-defined combinators* to the system by inheritance. A full description of this framework can be found in [27].

Integration of the framework with our \mathcal{X} implementation, also written in Java, was straightforward. Unfortunately, preliminary testing revealed the implementation was unable to traverse some of the larger term-graphs generated by our benchmarks (which can contain in excess of 300,000 nodes) resulting in stack-overflows. The reason for this was the heavy reliance on recursion due to the use of a modified Visitor design-pattern⁷. We chose to re-implement the framework taking an iterative approach instead. The set of strategy combinators (Definition 6.3) extended with ‘user-defined combinators’ allows recursive strategies to be built. These were implemented as cyclic graphs (following the implementation of JJTraveller).

To maintain the state of the strategy (i.e., to track how much of the strategy had been processed during an application), we used two stacks. The working details of these structures were exemplified in Example 6.6. Recall that the ‘combinator stack’ tracked the current position within the strategy combinator graph, and the ‘node stack’ tracked the node of term graph which the current combinator was being applied to. The approach allowed us to obtain an accurate measure of the cost of *traversing* the graph searching for redexes: it was a count of the number of (constant-time) stack operations, plus the number of node matchings, plus the cost of checking the side-conditions.

We chose to measure the running of (interpreted) λ -terms since these are well-known benchmarks [1, 19], and the efficiency of the various formalisms and abstract machines can be better compared. We can of course not confront their (published) run-time measurements because of differences in platforms and processor architectures.

We use the usual encoding for Church Numerals ($n = \lambda xy.x^n y$). In addition, we use the combinators $P = (\lambda z.(\lambda x.zxxx)(\lambda y.2(\lambda x.y(xl))n))l$, with n replaced by a chosen Church Numeral and $l = \lambda x.x$.

Note that we do not wish to consider the cost of encoding λ -terms to \mathcal{X} as we are interested in comparing the efficiency of the α -conversion mechanisms for the systems presented; because of this, we use Prawitz’s encoding (Definition 2.19). Our benchmarking results are listed in Figures 10 and 11. For each test case we record the following two measurements:

⁷Simply increasing the stack size of the JVM was not seen to be a scalable solution. Although recent Java implementations do include recursion optimisations, these mainly work on performance. Since our benchmarks will count atomic operations rather than measure time, our main concern is heap usage, which we can manage more efficiently with an iterative approach.

TestCase	rebind - GC		α -safe - GC		α -safe + GC		avoid capt. - GC	
	Search	Rewrite	Search	Rewrite	Search	Rewrite	Search	Rewrite
22II	0.861	0.163	0.264	0.0765	0.159	0.0421	0.0994	0.0330
222II	62.0	15.7	3.64	0.942	0.841	0.208	0.495	0.153
2222II	-	-	-	-	4650	715	1690	522
210II	0.372	0.0682	0.140	0.0401	0.0762	0.0214	0.0548	0.0194
2210II	2.40	0.471	0.649	0.165	0.269	0.0660	0.173	0.0541
22210II	225	56.7	10.6	1.84	1.53	0.305	0.805	0.236
P2	32.9	7.85	4.01	0.857	0.999	0.218	0.429	0.145
P3	51.4	12.5	5.54	1.11	1.20	0.256	0.487	0.165
P5	110	27.2	9.94	1.74	1.63	0.331	0.602	0.206
P10	429	109	31.4	4.01	2.98	0.519	0.890	0.309
P20	2240	575	143	11.6	6.71	0.895	1.47	0.514
P50	26100	6780	1550	58.4	26.3	2.02	3.19	1.13

Figure 10: CBV Results: Cost measured in units of 10^6 operations (to 3.s.f)

TestCase	rebind + GC		α -safe - GC		α -safe - GC		avoid capt. - GC	
	Search	Rewrite	Search	Rewrite	Search	Rewrite	Search	Rewrite
22II	0.914	0.156	0.314	0.0908	0.173	0.0439	0.0963	0.0322
222II	17.0	3.27	3.86	1.01	1.42	0.305	0.551	0.178
2222II	-	-	146000	29700	13100	1460	2490	805
210II	0.458	0.0819	0.164	0.0441	0.0868	0.0227	0.0559	0.0198
2210II	2.12	0.360	0.726	0.177	0.315	0.0742	0.151	0.0507
22210II	32.1	6.38	7.29	1.60	2.11	0.441	0.754	0.246
P2	540	138	22.5	3.02	3.51	0.582	0.902	0.308
P3	875	224	30.6	3.67	4.23	0.670	1.04	0.354
P5	2000	515	52.6	5.22	5.85	0.847	1.32	0.447
P10	8670	2240	152	10.5	11.0	1.29	2.02	0.679
P20	49600	12900	624	27.2	25.9	2.17	3.41	1.14
P50	626000	163000	6050	126	107	4.81	7.59	2.54

Figure 11: CBN Results: Cost measured in units of 10^6 operations (to 3.s.f)

Search Cost : a count of atomic operations involved in traversing the graph and searching for redexes, i.e. the number of push/pop stack operations to evaluate the strategy, plus the number of attempted matchings made between the rule heads and graph nodes, plus the cost of testing the side-condition.

Rewrite Cost : a count of the (more expensive) graph transforming operations, i.e. the number of nodes added and deleted plus the number of edges added and deleted.

A straightforward numerical comparison of costs suggests the following relationship of efficiency between α -conversion schemes, under either reduction strategy (CBN or CBV).

$$\text{Rebinding GC} < \alpha\text{-Safe - GC} < \alpha + \text{Safe GC} < \text{Capture Avoid GC}$$

Another observable trend is the linear relationship between the redex search cost and graph rewrite cost. As the size of the graph increases, the search cost increases since more operations are required to traverse the graph structure. Once a suitable redex (a cut) is found, the cost of reducing that cut is related to the size of its subterms since it must be propagated through them.

There is a significant difference between the cost of reduction under the rebinding scheme

versus the other schemes tested. In fact, our original attempt at implementing the rebinding solution was not lazy at all, and often resulted in memory requirements greater than the 2GiB limit. We recorded the size of the graph (number of node objects) as the reduction progressed for each test-case in order to gain insight into this vast requirement of system resources. Results from the eager strategy are shown in Figures 12(a) and 13(a). As explained in Section 5.1, the rebinding scheme works by destroying the sharing in portions of the term-graph as so to guarantee that no binder occurs more than once. This copying-out effect, seen as peaks in the graphs, shows an increase in the number of nodes whenever a cut is propagated through varying sizes of subterm causing all sharing to be destroyed. Looking closer at these graphs the cost of searching is also clearly visible. As the number of nodes in a graph increases, small horizontal ‘platforms’ can be observed. These regions represent pure search costs consisting only of traversal stack operations plus unsuccessful rule matches.

Switching to a lazy mechanism as detailed in Section 6.3, although still relatively expensive, kept the size of the graph low enough for many previously failed tests to run to completion. The results of the lazy strategy are shown in Figures 12(b) and 13(b). The CBV graph highlights nicely the copying out of each argument when it is supplied to a function.

We also investigated the variation of graph size under the α -conversion schemes; the results are displayed in Figures 12(c) and 13(c). The shapes of these graphs appear to be less random than those from the rebinding scheme, and we see that for a particular reduction strategy (CBN or CBV) the overall shape of the graphs are similar. This is an expected side effect of the reduction strategy `evalcut()`, designed to make α -conversions transparent from the point of view of the term being reduced. The reductions therefore only diverges at points where the need for α -conversions differ.

Conclusions and future work

We have studied various solutions for the problem of name capture in the context of a term-graph rewriting implementation of the λ -calculus. The first uses a rebinding technique, that required an extension of the syntax, with additional rules. The second and third change the reduction rules of λ , but without extending the signature. The second solution guarantees that (generated) term adhere to Barendregt’s convention, whereas the third checks for captivation of free names. We have measured the efficiency of all these solutions, and conclude that the latter, although syntactically very close to the second, is by far the best.

There are a number of questions still open that will be investigated in future work. First of all, we need to understand why the CBN measurement of the pure α -safe system are only for the Q terms so much worse than the CBV measurements. We aim to show that the reduction strategy used in the tool is normalising.

References

- [1] A. Asperti, C. Giovanetti, and A. Naletto. The Bologna Optimal Higher-Order Machine. *Journal of Functional Programming*, 6(6):763–810, 1996.
- [2] S. van Bakel, S. Lengrand, and P. Lescanne. The language λ : Circuits, Computations and Classical Logic. In M. Coppo, Elena Lodi, and G. Michele Pinna, editors, *Proceedings of Ninth Italian Conference on Theoretical Computer Science (ICTCS’05), Siena, Italy*, volume 3701 of *Lecture Notes in Computer Science*, pages 81–96. Springer Verlag, 2005.
- [3] S. van Bakel and P. Lescanne. Computation with Classical Sequents. *Mathematical Structures in Computer Science*, 18:555–609, 2008.
- [4] S. van Bakel and J. Raghunandan. Implementing λ . In *Electronic Proceedings of Second International Workshop on Term Graph Rewriting 2004 (TermGraph’04), Rome, Italy*, Electronic Notes in Theoretical Computer Science, 2005.

- [5] H. Barendregt. *The Lambda Calculus: its Syntax and Semantics*. North-Holland, Amsterdam, revised edition, 1984.
- [6] H.P. Barendregt, M.C.J.D. van Eekelen, J.R.W. Glauert, J.R. Kennaway, M.J. Plasmeijer, and M.R. Sleep. Term graph rewriting. In *Proceedings of PARLE, Parallel Architectures and Languages Europe*, Eindhoven, The Netherlands, volume 259-II of *Lecture Notes in Computer Science*, pages 141–158. Springer Verlag, 1987.
- [7] H.P. Barendregt, M.C.J.D. van Eekelen, J.R.W. Glauert, J.R. Kennaway, M.J. Plasmeijer, and M.R. Sleep. Towards an Intermediate Language based on Graph Rewriting. In *Proceedings of PARLE, Parallel Architectures and Languages Europe*, Eindhoven, The Netherlands, volume 259-II of *Lecture Notes in Computer Science*, pages 159–175. Springer Verlag, 1987.
- [8] E. Barendsen and S. Smetsers. Extending Graph Rewriting with Copying. In Hans Jürgen Schneider and Hartmut Ehrig, editors, *Proceedings of International Workshop ‘Graph Transformations in Computer Science’*, Dagstuhl, Germany, January 1993, volume 776 of *Lecture Notes in Computer Science*, pages 51–70. Springer Verlag, 1994.
- [9] E. Barendsen and S. Smetsers. Conventional and Uniqueness Typing in Graph Rewrite Systems. *Mathematical Structures in Computer Science*, 1996.
- [10] R. Bloo and K.H. Rose. Preservation of Strong Normalisation in Named Lambda Calculi with Explicit Substitution and Garbage Collection. In *CSN’95 – Computer Science in the Netherlands*, pages 62–72, 1995.
- [11] M. Bravenboer, K. Trygve Kalleberg, R. Vermaas, and E. Visser. *Stratego/XT Tutorial, Examples, and Reference Manual for Stratego/XT 0.16*. Department of Information and Computing Sciences, Universiteit Utrecht, Utrecht, The Netherlands, November 2005.
- [12] A. Church. A Note on the Entscheidungsproblem. *Journal of Symbolic Logic*, 1(1):40–41, 1936.
- [13] P.-L. Curien and H. Herbelin. The Duality of Computation. In *Proceedings of the 5th ACM SIGPLAN International Conference on Functional Programming (ICFP’00)*, volume 35.9 of *ACM Sigplan Notices*, pages 233–243. ACM, 2000.
- [14] G. Gentzen. Investigations into logical deduction. In *The Collected Papers of Gerhard Gentzen*. Ed M. E. Szabo, North Holland, 68ff (1969), 1935.
- [15] Stefano Guerrini. Sharing Implementations of Graph Rewriting Systems. *Electronic Notes in Theoretical Computer Science*, 127(5):113–132, 2005.
- [16] J. R. Kennaway, J. W. Klop, M. R. Sleep, and F. J de Vries. The adequacy of term graph rewriting for simulating term rewriting. In M.R. Sleep, M.J. Plasmeijer, and M.C.D.J. van Eekelen, editors, *Term Graph Rewriting: Theory and Practice*, pages 157–168. J. Wiley & Sons, 1993.
- [17] S.C. Kleene. *Introduction to Metamathematics*. North Holland, Amsterdam, 1952.
- [18] S. Lengrand. Call-by-value, call-by-name, and strong normalization for the classical sequent calculus. In B. Gramlich and S. Lucas, editors, *3rd Workshop on Reduction Strategies in Rewriting and Programming (WRS 2003)*, volume 86 of *Electronic Notes in Theoretical Computer Science*. Elsevier, 2003.
- [19] I. Mackie. Efficient lambda-Evaluation with Interaction Nets. In *RTA*, pages 155–169, 2004.
- [20] M. Parigot. An algorithmic interpretation of classical natural deduction. In *Proceedings of 3rd International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR’92)*, volume 624 of *Lecture Notes in Computer Science*, pages 190–201. Springer Verlag, 1992.
- [21] D. Prawitz. *Natural Deduction, A Proof-Theoretical Study*. Almqvist & Wiksell, 1965.
- [22] R. Sleep, M.J. Plasmeijer, and M.C.J.C van Eekelen, editors. *Term Graph Rewriting. Theory and Practice*. Wiley, 1993.
- [23] A.J. Summers. Interpretation of λ -calculus terms to \mathcal{X} in as suggested by Prawitz. Personal communication, 2007”.
- [24] C. Urban. *Classical Logic and Computation*. PhD thesis, University of Cambridge, October 2000.
- [25] C. Urban and G.M. Bierman. Strong normalisation of cut-elimination in classical logic. *Fundamenta Informaticae*, 45(1,2):123–155, 2001.
- [26] E. Visser and Z. Benaissa. A core language for rewriting. *Electronic Notes in Theoretical Computer Science*, 15, 1998.
- [27] J. Visser. Visitor Combination and Traversal Control. In *OOPSLA*, pages 270–282, 2001.
- [28] P. Wadler. Call-by-Value is Dual to Call-by-Name. In *Proceedings of the eighth ACM SIGPLAN international conference on Functional programming*, pages 189 – 201, 2003.
- [29] C.P. Wadsworth. *Semantics and pragmatics of the lambda calculus*. PhD thesis, Oxford University, 1971. Thesis CST-33-85.

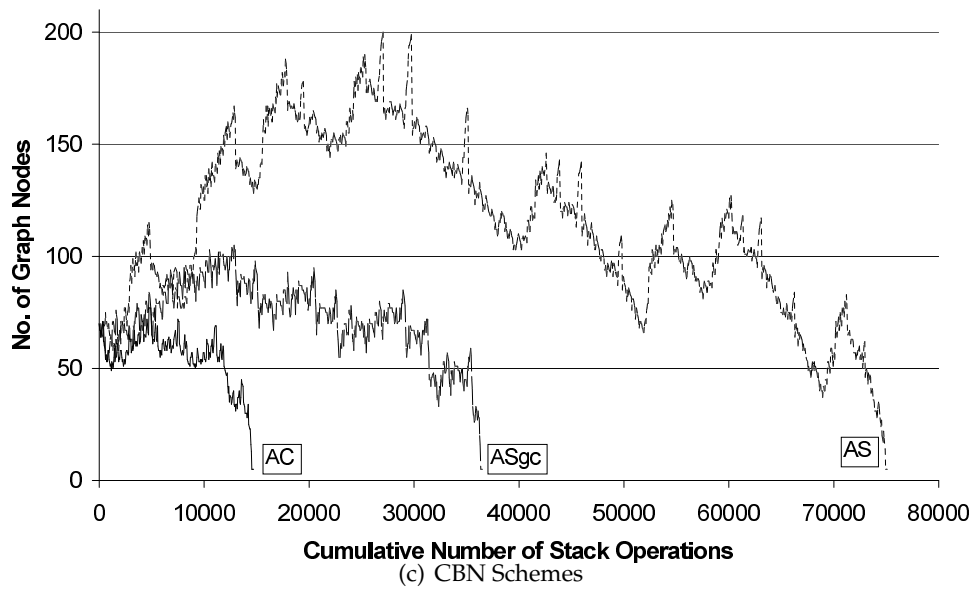
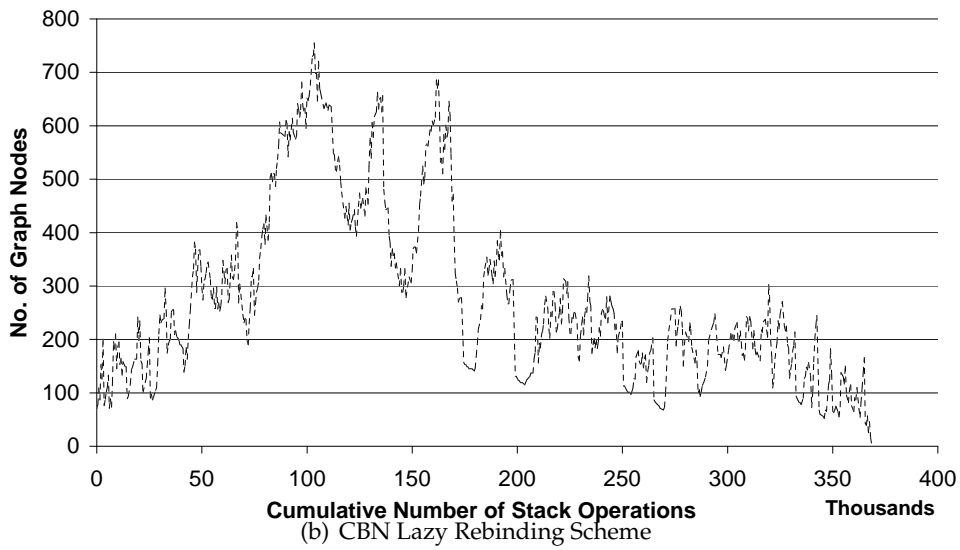
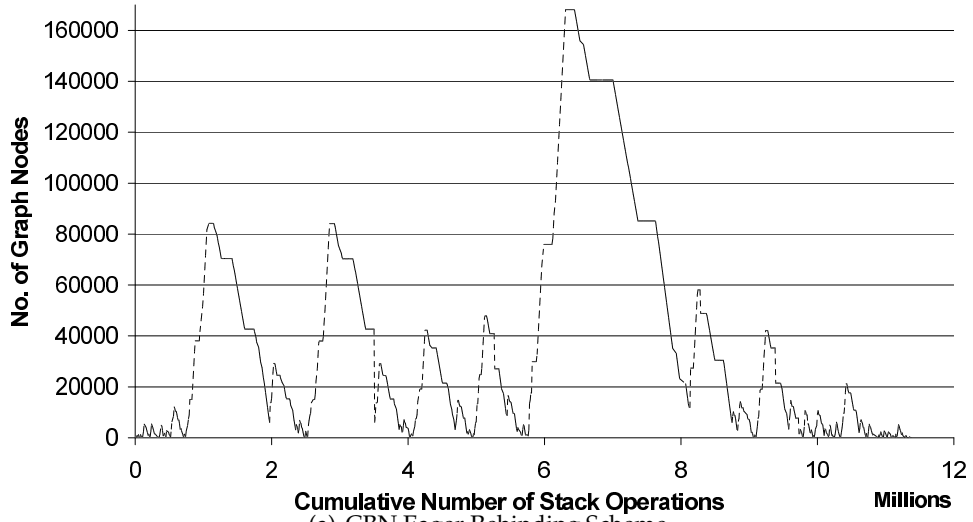


Figure 12: Variation of Graph Size over ‘time’ for reducing $\llbracket 222II \rrbracket_\delta$ under different α -conversion schemes (AS= α Safe-NoGC, ASgc= α Safe-GC, AC=AvoidCapture-GC)




Figure 13: Variation of Graph Size over 'time' for reducing $\llbracket 222II \rrbracket_\delta$ under different α -conversion schemes (AS= α Safe-NoGC, ASgc= α Safe-GC, AC=AvoidCapture-GC)