Sharing continuations: proofnets for languages with explicit control

Julia L. Lawall[*]
Harry G. Mairson[†]
Computer Science Department
Brandeis University
Waltham, Massachusetts 02454-9110
{jll,mairson}@cs.brandeis.edu

Abstract

We introduce graph reduction technology that implements functional languages with control, such as Scheme with `call/cc`, where continuations can be manipulated explicitly as values, and can be optimally reduced in the sense of Lévy. The technology is founded on *proofnets* for multiplicative-exponential linear logic, extending the techniques originally proposed by Lamping, where we adapt the continuation-passing style transformation to yield a new understanding of sharable values. Confluence is maintained by returning multiple answers to a (shared) continuation.

Proofnets provide a *concurrent* version of linear logic proofs, eliminating structurally irrelevant sequentialization, and ignoring asymmetric distinctions between inputs and outputs—dually, expressions and continuations. While Lamping's graphs and their variants encode an embedding of intuitionistic logic into linear logic, our construction implicitly contains an embedding of classical logic into linear logic.

We propose a family of translations, produced uniformly by beginning with a continuation-passing style semantics for the languages, employing standard codings into proofnets using call-by-value, call-by-name—or hybrids of the two—to locate proofnet *boxes*, and converting the proofnets to direct style. The resulting graphs can be reduced simply (cut elimination for linear logic), have a consistent semantics that is preserved by reduction (geometry of interaction, via the so-called *context semantics*), and allow shared, incremental evaluation of continuations (optimal reduction).

1 Introduction

Expressions and continuations are dual, separate but equal computational structures in a programming language. The former provides a value; the latter consumes it. Yet evaluating expressions is very familiar, while evaluating continuations is considered esoteric, even though both are made of the same stuff. The incorporation of continuations as first-class citizens in programming languages was not welcomed like the Emancipation Proclamation, but instead regarded warily as a kind of witchcraft, with implementation pragmatics that are ill-defined and unclear. If expressions and continuations are indeed dual, then so should be the technology of their implementation, and the flexibility with which we reason about them. Efficient evaluation of one should reveal dual strategies for evaluating the other. In short, everything we know about expressions we ought to know about continuations.

We take a significant step towards this equality by formulating a general version of graph reduction which implements the sharing and optimal incremental evaluation of both expressions and continuations, each evaluated using the same primitive operations. By founding our technology on generic tools from logic and programming language theory, encompassing the CPS transform and its relation to linear logic, we are

[*]Supported by NSF Grant EIA-9806718.
[†]Supported by NSF Grants CCR-9619638 and EIA-9806718, and the Tyson Foundation.

for the first time able to produce a family of related implementations with different efficiency characteristics. In bringing ideas from logical theory closer to implementation technology, we hope to make researchers think about the pragmatics of continuations in simple, novel and useful ways.

Our methodology is founded on proofnets[1] for multiplicative-exponential linear logic, following the beautiful insights of John Lamping [20], who realized Jean-Jacques Lévy's specification of correct, optimal reduction for the λ-calculus [21], and of Gonthier, Abadi, and Lévy, who reinterpreted Lamping's insights in the guise of Girard's geometry of interaction, and the related embedding of intuitionistic logic in linear logic [14, 15]. Linear logic [12, 13] provides an ideal substrate for the implementation of control operators, as it makes no asymmetric distinctions between inputs and outputs, or analogous expressions and continuations. We extend the optimal reduction technology to implement explicit control and sharing of continuations, essentially via an embedding of *classical* logic in linear logic, following a line of research beginning with Griffin and then Murthy [16, 24]. Our construction is based on continuation-passing style, but generates direct-style graphs in which continuations are implemented as sharable (*boxed*) values. This approach extends to implement most any functional language with abortive control operators whose semantics can be described in continuation-passing style—for example, Filinski's *symmetric λ-calculus* [9], and Parigot's *λ_μ-calculus* [26].[2]

What do optimality and correctness mean in a language with explicit control? To understand *optimality* and *sharing* in the context of continuations, consider the evaluation of a Scheme expression

```
([fun1]
   (call/cc (lambda (a) ([fun2]
      (call/cc (lambda (b) [exp])))))))
```

in the context of some complex continuation k. The expression [exp] can *implicitly* access the current continuation c, or *explicitly* access continuations a and b named by `a` and `b`. All of these continuations extend the continuation of the entire expression k. Optimality ensures that a, b, and c manipulate a single, shared copy of k, that b and c have a shared copy of a, and so on. If continuations are not copied in their entirety, duplicate work can be avoided as continuations are simplified.

A reduction strategy ρ is *correct* when for any strategy σ reducing expression E, if σ does not diverge during evaluation, neither does ρ. In the case of the pure λ-calculus, normal forms are unique (up to η-equivalence), so all correct strategies produce the same normal form, if one exists. However, control operators destroy the uniqueness of normal forms: let E be `(call/cc (lambda (k) ((lambda (x) 1) (k 2))))`; evaluating E returns 1 in a call-by-name strategy for Scheme, and returns 2 in a call-by-value strategy for Scheme. We can use this term to show that a correct evaluation strategy cannot simply choose one of these answers. Define context C as `(if (= [−] 1) 0` $\perp$`)` and C' as `(if (= [−] 1)` $\perp$ `0)`. Since $C[E]$ evaluates to 0 under call-by-name and diverges under call-by-value, a correct evaluation of E must return 1. But since $C'[E]$ evaluates to 0 under call-by-value and diverges under call-by-name, a correct evaluation of E must return 2. How can the answer be both 1 and 2?[3]

When faced with the seeming paradox of such a dilemma, the best place to go is right between the horns: a correct reduction strategy must return *all* answers produced by *any* non-divergent evaluation. (This alternative suggests yet new ways of implementing McCarthy's `amb` operator and nondeterministic computation.) Returning both 1 and 2 in the evaluation of E is *not* contradictory: it merely amounts to supplying both answers to a single *shared* continuation.

Technical contributions: The efficiency of optimal reduction is based on the incremental propagation of sharing nodes. Implementations of optimal reduction based on linear logic, as proposed by Gonthier, Abadi, and Lévy, and later by Asperti, use proofnet *boxes* to coordinate these interactions. Nevertheless, the boxing strategy only permits the sharing of values. To extend this technology to languages with control operators, the key technical question is: where do we put the boxes to allow the sharing of continuations?

[1] We use the phrases *proofnets*, *sharing graphs*, and *graphs* more or less interchangeably. Any typed sharing graph that remains connected under reduction is a proofnet. The ability to discard arguments in the λIK-calculus keeps sharing graphs from being proofnets, but we ignore this technicality, which can be repaired either by introducing the mix rule, or by more elaborate constructions (e.g., [17]).

[2] We have implemented these languages using the techniques in this paper. This work will be included in a later, extended version of the manuscript.

[3] Anomalies of this sort compelled Filinski to abandon "full" symmetry in his so-called *symmetric λ-calculus* [9].

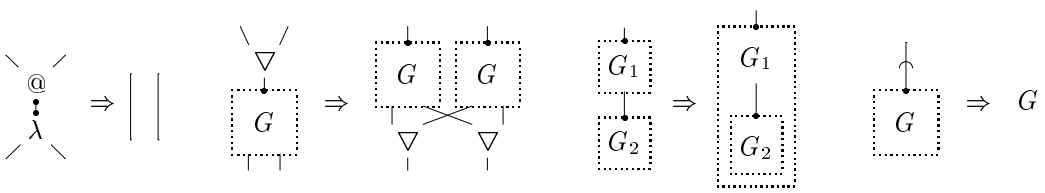

Figure 1: Global reduction rules

Our solution exploits the *continuation-passing style (CPS) transformation*. If a language with control operators can be translated into the pure λ-calculus using a CPS transformation, we can use existing technology to construct the graph of the CPS transformed term. The CPS translation of a term is more verbose than the original term, and more expensive to reduce to a normal form. We show that for a language with abortive continuations, the graphs of CPS terms can be mechanically converted back to direct style, maintaining the boxing of continuations induced by the CPS term. The transformation "rotates" principal ports of boxes so that continuations can be copied. We prove that this transformation does not change the underlying denotational semantics of the terms, as defined by the geometry of interaction. This approach can be applied to any variant of the CPS transformation, and any strategy for coding pure λ-terms as proofnets. The result is a family of possible translations into graphs, and these graphs can be optimally reduced in the sense of Lévy's labelled terms.

Traditionally, compiler optimizations have addressed sharing of expressions. The technology presented here provides a new systematic basis on which to optimize the sharing of continuations. We wish to emphasize that the Asperti-Mairson theorem on the complexity of parallel β-reduction does *not* say that such optimal reduction is inefficient—only that the parallel reduction cannot be considered a unit-cost operation [3].

In summary, all of the translations we outline possess a simple graph reduction on translated terms (cut elimination for linear logic), a consistent semantics that is preserved by reduction (geometry of interaction, via the so-called *context semantics* of Gonthier [14]), and a mechanism whereby continuations can be incrementally evaluated (optimal reduction). The situation of this technology within multiplicative-exponential linear logic ensures that the semantic characterization given is equivalent to the operational semantics of graph reduction. Viewing data types as games, and contexts (in the sense of Gonthier) as moves in a composite game, one immediately suspects that categories of games should provide the right kind of "more abstract" semantics for calculi with explicit control. Furthermore, full abstraction theorems for languages with control seem to be easily accessible, given the full completeness results that are known for linear logic.

2 Preliminaries

We briefly sketch the construction of graphs to implement λ-calculus; more details can be found elsewhere [1, 14]. Graphs are composed of wires and fixed-arity nodes, as well as *boxes*, which enclose subgraphs. The λ-calculus is encoded using apply nodes (@), lambda nodes (λ), sharing nodes ($\bigtriangledown$), weakening nodes ($\odot$), and croissants ($\frown$). A box allows a subgraph to be duplicated by a sharing node, or discarded by a weakening node. In terms of linear logic, a box transforms a graph of type A into a graph of type $!A$. When sharing is no longer required, a croissant can open the box, allowing interaction with the subgraph inside. The meaning of the other nodes should be intuitive.

One port of each node or box is designated as the *principal port*. Other ports are *auxiliary ports*. Reduction takes place when two graph constructs are connected at their principal ports. A box can also interact with another box at its auxiliary port. Global reduction rules are shown in Figure 1, where black dots indicate principal ports. Graphs can also be reduced by local reduction rules, described elsewhere [1, 14]. Local reduction of the graph of a lambda term implements Lévy's *optimal reduction* [21].

The big question in implementing λ-calculus within this framework is where to put the boxes to allow the unrestricted sharing of values. We mention two commonly used boxing disciplines; see [22] for others. The *call-by-value (CBV) coding* boxes the graph of every λ-abstraction. Correspondingly, a croissant is placed on

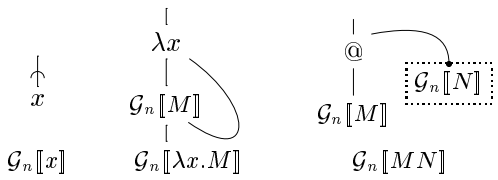

Figure 2: CBN coding of lambda terms

the function position of every apply node. The *call-by-name (CBN) coding* boxes the graph of the argument of every function application. Correspondingly, a variable reference is implemented by a croissant. Figure 2 illustrates the CBN coding of the λ-calculus, which we will use in this paper. Our results are equally applicable to the CBV coding, and to any other consistent boxing strategy.

Regardless of the boxing strategy, the constraints of linear logic typing ensure the following:

Proposition 1 *Boxes never get in the way of β reductions.*

As a consequence, optimal (local) reduction reduces any two graphs with the same arrangement of apply, lambda, and sharing nodes in the same way.

These two boxing strategies amount to Curry-Howard style embeddings of intuitionistic logic in linear logic. Let $\overline{A}$ be an interpretation of an intuitionistic implicational formula A as a linear formula. The call-by-value graph coding of a function with type $A \to B$ gives a graph of type $!(\overline{A} \multimap \overline{B})$ (box the whole function), and the call-by-name graph coding gives a graph with type $!\overline{A} \multimap \overline{B}$ (box the argument of the function). We suggest boxing schemes that yield other embeddings.

Asperti has proposed some optimizations to these boxing strategies [1]. The simplest is the application of the following rewrite rule to the translation of a lambda term:

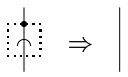

We apply this optimization, without comment, throughout the paper.

Implementing control operators: The λ-calculus codings that we have described encode a function application by connecting the right auxiliary port of an apply node to the graph of the argument, and the left auxiliary port to the context of the application, i.e., to the continuation. Correspondingly, the right auxiliary port of a lambda node leads to the argument, and the left auxiliary port to the continuation. Thus, the graph structure makes the argument and the continuation equally accessible. Because the λ-calculus can only express the sharing of arguments, via parameter binding, the boxing strategies only ensure that the value of the argument is enclosed in a box. Control operators such as Scheme's `call/cc`, however, introduce the possibility to name, and thus duplicate and discard, the continuation.

3 Continuations in the lambda calculus

We now derive a family of graph encodings for terms in the lambda calculus with `call/cc` from encodings of the corresponding CPS terms. For pure lambda terms, this approach produces graphs with the same arrangement of lambda, apply, and sharing nodes as previous translations, and thus such graphs reduce to normal form in the optimal number of beta steps, as given by Lévy's specification.

3.1 The CPS transformation

Continuation-passing style (CPS) is a style of programming in which the continuation at each point is represented explicitly as a function. Because the continuation function makes explicit the remaining computation at the current program point, a CPS term necessarily specifies an evaluation order. A lambda calculus program can be converted to CPS automatically using a CPS transformation. Furthermore, the control operator

$$
\begin{array}{rcl}
\mathcal{C}_v[\![x]\!] & \equiv & \lambda k.kx \\
\mathcal{C}_v[\![\lambda x.M]\!] & \equiv & \lambda k.k(\lambda x.\lambda k.\mathcal{C}_v[\![M]\!]k) \\
\mathcal{C}_v[\![MN]\!] & \equiv & \lambda k.\mathcal{C}_v[\![M]\!](\lambda v.\mathcal{C}_v[\![N]\!](\lambda w.vwk)) \\
\mathcal{C}_v[\![\texttt{call/cc}]\!] & \equiv & \lambda k.k(\lambda f.\lambda k.f(\lambda v.\lambda c.kv)k)
\end{array}
$$

Figure 3: CBV CPS transformation of the lambda calculus, including `call/cc`

$$
\begin{array}{rcl}
\mathcal{C}_n[\![x]\!] & \equiv & \lambda k.xk \\
\mathcal{C}_n[\![\lambda x.M]\!] & \equiv & \lambda k.k(\lambda x.\lambda k.\mathcal{C}_n[\![M]\!]k) \\
\mathcal{C}_n[\![MN]\!] & \equiv & \lambda k.\mathcal{C}_n[\![M]\!](\lambda v.v(\mathcal{C}_n[\![N]\!])k) \\
\mathcal{C}_n[\![\texttt{call/cc}]\!] & \equiv & \lambda k.k(\lambda f.\lambda k.f(\lambda a.a(\lambda q.q(\lambda v.\lambda c.vk))k))
\end{array}
$$

Figure 4: CBN CPS transformation of the lambda calculus, including `call/cc`

`call/cc` can be translated into CPS. Typically a CPS transformation encodes a CBV or CBN evaluation order, however any consistent mixture is possible [18].

Plotkin's CBV and CBN CPS transformations, extended with the translation of `call/cc`, are shown in Figures 3 and 4, respectively [27]. For typed terms, these transformations induce a corresponding transformation on types. Define $\alpha^* \equiv \alpha$ for any base type α (including $\bot$) and $(\alpha \to \beta)^* \equiv \alpha^* \to \neg\neg\beta^*$ where $\neg\tau \equiv \tau \to \bot$; then the CBV CPS transformation maps a derivation $\{x : \sigma \in \Gamma\} \vdash E : \tau$ to $\{x : \sigma^* \mid x : \sigma \in \Gamma\} \vdash \mathcal{C}_v[\![E]\!] : \neg\neg\tau^*$. Similarly, define $\alpha^\star \equiv \alpha$, and $(\alpha \to \beta)^\star \equiv \neg\neg\alpha^\star \to \neg\neg\beta^\star$; then the CBN CPS transformation maps the same derivation to $\{x : \neg\neg\sigma^\star \mid x : \sigma \in \Gamma\} \vdash \mathcal{C}_n[\![E]\!] : \neg\neg\tau^\star$.

Replacing the lambda terms produced by a CPS transformation by the corresponding graphs gives a translation of terms into graphs in which the continuation is accessible as a sharable value. Figure 5 presents the graph translation corresponding to the CBV CPS transformations. We have used the CBN boxing strategy.[4] The translation corresponding to the CBN CPS transformation is similar. Only $\bot$ types are indicated.

This implementation strategy, while straightforward, is not satisfactory. The CPS transformation introduces lambda and apply nodes that are not part of the original term. Thus, optimal reduction of the resulting graph is not a strategy for reducing the original term using the minimal number of β steps. Indeed, the number of β steps is affected by the CPS transformation chosen. Furthermore, the graph translation does not exploit the symmetry between the left side of a lambda or apply node, which connects to the continuation of a function application, and the right side, which connects to the argument. The CPS encoding does, however, produce a graph in which continuations are consistently boxed. Thus, we would like a graph translation that produces the same arrangement of lambda and apply nodes for pure lambda terms as the translations defined in Section 2, while retaining the boxing of continuations suggested by the CPS transformation.

3.2 The DS transformation on graphs

Essentially, we would like to eliminate the lambda and apply nodes that construct and manipulate continuations. To simplify the graph, we exploit the $\bot$ return type of every continuation and continuation abstraction. In a CPS program, we are not interested in the result of type $\bot$, but instead in the value passed to the initial continuation. We can show that the computation of a value of a non-$\bot$ type cannot depend on an edge transmitting a value of $\bot$ type. The simple conclusion is to remove all of these edges. Because this transformation eliminates continuations from CPS graphs, we refer to it as the *direct-style (DS) transformation.*

[4] With the Asperti optimizations, the CBN boxing strategy uses fewer boxes and croissants to encode CPS terms than the CBV translation. Our results can be adapted to the CBV translation, or to any consistent combination of these translations. Different translations induce different embeddings into linear logic.

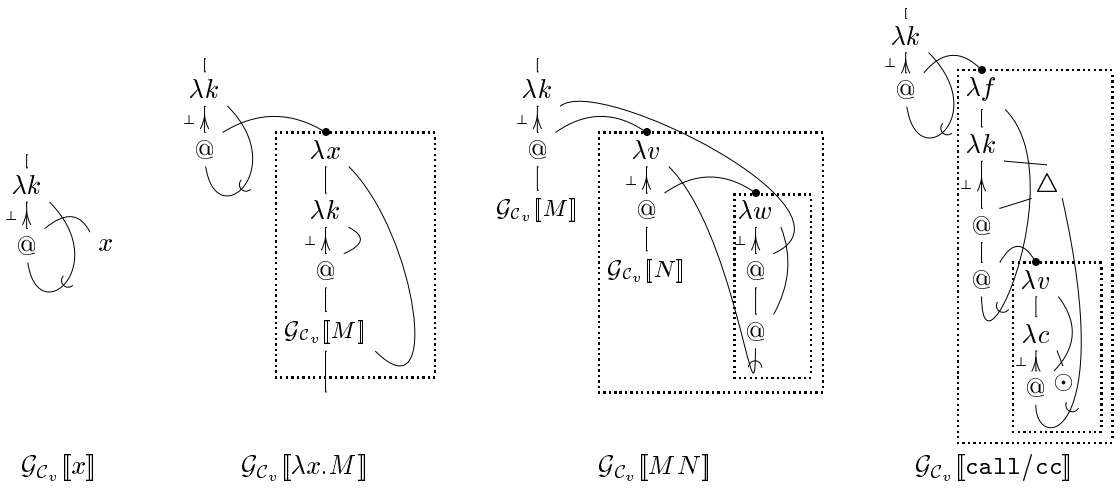

Figure 5: Graph translation based on the CBV CPS transformation

Figure 6: DS graph transformation rules

Removing edges from the graph of course affects the nodes incident upon these edges. Figure 6 shows transformation rules sufficient to treat CPS terms. $\perp$-typed edges in other positions are treated similarly.

While the graph codings of Section 2 were inspired by embeddings of intuitionistic logic into linear logic, they can equally well implement untyped terms. Here, however, we do require that $\perp$-typed values are used consistently.

The results of applying the DS transformation to the graph translations based on the CBV and CBN CPS transformations are shown in Figures 7 and 8, respectively. We refer to these translations as the $\text{CBV}_{\text{CPS/N}}$ and $\text{CBN}_{\text{CPS/N}}$ translations, respectively. Both achieve our goals: The arrangement of apply and lambda nodes in the translation of a pure lambda term is identical to that of the codings presented in Section 2, and continuations are boxed, allowing them to be duplicated or discarded.

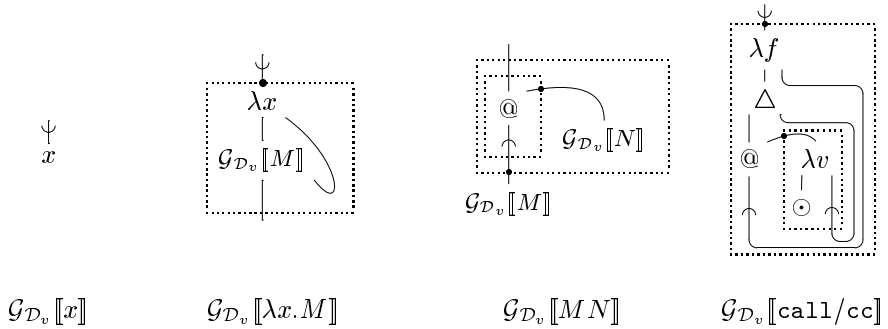

Figure 7: DS graph translation derived from the CBV CPS transformation

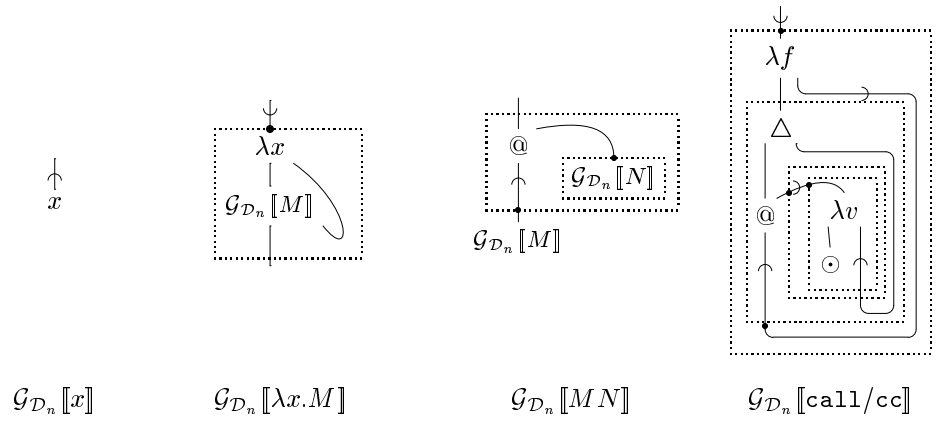

Figure 8: DS graph translation derived from the CBN CPS transformation

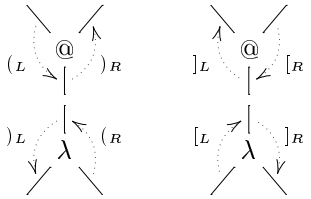

Figure 9: Edge labels

3.3 Correctness of the DS transformation

A graph can be viewed as a set of apply and lambda nodes connected by edges that may contain sharing information, as controlled by sharing nodes, croissants, and box boundaries. Because the DS transformation only modifies (i.e. eliminates) apply and lambda nodes, it does not directly affect this sharing information. Consider the path between two edges that are ultimately connected by reduction. We prove that the DS transformation maintains the way in which such a path traverses apply and lambda nodes. Because the arrangement of the other nodes is not modified by the DS transformation, they continue to behave in the same way as well.

Throughout, we assume the graph is simply typed.

Definition 1 *A* path *is a directed sequence of connected edges, labelled as shown in Figure 9, where for each node traversed on the path, the two edges incident on the node connect respectively to the principal port, and to an auxiliary port.*

Definition 2 *A* well-balanced path *is a path whose label is described by the following grammar, where* 1 *is the empty word:*

$$B \quad ::= \quad 1 \mid (_L B)_L \mid (_R B)_R \mid [_L B]_L \mid [_R B]_R \mid BB$$

An unbalanced path *is path that is not well-balanced and whose label is a subword of the labels derivable from* B*.*

This definition of well-balanced path generalizes that of Asperti and Laneve [2] to include paths that cannot occur in the translation of an ordinary lambda term, but can occur in the image of the DS transformation. Note that some paths are neither well-balanced nor unbalanced, for example a path that enters an

apply node on the left and immediately exits the next lambda node on the right, with label $(_L]_R$. Unbalanced paths describe correct information flow (in the sense of the geometry of interaction) in a well-typed graph, but cannot reduce to form a beta redex.

Proposition 2 *Let p be a path that is converted to an edge by the DS transformation. Then, the label of p is a subword of the labels derivable from B (i.e. p is either balanced or unbalanced).*

The four kinds of left (right) parentheses only serve to eliminate paths that do not describe correct information flow in a well-typed graph. Thus, in discussing balanced and unbalanced paths, we can simply use left and right parentheses.

We can show the following by induction on the structure of a path in a well-typed graph:

Proposition 3 *A matching pair of parentheses in the label of a path corresponds to an apply node and a lambda node of the same type.*

Proposition 4 *In a well-balanced path, the first and last edges have identical types.*

Because of this proposition, we can unambiguously refer to the type of a path as the type of its first edge.

Proposition 5 *In a well-balanced path not of type $\perp$, no edge has type $\perp$.*

We can apply the DS transformation to a path p by applying the DS transformation to every node along the path. We now show that the DS transformation maps a well-balanced path into a well-balanced path.

Lemma 1 *Let p be a well-balanced path of non-$\perp$ type. Then $\mathcal{D}(p)$ is a well-balanced path.*

Proof: By Proposition 5, the path remains connected after the DS transformation. We proceed by induction on the label of p; the only interesting case is a path p with a label of the form (B). In this case, the subpath with the label B connects the principal ports of a lambda and apply node, and the outer pair of parentheses represents traversal of these nodes. By Proposition 3, the apply and lambda node have the same type, and are thus both either preserved or eliminated by the DS transformation. The subpath labelled B continues to be well-balanced by induction. □

Similarly, one expects that if p is an unbalanced path, then $\mathcal{D}(p)$ is either not a path or is unbalanced, but this is not true. Consider a path with label $]_R$ (enter a lambda node at its principal port, exiting to the right). If the lambda node identifies a continuation, it is eliminated by the DS transformation, and the unbalanced path is converted to a balanced path of weight 1.

Lemma 2 *Let p be an unbalanced path connecting the principal ports of an apply node a and a lambda node l in a well-typed graph G, where a and l are not removed by the DS transformation. Then $\mathcal{D}(p)$ is an unbalanced path.*

Proof: Let L be the label of p. We can extend p to an unbalanced path with label $L' = ({}^a L)^l$, where $({}^a$ and $)^l$ indicate traversals of the nodes a and l. Since p is unbalanced, either $L = B)^{\overline{l}} X$ or $L = X({}^{\overline{a}} B$. In the first case, $p = bq$, where b is well-balanced with label B, and q is unbalanced with label $)^{\overline{l}} X$. Then, $L' = ({}^a B)^{\overline{l}} X)^l$. Well-typing ensures that the lambda node corresponding to $)^{\overline{l}}$ is not removed by the DS transformation (Proposition 3), so $\mathcal{D}(q)$ is unbalanced. Since $\mathcal{D}(p) = \mathcal{D}(b)\mathcal{D}(q)$, and $\mathcal{D}(b)$ is well-balanced (Lemma 1), $\mathcal{D}(p)$ is unbalanced. The case where $L = X({}^{\overline{a}} B$ is similar. □

Lemma 3 *In an unbalanced path p, either the first or last edge has arrow type.*

Proof: Let L be the label of p. Since p is unbalanced, either $L = B)X$ or $L = X(B$. Consider $B)X$. By Proposition 5, the type of this path is the same as type the of $)X$. A path of this form begins with the edge incident on the root of a lambda node or the function port of an apply node. Such an edge has arrow type. The case for $X(B$ is similar; here the last edge has arrow type. □

Theorem 1 (Soundness) *Let G be the graph of a pure λ-term. Then the diagram*

$$\begin{array}{ccc} G & \longrightarrow & G' \\ \downarrow{\scriptstyle \mathcal{D}} & & \downarrow{\scriptstyle \mathcal{D}} \\ \mathcal{D}(G) & \longrightarrow & \mathcal{D}(G') \equiv G'' \end{array}$$

commutes, where:

1. *(The top path can be simulated by the bottom path): If G reduces to G' by global reductions, then $\mathcal{D}(G)$ reduces to $\mathcal{D}(G')$ by global reductions.*
2. *(The bottom path can be simulated by the top path): If $\mathcal{D}(G)$ reduces to G'' by global reductions, then there is some G' such that G reduces to G' by global reductions, and $G'' \equiv \mathcal{D}(G')$.*

Proof: To prove that the top path can be simulated by the bottom path, we use induction on the number of reduction steps in the top path. Observe that the effect of the DS transformation on a node is completely local, determined only by the types of edges incident on the node. Thus, the effect of the DS transformation on source nodes not affected by the reduction step is the same in both the source graph and the reduced graph; we need only consider the effect of the DS transformation on the nodes involved in the reduction step. Consider each possible redex in the source graph:

- A β-redex with function type $A \to B$, where $B \neq \bot$. Since the redex is not affected by the DS transformation, we can perform the same reduction step in $\mathcal{D}(G)$, and the resulting graph has the form $\mathcal{D}(G')$.
- A β-redex with function type $A \to \bot$. Reducing a beta redex creates an edge of the argument type (connecting the argument to the occurrence of the parameter) and an edge of the return type (connecting the result of the body to the context of the application). Thus, reducing a redex with function type $A \to \bot$ creates an edge of type A and an edge of type $\bot$ in G'. The edge of type $\bot$ is then eliminated by the DS transformation. Applying the DS transformation to G directly also eliminates the $\bot$-typed edges and connects the A-typed edges, thus having the same effect as beta reduction.
- Box duplication, absorption, croissant-box interaction. In a λ-term, no box has an auxiliary port of type $\bot$; in the source terms created by CPS conversion, no box has a root of type $\bot$. Thus a box duplication, box absorption, or box opening in the top path is simulated by the identical operation in the bottom path; these graph operations commute with the DS transformation.

To show that reductions in the bottom path can be simulated by reductions in the top path, we proceed by induction on the number of reduction steps from $\mathcal{D}(G)$ to G''. Since the DS transformation only converts λ- and apply nodes to edges, an edge $\mathcal{D}(p)$ between two interaction ports in $\mathcal{D}(G)$ corresponds to a path p in G consisting of a sequence of lambda and apply nodes. If p is not an edge, we show that p must β-reduce to one; thus, a single reduction step in $\mathcal{D}(G)$ is simulated by a sequence of β-steps in G, followed by the same reduction step as performed in $\mathcal{D}(G)$. Because p consists of only λ- and apply nodes, it suffices to show that p is a balanced path. Consider each possible redex in $\mathcal{D}(G)$:

- A β-redex. Because $\mathcal{D}(p)$ is a single edge, it is a balanced path. $\mathcal{D}(p)$ also connects an apply node and a lambda node, not removed by the DS transformation. By Lemma 2, p cannot be an unbalanced path. Thus, by Proposition 2, p must be a balanced path.

- Box duplication, box absorption, croissant-box interaction. In all of these cases, the type of the edge $\mathcal{D}(p)$ between the interaction ports is $!A$. Because the boxes, croissants, and sharing nodes are not affected by the DS transformation, the first and last edges of p must also have type $!A$. By Lemma 3, p cannot be unbalanced. Thus, by Proposition 2, p must be balanced.

□

3.4 Embeddings of classical logic

Each of our proofnet implementations implicitly encodes, via an extended Curry-Howard correspondence, an embedding of classical logic in multiplicative-exponential linear logic (MELL). This family of encodings results from the mix-and-match of standard double-negation embeddings of classical logic into intuitionistic logic, composed with embeddings of intuitionistic logic into MELL. We discuss these constructions for minimal implicational logic with $\neg\neg$-elimination.

Let $[\alpha \to \beta] \equiv !{[\alpha]} \multimap [\beta]$ be the Girard translation of intuitionistic implication in linear logic; similarly, let $\langle \alpha \to \beta \rangle \equiv !(\langle \alpha \rangle \multimap \langle \beta \rangle)$ be the Gonthier-Abadi-Lévy translation (see [14]). By standard linear logic identities, $[\neg\neg\tau] = ?![\tau]$ and $\langle \neg\neg\tau \rangle = !?\langle \tau \rangle$, where ! and ? are the (dual) exponential modalities.

What does this mean in terms of graph reduction? When a subgraph G with root typed $?!\alpha$ is substituted into a sharable context C, the ? marks a croissant at the root of G that breaks the box around C, which then shares the value of type $!\alpha$. Dually, if G has type $!?\beta$, the context has type $?!(\beta)^{\perp}$, and the protocol for box opening and sharing reverses the role of context and value.

Recall $(\alpha \to \beta)^* \equiv \alpha^* \to \neg\neg\beta^*$ in the translation of $\alpha \to \beta$ by the CBV CPS transformation, and $(\alpha \to \beta)^\star \equiv \neg\neg\alpha^\star \to \neg\neg\beta^\star$ is the translation induced by the CBN CPS transformation. The CPS translations of a function (classical proof) $E : \alpha \to \beta$ result in a function of type $\neg\neg(\alpha \to \beta)^*$ or $\neg\neg(\alpha \to \beta)^\star$; we then have

$$\begin{aligned}
[\neg\neg(\alpha \to \beta)^*] &= ?!(!\,[\alpha^*] \multimap ?!\,[\beta^*]) \\
[\neg\neg(\alpha \to \beta)^\star] &= ?!(!?!\,[\alpha^\star] \multimap ?!\,[\beta^\star]) \\
\langle \neg\neg(\alpha \to \beta)^* \rangle &= !?!(\langle \alpha^* \rangle \multimap !?\langle \beta^* \rangle) \\
\langle \neg\neg(\alpha \to \beta)^\star \rangle &= !?!(!?\langle \alpha^\star \rangle \multimap !?\langle \beta^\star \rangle)
\end{aligned}$$

Other variants of this mix-and-match style are possible. Fewer modalities means fewer boxes and a greater implementation efficiency. To translate classical implicational proofs, we also need to consider $\neg\neg$-elimination, treated by a cousin of `call/cc`, Felleisen's C operator, where classical proof $M : \neg\neg\tau$ implies $\mathsf{C}M : \tau$. We then have (see [24]):

$$\begin{aligned}
\mathcal{C}_v[\![\mathsf{C}M]\!] &\equiv \lambda k.\mathcal{C}_v[\![M]\!](\lambda m.m(\lambda v.\lambda h.kv)(\lambda x.x)) : \neg\neg\tau^* \\
\mathcal{C}_n[\![\mathsf{C}M]\!] &\equiv \lambda k.\mathcal{C}_n[\![M]\!](\lambda m.m(\lambda g.g(\lambda v.\lambda h.vk))(\lambda x.x)) : \neg\neg\tau^\star
\end{aligned}$$

where

$$\begin{aligned}
\mathcal{C}_v[\![M]\!] &: \alpha_v \equiv \neg\neg((\tau^* \to \neg\neg\perp) \to \neg\neg\perp) \\
\mathcal{C}_n[\![M]\!] &: \alpha_n \equiv \neg\neg(\neg\neg(\neg\neg\tau^\star \to \neg\neg\perp) \to \neg\neg\perp)
\end{aligned}$$

How does graph reduction implement $\neg\neg$-elimination? Consider the CBV CPS translation composed with the Girard translation. The DS transformation of $\mathcal{C}_v[\![\mathsf{C}M]\!]$ appears in Figure 10; the argument M has type $[\alpha_v]$. Continuation k is packaged as a function (the subgraph rooted at λv) which when applied, connects a null value (the plug to the left of node λv) to the current continuation (the other plug in Figure 10, of type $?!\perp$), and links the argument (of type $?![\tau^*]$) to k.

4 Related work

We consider three areas of related work: other CPS transformations, other approaches to converting CPS programs back to direct style, and other connections between control operators and linear logic.

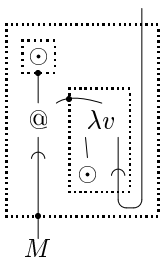

$\mathcal{G}_{\mathcal{D}_v}[\![\mathsf{C}M]\!]$

Figure 10: DS graph translation of $\mathsf{C}M$ derived from the CBV CPS transformation

Optimizing the CPS transformation: The Plotkin CPS transformations create many "administrative" redexes involving the application of a continuation or continuation abstraction [27]. Our DS transformation converts administrative redexes into box-croissant redexes, adding a bureaucratic cost to optimal reduction [1, 23]. The optimized CPS transformations in [8, 28] could then lead to more efficient implementations.

Converting CPS programs back to direct style: Danvy first investigated the problem of converting a CPS program back to a direct-style [6], later extended with Lawall. The conversions were only on terms that could be output by CPS transformation. Our DS transformation also relies on a uniform, but weaker property: values of type $\bot$ must occur consistently, and do not contribute to the final result. At the extreme, our DS transformation is simply the identity on the graphs of DS terms.

Relating languages with control operators to linear logic: Nishizaki also investigated encodings of lambda calculus plus `call/cc` in proofnets [25]. He showed that normalization of these proofnets is complete with respect to normalization in the term language. He began by adding modalities in an ad hoc manner (induced mechanically by our $\mathrm{CBV}_{\mathrm{CPS/N}}$ translation) to the type $!A \multimap B$, allowing sharing of both values and continuations. His more complex proofnets are an optimized variant of our $\mathrm{CBV}_{\mathrm{CPS/N}}$ translation, eliminating some box croissant interactions corresponding to administrative redexes. Because Nishizaki derived a translation from the types rather than from the semantics, he had to prove that the resulting graphs model the semantics of the language. The correctness of our approach relies only on the correctness of the CPS transformation, and on the correctness of the DS transformation on graphs, which is independent of the language being implemented.

In a sequel to his earlier work on symmetric lambda calculus, Filinski used linear logic as a tool for understanding continuations [10]. Some of the linear types he proposed for continuations appear in our codings, the most common being $?!\alpha$, resulting from the DS transformation of a graph with type $(\alpha \multimap \bot) \multimap \bot$. Griffin, and later Murthy, showed the relation between so-called $\neg\neg$-embeddings of classical logic in intuitionistic logic, and the implementation of control operators [16, 24]. In particular, they showed how varieties of the CPS transformation provide the constructive content of such embeddings. We further translate such terms into proofnets in direct style, eliminating the administrative redexes. The result is a family of constructive embeddings of classical logic into linear logic.

5 Future work

Since the continuation created by `call/cc` is abortive, a term containing `call/cc` can reduce to different normal forms; its CPS counterpart, like all pure λ-terms, has only one normal form. Because DS transformation produces boxes according to the CPS transformation providing its input, it should be possible to identify, among the shared normal forms it can return, the answer that would have been produced by the CPS-converted input. For example, the CBV CPS conversion of `call/cc` $\lambda k.((\lambda x.1)(k2))$ reduces to $\lambda k.(\lambda c.k2)(\lambda r.k1)$. Suppose evaluation begins with the expression $k1$. If k is an abortive continuation, the entire context is discarded, and the result is 1. If k is an ordinary function, the result of $k1$ is returned to the enclosing context, where it is eventually discarded by the outer beta redex. Graph reduction gives

both answers—but we should be able to distinguish, by box placement, which is the CBV answer. We leave further analyses of these observations to future work.

Efficiently managing the reified continuation is a significant problem in implementing languages with control operators [5, 19]. Proofnet implementations suggest the possibility of evaluating programs containing control operators using optimal reduction, with a minimal copying of shared values. Nevertheless, we are exchanging the savings of optimal reduction for the overhead of box management. Further experiments are needed to understand whether the exchange is cost-effective, and if it can be further optimized by better box technology.

Our proofnet technology might be extended to languages with functional control operators, such as Danvy and Filinski's `shift` and `reset`, and Sitaram and Felleisen's `control` and `prompt` [7, 30]. While `shift` and `reset` are defined in terms of a CPS transformation, they do not have return type $\perp$; the DS transformation is then inapplicable. Both `shift` and `reset`, and `control` and `prompt` can be defined in terms of `call/cc` and a reference cell [11, 30]. Bawden has shown how to implement reference cells using sharing graphs [4], so this strategy may still lead to an interesting proofnet implementation.

6 Conclusions

We have shown how to implement various languages with explicit control using graph reduction, where the structure of the graphs are proofnets from linear logic. The principal technical difficulty in such codings is the location of boxes, which indicate what computations can be shared. Rather than specifying a fixed scheme for locating boxes, we have introduced a very general methodology based on the CPS transform. Different versions of CPS, followed by our DS transform on graphs, produce a wide range of consistent schemes for locating boxes in proofnets. The noble art of linear decorating, to repeat the phrase of Schellinx [29], has been replaced by a factory.

The theoretical foundation of our implementation technology means that we have a consistent semantics provided by the geometry of interaction, and a means of incrementally evaluating continuations via optimal evaluation. The codings may clarify full abstraction theorems for languages with explicit control, given the full completeness results that are known for linear logic. But the genuine progress reflected in the presented techniques is the technology transfer of logic and proofnets to the mundane algorithmics of implementation. The pragmatics of double negation in logic, for example, is just *packaging:* the boxing of sharable data so that they can interact with each other. Further, implementation improvements amount to a better understanding of where to put boxes. A generation of compiler writers has spent considerable effort optimizing the efficiency of sharing expressions. We have presented a systematic basis on which to optimize the sharing of continuations, providing new territory for similar efficiency improvements.

Acknowledgments. We thank Alan Bawden and Olivier Danvy for commenting on a draft of this paper.

References

[1] A. Asperti. $\delta \circ ! \epsilon = 1$: Optimizing optimal λ-calculus implementations. In *Rewriting Techniques and Applications*, pages 102–116, Kaiserslautern, Germany, 1995.

[2] A. Asperti and C. Laneve. Paths, computations and labels in the lambda-calculus. In *Rewriting Techniques and Applications, 5th International Conference*, volume 690, pages 152–167. Lecture Notes in Computer Science, 1993.

[3] A. Asperti and H.G. Mairson. Parallel beta reduction is not elementary recursive. In *Conference Record of POPL '98: The 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 303–315, San Diego, California, 19–21 January 1998.

[4] A. Bawden. Implementing distributed systems using linear naming. TR 1627, MIT AI Lab, March 1993. (PhD thesis, originally titled *Linear Graph Reduction: Confronting the Cost of Naming*. MIT, May 1992).

[5] C. Bruggeman, O. Waddell, and R.K. Dybvig. Representing control in the presence of one-shot continuations. In *Proceedings of the ACM SIGPLAN'96 Conference on Programming Language Design and Implementation*, volume 31(5), pages 99–107, Philadephia, Pennsylvania, May 1996. SIGPLAN Notices.

[6] O. Danvy. Back to direct style. *Science of Computer Programming*, 22(3):183–195, 1994.

[7] O. Danvy and A. Filinski. Abstracting control. In *Proceeding of the 1990 ACM Conference on Lisp and Functional Programming*, pages 151–160, Nice, France, June 1990. ACM Press.

[8] O. Danvy and A. Filinski. Representing control: A study of the CPS transformation. *Mathematical Structures in Computer Science*, 4:361–391, 1992.

[9] A. Filinski. Declarative continuations and categorical duality. Technical Report 89/11, University of Copenhagen, 1989. Masters Thesis.

[10] A. Filinski. Linear continuations. In *Conference Record of the Nineteenth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 27–38, Albuquerque, New Mexico, January 1992.

[11] A. Filinski. Representing monads. In *Conference Record of POPL '94: 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 446–457, Portland, Oregon, January 1994.

[12] J.-Y. Girard. Linear logic. *Theoretical Computer Science*, 46:1–102, 1986.

[13] J.-Y. Girard. Linear logic: Its syntax and semantics. In J.-Y. Girard, Y. Lafont, and L. Regnier, editors, *Advances in Linear Logic*, pages 1–42. Cambridge University Press, 1995. Proceedings of the Workshop on Linear Logic, Ithaca, New York, June 1993.

[14] G. Gonthier, M. Abadi, and J.-J. Lévy. The geometry of optimal lambda reduction. In *Conference record of the Nineteenth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 15–26, Albuquerque, New Mexico, January 1992.

[15] G. Gonthier, M. Abadi, and J.-J. Lévy. Linear logic without boxes. In *Proceedings of the Seventh Annual Symposium on Logic in Computer Science*, pages 223–234, June 1992.

[16] T. Griffin. A formulae-as-types notion of control. In *Conference Record of the Seventeenth Annual ACM Symposium on Principles of Programming Languages*, pages 47–58, San Francisco, California, January 1990.

[17] S. Guerrini, S. Martini, and A. Masini. Proof nets, garbage, and computations. In *Typed Lambda Calculi and Applications*, volume 1210, pages 181–195, Nancy, France, April 1997. Lecture Notes in Computer Science.

[18] J. Hatcliff and O. Danvy. A generic account of continuation-passing styles. In *Conference Record of POPL '94: 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 458–471, Portland, Oregon, January 1994.

[19] R. Hieb, R.K. Dybvig, and C. Bruggeman. Representing control in the presence of first-class continuations. In *Proceedings of the ACM SIGPLAN'90 Conference on Programming Language Design and Implementation*, volume 25(6), pages 66–77, White Plains, New York, June 1990. SIGPLAN Notices.

[20] J. Lamping. An algorithm for optimal lambda calculus reduction. In *Conference Record of the Seventeenth Annual ACM Symposium on Principles of Programming Languages*, pages 16–30, San Francisco, California, January 1990.

[21] J.-J. Lévy. *Optimal Reductions in the Lambda-Calculus*, pages 159–191. Academic Press, 1980.

[22] I. Mackie. *The Geometry of Implementation*. PhD thesis, University of London, September 1994.

[23] I. Mackie. YALE: Yet another lambda evaluator based on interaction nets. In *Proceedings of the ACM SIGPLAN International Conference on Functional Programming*, pages 117–128, Baltimore, Maryland, September 1998.

[24] C.R. Murthy. *Extracting Constructive Content from Classical Proofs*. PhD thesis, Cornell University, August 1990.

[25] S. Nishizaki. Programs with continuations and linear logic. *Science of Computer Programming*, 21(2):165–190, 1993.

[26] M. Parigot. Lambda-mu-calculus: an algorithmic interpretation of classical natural deduction. In A. Voronkov, editor, *Logic Programming and Automated Reasoning: International Conference LPAR '92 Proceedings, St. Petersburg, Russia*, pages 190–201, Berlin, DE, 1992. Springer-Verlag.

[27] G. D. Plotkin. Call-by-name, call-by-value and the lambda calculus. *Theoretical Computer Science*, 1:125–159, 1975.

[28] A. Sabry and M. Felleisen. Reasoning about programs in continuation-passing style. *Lisp and Symbolic Computation*, 6(3/4):289–360, November 1993.

[29] H. Schellinx. *The Noble Art of Linear Decorating*. PhD thesis, Institute for Logic, Language and Computation, University of Amsterdam, 1994.

[30] D. Sitaram and M. Felleisen. Control delimiters and their hierarchies. *Lisp and Symbolic Computation*, 3(1):67–99, January 1990.