The Verse Calculus: A Core Calculus for Deterministic Functional Logic Programming (Extended Version)

LENNART AUGUSTSSON, Epic Games, Sweden
JOACHIM BREITNER, Unaffiliated, Germany
KOEN CLAESSEN, Epic Games, Sweden
RANJIT JHALA, Epic Games, USA
SIMON PEYTON JONES, Epic Games, United Kingdom
OLIN SHIVERS, Epic Games, USA
GUY L. STEELE JR., Oracle Labs, USA
TIM SWEENEY, Epic Games, USA

Functional logic languages have a rich literature, but it is tricky to give them a satisfying semantics. In this paper we describe the Verse calculus, $\mathcal{VC}$, a new core calculus for deterministic functional logic programming. Our main contribution is to equip $\mathcal{VC}$ with a small-step rewrite semantics, so that we can reason about a $\mathcal{VC}$ program in the same way as one does with lambda calculus; that is, by applying successive rewrites to it. We also show that the rewrite system is confluent for well-behaved terms.

This is an extended version (with appendices) of the paper in the Proceedings of the International Conference on Functional Programming (2023).

CCS Concepts: • **Theory of computation** → **Equational logic and rewriting**; *Proof theory*; **Rewrite systems**; *Grammars and context-free languages*; • **Software and its engineering** → **Syntax**; **Semantics**; **Functional languages**; **Constraint and logic languages**; **Multiparadigm languages**.

Additional Key Words and Phrases: choice operator, confluence, declarative programming, evaluation strategy, even/odd problem, functional programming, lambda calculus, lenient evaluation, logic programming, logical variables, normal forms, rewrite rules, skew confluence, substitution, unification, Verse calculus, Verse language

1 INTRODUCTION

Functional logic programming languages add expressiveness to functional programming by introducing logical variables, equality constraints among those variables, and choice to allow multiple

Authors' addresses: Lennart Augustsson, Epic Games, Sweden, lennart.augustsson@epicgames.com; Joachim Breitner, Unaffiliated, Germany, mail@joachim-breitner.de; Koen Claessen, Epic Games, Sweden, koen.claessen@epicgames.com; Ranjit Jhala, Epic Games, USA, ranjit.jhala@epicgames.com; Simon Peyton Jones, Epic Games, United Kingdom, simonpj@epicgames.com; Olin Shivers, Epic Games, USA, olin.shivers@epicgames.com; Guy L. Steele Jr., Oracle Labs, USA, guy.steele@oracle.com; Tim Sweeney, Epic Games, USA, tim.sweeney@epicgames.com.

alternatives to be explored. Here is a tiny example:

$$\exists x\,y\,z.\ x = \langle y, 3\rangle;\ x = \langle 2, z\rangle;\ y$$

This expression introduces three logical (or existential) variables x, y, z, constrains them with two equalities ($x = \langle y, 3\rangle$ and $x = \langle 2, z\rangle$), and finally returns y. The only solution to the two equalities is $y = 2$, $z = 3$, and $x = \langle 2, 3\rangle$; so the result of the whole expression is 2.

Functional logic programming has a long history and a rich literature [Antoy and Hanus 2010]. But it is somewhat tricky for programmers to *reason* about functional logic programs: they must think about logical variables, needed narrowing, unification, and the like. This contrasts with functional programming, where one can say "just apply rewrite rules, such as β-reduction, let-inlining, and case-of-known-constructor." We therefore seek *a precise expression of functional logic programming as a term-rewriting system*, to give us both a formal semantics (via small-step reductions), and a powerful set of equivalences that programmers can use to reason about their programs, and that compilers can use to optimize them.

We make two main contributions. **Our first contribution** is a new core calculus for functional logic programming, the Verse calculus or $\mathcal{VC}$ for short (Section 2). Like any functional logic language, $\mathcal{VC}$ supports logical variables, equalities, and choice, but it is distinctive in several ways:

- *Natively higher order.* $\mathcal{VC}$ directly supports higher-order functions, just like the lambda calculus. Indeed, *every lambda calculus program is a $\mathcal{VC}$ program*. In contrast, most of the functional logic literature is rooted in a first-order world, and addresses higher-order features via an encoding called defunctionalization [Reynolds 1972; Hanus 2013, 3.3].
- *Deterministic, with native encapsulation.* $\mathcal{VC}$ is *deterministic*, in the sense that when an expression yields more than one value (as is often the case in functional logic programs), those values are returned in a well-specified order. This makes it easy to solve the notoriously tricky issue [Braßel et al. 2004a,b] of how to *encapsulate* the result of a search as a data structure, using the **all** operator (see Section 2.6). It opens up a new approach to dealing with so-called "flexible" *vs.* "rigid" variables (see Section 2.5). It supports an elegant economy of concepts: for example, there is just one equality (other languages may have a suspending equality and a narrowing equality), and conditional expressions are driven by failure rather than booleans (see Section 2.5). On the other hand, it pretty much rules out *laziness* (see Section 3.6) and parallel first-come first-returned search strategies. Most other functional-logic languages (Curry [Hanus et al. 2016] is the brand leader in this design space) are non-deterministic by design; $\mathcal{VC}$ explores a different (and less well-examined) part of the design space.

Our second contribution is to equip $\mathcal{VC}$ with a *small-step term-rewriting semantics* (see Section 3). We said that the lambda calculus is a subset of $\mathcal{VC}$, so it is natural to give its semantics using rewrite rules, just as for the lambda calculus. That seems challenging, however, because logical variables and unification involve sharing and non-local communication. How can that be expressed in a rewrite system?

Happily, we can build on prior work: exactly the same difficulty arises with call-by-need in the lambda calculus. For a long time, the only semantics of call-by-need that was faithful to its sharing semantics (in which thunks are evaluated at most once) was an operational semantics that sequentially threads a global heap through execution [Launchbury 1993]. But then Ariola *et al.*, in a seminal paper, showed how to *reify the heap into the term itself*, and thereby build a rewrite system that is completely faithful to lazy evaluation [Ariola et al. 1995]. Inspired by their idea, we present a new rewrite system for functional logic programs that reifies logical variables, unification, and choice into the term itself, and replaces non-deterministic search with

a (deterministic) tree of successful results. In $\mathcal{VC}$ the choices are "laid out in space" (in the syntax of the term) rather than, as is more typical, "laid out in time" (via non-deterministic rewrites and backtracking).

Integers	k
Variables	x, y, z, f, g
Programs	$p ::= \textbf{one}\{e\}$ where $\text{fvs}(e) = \emptyset$
Expressions	$e ::= v \mid eq;\ e \mid \exists x.\ e \mid \textbf{fail} \mid e_1 \| e_2 \mid v_1\ v_2 \mid \textbf{one}\{e\} \mid \textbf{all}\{e\}$
	$eq ::= e \mid v = e$ Note: "eq" is pronounced "expression or equation"
Values	$v ::= x \mid hnf$
Head values	$hnf ::= k \mid op \mid \langle v_1, \cdots, v_n \rangle \mid \lambda x.\ e$
Primops	$op ::= \textbf{gt} \mid \textbf{add}$

Concrete syntax: "$\|$" and ";" are right-associative.

"$=$" binds more tightly than ";".

"λ" and "$\exists$" each scope as far to the right as possible.

For example, $(\lambda y.\ \exists x.\ x = 1;\ x + y)$ means $(\lambda y.\ (\exists x.\ ((x = 1);\ (x + y))))$.

Parentheses may be used freely to aid readability and override default precedence.

$\text{fvs}(e)$ means the free variables of e; in $\mathcal{VC}$, λ and $\exists$ are the only binders.

Desugaring of extended expressions

$$
\begin{array}{rcll}
e_1 + e_2 & \text{means} & \textbf{add}\langle e_1, e_2 \rangle & \\
e_1 > e_2 & \text{means} & \textbf{gt}\langle e_1, e_2 \rangle & \\
\exists x_1\ x_2 \cdots x_n.\ e & \text{means} & \exists x_1.\ \exists x_2.\ \cdots \exists x_n.\ e & \\
x := e_1;\ e_2 & \text{means} & \exists x.\ x = e_1;\ e_2 & \\
e_1\ e_2 & \text{means}^\dagger & f := e_1;\ x := e_2;\ f\ x & f, x \text{ fresh} \\
\langle e_1, \cdots, e_n \rangle & \text{means}^\dagger & x_1 := e_1;\ \cdots;\ x_n := e_n;\ \langle x_1, \cdots, x_n \rangle & x_i \text{ fresh} \\
e_1 = e_2 & \text{means}^\ddagger & x := e_1;\ x = e_2;\ x & x \text{ fresh} \\
\lambda \langle x_1, \cdots, x_n \rangle.\ e & \text{means} & \lambda p.\ \exists x_1 \cdots x_n.\ p = \langle x_1, \cdots, x_n \rangle;\ e & p \text{ fresh}, n \geqslant 0 \\
\textbf{if}\ (\exists x_1 \cdots x_n.\ e_1)\ \textbf{then}\ e_2\ \textbf{else}\ e_3 & \text{means} & (\textbf{one}\{(\exists x_1 \cdots x_n.\ e_1;\ \lambda \langle \rangle.\ e_2) \| (\lambda \langle \rangle.\ e_3)\}) \langle \rangle & \\
\end{array}
$$

†Apply this rule only if at least one of the e_i is not a value v.

‡Apply this rule only if either (i) e_1 is not a value v, or (ii) $e_1 = e_2$ is not to the left of a ";".

Fig. 1. $\mathcal{VC}$: **Syntax**

As an example of rewriting in action, the expression above can be rewritten as follows[1]:

$$
\begin{array}{lll}
& \exists x\,y\,z.\ x = \langle y, 3 \rangle;\ x = \langle 2, z \rangle;\ y & \longrightarrow \{\textsc{subst}\} & \exists x\,y\,z.\ \langle 2, z \rangle = \langle y, 3 \rangle;\ x = \langle 2, z \rangle;\ y \\
\longrightarrow \{\textsc{eqn-elim}\} & \exists y\,z.\ \langle 2, z \rangle = \langle y, 3 \rangle;\ y & \longrightarrow \{\textsc{u-tup}\} & \exists y\,z.\ 2 = y;\ z = 3;\ y \\
\longrightarrow \{\textsc{eqn-elim}\} & \exists y.\ 2 = y;\ y & \longrightarrow \{\textsc{hnf-swap}\} & \exists y.\ y = 2;\ y \\
\longrightarrow \{\textsc{subst}\} & \exists y.\ y = 2;\ 2 & \longrightarrow \{\textsc{eqn-elim}\} & 2 \\
\end{array}
$$

Rules may be applied anywhere they match, including under binders, again just like the lambda calculus. This freedom only makes sense, however, if each term ultimately reduces to a unique value, regardless of its reduction path, so we show that $\mathcal{VC}$ is *confluent*, in Section 4.

As always with a calculus, the idea is that $\mathcal{VC}$ distills the essence of (deterministic) functional logic programming. Each construct does just one thing, and $\mathcal{VC}$ cannot be made smaller without losing key features. We are working on Verse, a new general-purpose programming language, built directly on $\mathcal{VC}$; indeed, our motivation for developing $\mathcal{VC}$ is practical rather than theoretical. No single aspect of $\mathcal{VC}$ is unique, but we believe that their combination is particularly harmonious and orthogonal. We discuss design alternatives in Section 5 and related work in Section 6.

---

[1]The rule names after each arrow come from Fig. 3, to be discussed in Section 3; they are given here just for reference.

2   THE VERSE CALCULUS, INFORMALLY

We begin by presenting the Verse calculus, $\mathcal{VC}$, informally. We describe its rewrite rules precisely in Section 3. The (abstract) syntax of $\mathcal{VC}$ is given in Fig. 1. It has a very conventional sub-language that is just the lambda calculus with some built-in operations and tuples as data constructors:

- *Values.* A value v is either a variable x or a head-normal form *hnf*. In $\mathcal{VC}$, a variable counts as a value because in a functional logic language an expression may evaluate to an as-yet-unknown logical variable. A head-normal form is a conventional value: a built-in constant k, an operator *op*, a tuple, or a lambda. Our tiny calculus offers only integer constants k and two illustrative integer operators *op*, namely **gt** and **add**.
- *Expressions e* include values v, and applications $v_1\ v_2$; we will introduce the other constructs as we go. For clarity, we often write $v_1(v_2)$ rather than $v_1\ v_2$ when v_2 is not a tuple.
- A term *eq* is either an ordinary expression e, or an *equation* $v = e$; this syntax ensures that equations can only occur to the left of a "; " (Section 2.1).
- A *program*, p, contains a closed expression from which we extract one result using **one** (see Section 2.5)—unless the expression fails, in which case the program fails (Section 2.2).

The formal syntax for e allows only applications of *values* to *values*, $(v_1\ v_2)$, but the desugaring rules in Fig. 1 show how to desugar more applications $(e_1\ e_2)$. This restriction is not fundamental; it simply reduces the number of rewrite rules we need[2].

Modulo this desugaring, every lambda calculus term is a $\mathcal{VC}$ term and has the same semantics. Just like the lambda calculus, $\mathcal{VC}$ is untyped; adding a type system is an excellent goal but is the subject of another paper.

Expressions also include two other key collections of constructs: logical variables with the use of equations to perform unification (Section 2.1), and choice (Section 2.2). The details of choice and unification, and especially their interaction, are subtle, so this section does a lot of arm-waving. But fear not: Section 3 spells out the precise details. We only have space to describe one incarnation of $\mathcal{VC}$; Section 5 explores some possible alternative design choices.

2.1   Logical Variables and Equations

The Verse calculus includes first-class *logical variables* and *equations* that constrain their values. You can bring a fresh logical variable into scope with $\exists$, constrain a value to be equal to an expression with an equation $v = e$, and compose expressions in sequence with $eq; e$ (see Fig. 1). As an example, what might be written **let** $x = e_1$ **in** e_2 in a conventional functional language can be written $\exists x.\ x = e_1;\ e_2$ in $\mathcal{VC}$. The syntax carefully constrains both the form of equations and where they can appear: an equation $(v = e)$ always equates a *value* v to an expression e; and an equation can appear only to the left of a "; " (see *eq* in Fig. 1). The desugaring rules in Fig. 1 rewrite a general equation $e_1 = e_2$ (where e_1 is not a value) into equations of this simpler form.

A program executes by solving its equations, using the process of unification. For example,

$$\exists x\,y\,z.\ x = \langle y, 3\rangle;\ x = \langle 2, z\rangle;\ y$$

is solved by unifying x with $\langle y, 3\rangle$ and with $\langle 2, z\rangle$; that in turn unifies $\langle y, 3\rangle$ with $\langle 2, z\rangle$, which unifies y with 2 and z with 3. Finally, 2 is returned as the result. Note carefully that, as in any declarative language, *logical variables are not mutable*; a logical variable stands for a single, immutable value. We use "$\exists$" to bring a fresh logical variable into scope, because we really mean "there exists an x such that ...".

High-level functional languages usually provide some kind of pattern matching; in such a language, we might define *first* by $first\langle a, b\rangle = a$. Such pattern matching is typically desugared to

---

[2]This is a common pattern, often called "administrative normal form", or ANF [Sabry and Felleisen 1992]

more primitive **case** expressions, but in $\mathcal{VC}$ we do not need **case** expressions: unification does the job. For example we can define *first* like this:

$$first := \lambda p. \exists a\, b.\ p = \langle a, b \rangle;\ a$$

For convenience, we allow ourselves to write a term like $first\langle 2, 5 \rangle$, where we define the library function *first* separately with ":=". Formally, you can imagine each example e being wrapped with a binding for *first*, thus $\exists first.\ first = ...;\ e$, and similarly for other library functions.

This way of desugaring pattern matching means that the input to *first* is not required to be fully determined when the function is called. For example:

$$\exists x\, y.\ x = \langle y, 5 \rangle;\ 2 = first(x);\ y$$

Here $first(x)$ evaluates to y, which we then unify with 2. Another way to say this is that, as usual in logic programming, we may constrain the *output* of a function (here $2 = first(x)$), and thereby affect its *input* (here $\langle y, 5 \rangle$).

Although ";" is called "sequencing," the order of that sequence is immaterial for equations that do not contain choices (see Section 2.2 for the latter caveat). For example, consider $(\exists x\, y.\ x = 3 + y;\ y = 7;\ x)$. The sub-expression $3 + y$ is stuck until y gets a value. In $\mathcal{VC}$, we can unify x only with a *value*—we will see why in Section 2.2—and hence the equation $x = 3 + y$ is also stuck. No matter! We simply leave it and try some other equation. In this case, we can make progress with $y = 7$, and that in turn unlocks $x = 3 + y$ because now we know that y is 7, so we can evaluate $3 + 7$ to 10 and unify x with that. The idea of leaving stuck expressions aside and executing other parts of the program is called *residuation* [Hanus 2013][3], and is at the heart of our mantra "just solve the equations."

2.2 Choice

In conventional functional programming, an expression evaluates to a single value. In contrast, a $\mathcal{VC}$ expression evaluates to zero, one, or many values; or it can get stuck, which is different from producing zero values. The expression **fail** yields no values; a value v yields one value; and the choice $e_1 \mid e_2$ yields all the values yielded by e_1 followed by all the values yielded by e_2. Order is maintained and duplicates are not eliminated; we shall see why in Section 2.8. In short, an expression yields a *sequence* of values, not a bag, and certainly not a set.

The equations we saw in Section 2.1 can fail, if the arguments are not equal, yielding no results. Thus $3 = 3$ succeeds, while $3 = 4$ fails, returning no results. In general, we use "fail" and "returns no results" synonymously.

What if the choice was not at the top level of an expression? For example, what does $\langle 3, (7 \mid 5) \rangle$ mean? In $\mathcal{VC}$, it does *not* mean a pair with some kind of multi-value in its second component. Indeed, as you can see from Fig. 1, this expression is syntactically ill-formed. We must use the desugaring rules of Fig. 1, which give a name to that choice, thus: $\exists x.\ x = (7 \mid 5);\ \langle 3, x \rangle$. Now the expression is syntactically legal, but what does it mean? In $\mathcal{VC}$, a variable is never bound to a multi-value. Instead, x is successively bound to 7, and then to 5, like this:

$$\exists x.\ x = (7 \mid 5);\ \langle 3, x \rangle \quad \longrightarrow \quad (\exists x.\ x = 7;\ \langle 3, x \rangle) \mid (\exists x.\ x = 5;\ \langle 3, x \rangle)$$

We duplicate the context surrounding the choice, and "float the choice outwards". The same thing happens when there are multiple choices. For example:

$$\exists x\, y.\ x = (7 \mid 22);\ y = (31 \mid 5);\ \langle x, y \rangle \quad \text{yields the sequence } \langle 7, 31 \rangle, \langle 7, 5 \rangle, \langle 22, 31 \rangle, \langle 22, 5 \rangle$$

Notice that the order of the two equations now *is* significant:

$$\exists x\, y.\ y = (31 \mid 5);\ x = (7 \mid 22);\ \langle x, y \rangle \quad \text{yields the sequence } \langle 7, 31 \rangle, \langle 22, 31 \rangle, \langle 7, 5 \rangle, \langle 22, 5 \rangle$$

---

[3]Hanus did not invent the terms "residuation" and "narrowing," but his survey is an excellent introduction and bibliography.

Readers familiar with list comprehensions in Haskell and other languages will recognize this nested-loop pattern, but here it emerges naturally from choice as a deeply built-in primitive, rather than being a special construct for lists.

Just as we never bind a variable to a multi-value, we never bind it to **fail** either; rather we iterate over zero values, and that iteration of course returns zero values. So:

$$\exists x.\ x = \textbf{fail};\ 33 \quad \longrightarrow \quad \textbf{fail}$$

2.3   Mixing Choice and Equations

In the last section, we discussed what happens if there is a choice in the right-hand side (RHS) of an equation. What if we have equations under choice? For example:

$$\exists x.\ (x = 3;\ x + 1) \mathbin{|} (x = 4;\ x + 4)$$

Intuitively, "either unify x with 3 and yield $x + 1$, or unify x with 4 and yield $x + 4$". But there is a problem: so far we have said only "a program executes by solving its equations" (Section 2.1). Here, we can see two equations, $(x = 3)$ and $(x = 4)$, which are mutually contradictory, so clearly we need to refine our notion of "solving." The answer is pretty clear: in a branch of a choice, solve the equations in that branch to get the values for some logical variables, *and propagate those values to occurrences in that branch (only)*. Occurrences of that variable outside the choice are unaffected. We call this *local propagation*. This local-propagation rule would allow us to reason thus:

$$\exists x.\ (x = 3;\ x + 1) \mathbin{|} (x = 4;\ x + 4) \longrightarrow \exists x.\ (x = 3;\ 4) \mathbin{|} (x = 4;\ 8)$$

Are we stuck now? No, we can float the choice out as before[4],

$$\exists x.\ (x = 3;\ 4) \mathbin{|} (x = 4;\ 8) \longrightarrow (\exists x.\ x = 3;\ 4) \mathbin{|} (\exists x.\ x = 4;\ 8)$$

and now it is apparent that the sole occurrence of x in each $\exists$ is the equation $(x = 3)$, or $(x = 4)$ respectively; so we can drop the $\exists$ and the equation, yielding $(4 \mathbin{|} 8)$.

2.4   Pattern Matching and Narrowing

We remarked in Section 2.1 that we can desugar the pattern matching of a high-level language into equations. But what about multi-equation pattern matching, such as this definition in Haskell:

append [ ]   *ys* = *ys*
append $(x : xs)$ *ys* = $x :$ *append xs ys*

If pattern matching on the first equation fails, we want to fall through to the second. Fortunately, choice allows us to express this idea directly, where we use the empty tuple $\langle\rangle$ to represent the empty list and pairs to represent cons cells (see Fig. 1 to desugar the pattern-matching lambda):

$$\textit{append} := \lambda\langle xs, ys\rangle.\ ((xs = \langle\rangle;\ ys) \mathbin{|} (\exists x\, xr.\ xs = \langle x, xr\rangle;\ \langle x, \textit{append}\langle xr, ys\rangle\rangle))$$

If xs is $\langle\rangle$, the left-hand choice succeeds, returning ys; and the right-hand choice fails (by attempting to unify $\langle\rangle$ with $\langle x, xr\rangle$). If xs is of the form $\langle x, xr\rangle$, the right-hand choice succeeds, and we make a recursive call to *append*. Finally, if xs is built with head-normal forms other than the empty tuple and pairs, both choices fail, and *append* returns no results at all.

This approach to pattern matching is akin to *narrowing* [Hanus 2013]. Suppose $\textit{single} = \langle 1, \langle\rangle\rangle$, a singleton list whose only element is 1. Consider the call $\exists zs.\ \textit{append}\langle zs, \textit{single}\rangle = \textit{single};\ zs$. The call to *append* expands into a choice:

$$(zs = \langle\rangle;\ \textit{single}) \quad \mathbin{|} \quad (\exists x\, xr.\ zs = \langle x, xr\rangle;\ \langle x, \textit{append}\langle xr, \textit{single}\rangle\rangle)$$

---

[4]Indeed, we could have done so first, had we wished.

which amounts to exploring the possibility that zs is headed by $\langle\rangle$ or a pair—the essence of narrowing. It should not take long to reassure yourself that the program evaluates to $\langle\rangle$, effectively running *append* backwards in the classic logic-programming manner.

This example also illustrates that $\mathcal{VC}$ allows an equation (for *append*) that is recursive. As in any functional language with recursive bindings, you can go into an infinite loop if you keep fruitlessly inlining the function in its own right-hand side. It is the business of an *evaluation strategy* to do only rewrites that make progress toward a solution (Section 3.7).

2.5   Conditionals and one

Every source language will provide a conditional, such as **if** $(x=0)$ **then** e_2 **else** e_3. But what is the equality operator in $(x=0)$? One possibility, adopted by Curry [Antoy and Hanus 2021, §3.4], is this: there is one "=" for equations (as in Section 2.1), and another, say "==", for testing equality (returning a boolean with constructors *True* and *False*). $\mathcal{VC}$ takes a different, more minimalist position, following the lead of Icon (see Section 6.7). In $\mathcal{VC}$, *there is just one equality operator*, written "=" just as in Section 2.1. The expression **if** $(x=0)$ **then** e_2 **else** e_3 tries to unify x with 0. If that succeeds (yields one or more values), the **if** returns e_2; otherwise it returns e_3. There are no data constructors *True* and *False*; instead failure (returning zero values) plays the role of falsity.

But something is terribly wrong here. Consider $\exists x\, y.\ y = ($**if** $(x=0)$ **then** 3 **else** 4$);\ x = 7;\ y$. Presumably this is meant to set x to 7, test whether it is equal to 0 (it is not), and unify y with 4. But what is to stop us instead unifying x with 0 (via $(x=0)$), unifying y with 3, and then failing when we try to unify x with 7? Not only is that not what we intended, but it also looks very non-deterministic: the result is affected by the order in which we did unifications.

To address this, we give **if** a special property: in the expression **if** e_1 **then** e_2 **else** e_3, equations inside e_1 (the condition of the **if**) can only unify variables bound inside e_1; variables bound outside e_1 are called "rigid." So in our example, the x in $(x=0)$ is rigid and cannot be unified. Instead, the **if** is stuck, and we move on to unify $x=7$. That unblocks the **if** and all is well.

In fact, $\mathcal{VC}$ desugars the three-part **if** into something simpler, the unary construct **one**$\{e\}$. Its specification is this: if e fails, **one**$\{e\}$ fails; otherwise **one**$\{e\}$ returns the first of the values yielded by e. Now, **if** e_1 **then** e_2 **else** e_3 can (nearly) be re-expressed like this:

$$\mathbf{one}\{(e_1;\ e_2)\ \big|\ e_3\}$$

This isn't right yet, but the idea is this: if e_1 fails, the first branch of the choice fails, so we get e_3; if e_1 succeeds, we get e_2, and the outer **one** will select it from the choice. But what if e_2 or e_3 *themselves* fail or return multiple results? Here is a better translation, the one given in Fig. 1[5], which wraps the **then** and **else** branches within thunks[6]:

$$(\mathbf{one}\{(e_1;\ (\lambda\langle\rangle.\ e_2))\ \big|\ (\lambda\langle\rangle.\ e_3)\})\langle\rangle$$

The argument of **one** reduces to either $(\lambda\langle\rangle.\ e_2)\ \big|\ (\lambda\langle\rangle.\ e_3)$ or $(\lambda\langle\rangle.\ e_3)$ depending on whether e_1 succeeds or fails, respectively; **one** then picks the first value, that is, $\lambda\langle\rangle.\ e_2$ if e_1 succeeded or $\lambda\langle\rangle.\ e_3$ if e_1 failed, and applies it to $\langle\rangle$. As a bonus, provided we do no evaluation under a lambda, then e_2 and e_3 will remain unevaluated until the choice is made, just as we expect from a conditional.

We use the same local-propagation rule for **one** that we do for choice (Section 2.3). This, together with the desugaring for **if** into **one**, gives the "special property" of **if** described above.

---

[5]The translation in the figure also allows variables bound in the condition to scope over the **then** branch.
[6]Using thunks for the branches of a conditional is another very old idea; for example, see [Steele Jr. 1978, p. 54].

2.6 Tuples and all

The main data structure in $\mathcal{VC}$ is the *tuple*. A tuple is a finite sequence of values, $\langle v_1, \cdots, v_n \rangle$, where $n \geqslant 0$. A tuple can be used like a function: indexing is simply function application with the argument being integers from 0 and up. Indexing out of range is **fail**, as is indexing with a non-integer value. For example, $t := \langle 10, 27, 32 \rangle$; $t(1)$ reduces to 27 and $t := \langle 10, 27, 32 \rangle$; $t(3)$ reduces to **fail**.

What if we apply a tuple to a choice, *e.g.*, $\langle 10, 27, 32 \rangle (1 \mathbin{\|} 0 \mathbin{\|} 1)$? First we must desugar the application to the form $(v_1\ v_2)$, which is all $\mathcal{VC}$ permits (see Fig. 1), giving $x := (1 \mathbin{\|} 0 \mathbin{\|} 1)$; $\langle 10, 27, 32 \rangle(x)$, which readily reduces to $(27 \mathbin{\|} 10 \mathbin{\|} 27)$.

Tuples can be constructed by collecting all the results from a multi-valued expression, using the **all** construct: if e reduces to $(v_1 \mathbin{\|} \cdots \mathbin{\|} v_n)$, where $n \geqslant 2$, then **all**$\{e\}$ reduces to the tuple $\langle v_1, \cdots, v_n \rangle$; **all**$\{v\}$ produces the singleton tuple $\langle v \rangle$; and **all**$\{$**fail**$\}$ produces the empty tuple $\langle \rangle$. Note that $\mathbin{\|}$ is associative, which means that we can think of a sequence or tree of binary choices as really being a single n-way choice.

You might think that tuple indexing would be stuck until we know the index, but in $\mathcal{VC}$, the application of a tuple to a value rewrites to a choice of all the possible values of the index. For example, $t := \langle 10, 27, 32 \rangle$; $\exists i.\ t(i)$ looks stuck because we have no value for i, but this expression actually rewrites (via rule APP-TUP in Section 3.1) to:

$$\exists i.\ (i = 0;\ 10) \mathbin{\|} (i = 1;\ 27) \mathbin{\|} (i = 2;\ 32)$$

which (as we will see in Section 3) simplifies to just $(10 \mathbin{\|} 27 \mathbin{\|} 32)$. So **all** allows a choice to be reified into a tuple, and $(\exists i.\ t(i))$ allows a tuple to be turned back into a choice. The idea of rewriting a call of a function with a finite domain into a finite choice is called "narrowing" in the literature.

Do we even need **one** as a primitive construct, given that we have **all**? Can we not use $(\textbf{all}\{e\})(0)$ instead of **one**$\{e\}$? Indeed, they behave the same if e fully reduces to finitely many choices of values. But **all** really requires *every* arm of the choice tree to resolve to a value before proceeding, while **one** only needs the *first* choice to be a value. So, supposing that *loop* is a non-terminating function, **one**$\{1 \mathbin{\|} loop\langle\rangle\}$ can reduce to 1, while $(\textbf{all}\{1 \mathbin{\|} loop\langle\rangle\})(0)$ loops.

2.7 Programming in Verse

$\mathcal{VC}$ is a fairly small language, but it is quite expressive. For example, we can define the typical list functions one would expect from functional programming by using the duality between tuples and choices, as seen in Fig. 2. A tuple can be turned into choices by indexing with a logical variable i. Conversely, choices can be turned into a tuple using **all**. The choice operator "$\mathbin{\|}$" serves as both *cons* and *append* for choices; the corresponding operations for tuples are defined in Fig. 2. Partial functions, *e.g.*, *head*, will **fail** when the argument is outside of the domain.

Mapping a multi-valued function over a tuple is somewhat subtle. With *flatMap* the choices are flattened in the resulting tuple, *e.g.*, *flatMap*$\langle(\lambda x.\ x \mathbin{\|} x + 10), \langle 2, 3 \rangle\rangle$ reduces to $\langle 2, 12, 3, 13 \rangle$, whereas *map* keeps the choices. For example:

$$map\langle(\lambda x.\ x \mathbin{\|} x + 10), \langle 2, 3 \rangle\rangle \longrightarrow \langle(\lambda x.\ x \mathbin{\|} x + 10)(2), (\lambda x.\ x \mathbin{\|} x + 10)(3)\rangle \longrightarrow$$
$$\langle 2 \mathbin{\|} 12, 3 \mathbin{\|} 13 \rangle \longrightarrow \langle 2, 3 \rangle \mathbin{\|} \langle 2, 13 \rangle \mathbin{\|} \langle 12, 3 \rangle \mathbin{\|} \langle 12, 13 \rangle$$

Pattern matching for function definitions is simply done by unification of ordinary expressions; see the desugaring of pattern-matching lambda in Fig. 1. This in turn means that we can use ordinary abstraction mechanisms for patterns. For example, here is a function, *fcn*, that could be called as follows: *fcn*$\langle 88, 1, 99, 2 \rangle$.

$$fcn(t) := \exists x\,y.\ t = \langle x, 1, y, 2 \rangle;\ x + y$$

If we want to give a name to the pattern, it is simple to do so:

$$pat\langle v, w \rangle := \langle v, 1, w, 2 \rangle; \qquad fcn(t) := \exists x\,y.\ t = pat\langle x, y \rangle;\ x + y$$

$$
\begin{aligned}
head(xs) &:= xs(0) \\
tail(xs) &:= \mathbf{all}\{\exists i.\, i > 0;\ xs(i)\} \\
cons\langle x, xs\rangle &:= \mathbf{all}\{x \mid \exists i.\, xs(i)\} \\
append\langle xs, ys\rangle &:= \mathbf{all}\{(\exists i.\, xs(i)) \mid (\exists i.\, ys(i))\} \\
flatMap\langle f, xs\rangle &:= \mathbf{all}\{\exists i.\, f(xs(i))\} \\
map\langle f, xs\rangle &:= \mathbf{if}\ x := head(xs)\ \mathbf{then}\ cons\langle f(x), map\langle f, tail(xs)\rangle\rangle\ \mathbf{else}\ \langle\rangle \\
filter\langle p, xs\rangle &:= \mathbf{all}\{\exists i.\, x := xs(i);\ \mathbf{one}\{p(x)\};\ x\} \\
find\langle p, xs\rangle &:= \mathbf{one}\{\exists i.\, x := xs(i);\ \mathbf{one}\{p(x)\};\ x\} \\
some\langle p, xs\rangle &:= \mathbf{one}\{\exists i.\, p(xs(i))\} \\
zip\langle xs, ys\rangle &:= \mathbf{all}\{\exists i.\, \langle xs(i), ys(i)\rangle\}
\end{aligned}
$$

Desugaring of function definitions

$$
\begin{aligned}
f(x) := e &\quad \text{means} \quad f := \lambda x.\, e \\
f\langle x, y\rangle := e &\quad \text{means} \quad f := \lambda\langle x, y\rangle.\, e
\end{aligned}
$$

Fig. 2. Functions on tuples, analogous to list or array functions in some other languages

Patterns are truly first-class, going well beyond what can be done with, say, pattern synonyms in Haskell. For example, pat could be *computed*, like this:

$$pat\langle a, v, w\rangle := \mathbf{if}\ a = 0\ \mathbf{then}\ \langle v, 1, w, 2\rangle\ \mathbf{else}\ \langle 1, 1, w, v\rangle$$

so that the pattern depends on the value of a.

2.8   For Loops

The expression $\mathbf{for}(e_1)\ \mathbf{do}\ e_2$ will evaluate e_2 for each of the choices in e_1, rather like a list comprehension in languages like Haskell or Python. The scoping is peculiar[7] in that variables bound in e_1 also scope over e_2. So, for example, $\mathbf{for}(x := (2 \mid 3 \mid 5))\ \mathbf{do}\ (x + 1)$ will reduce to the tuple $\langle 3, 4, 6\rangle$.

Like list comprehension, $\mathbf{for}$ supports filtering; in $\mathcal{VC}$, this falls out naturally by just using a possibly failing expression in e_1. So, $\mathbf{for}(x := (2 \mid 3 \mid 5);\ x > 2)\ \mathbf{do}\ (x + 1)$ reduces to $\langle 4, 6\rangle$. Nested iteration in a $\mathbf{for}$ works as expected and requires nothing special. So, $\mathbf{for}(\exists x\, y.\, x = (10 \mid 20);\ y = (1 \mid 2 \mid 3))\ \mathbf{do}\ (x + y)$ reduces to $\langle 11, 12, 13, 21, 22, 23\rangle$.

Just as $\mathbf{if}$ is defined in terms of the primitive $\mathbf{one}$ (Section 2.5), we can define $\mathbf{for}$ in terms of the primitive $\mathbf{all}$. Again, we have to be careful when e_2 itself fails or produces multiple results; simply writing $\mathbf{all}\{e_1;\ e_2\}$ would give the wrong semantics. So we put e_2 within a lambda expression, and apply each element of the tuple to $\langle\rangle$ afterwards, using a map function (as defined in Fig. 2):

$$\mathbf{for}(\exists x_1 \cdots x_n.\, e_1)\ \mathbf{do}\ e_2 \quad \text{means} \quad v := \mathbf{all}\{\exists x_1 \cdots x_n.\, e_1;\ \lambda\langle\rangle.\, e_2\};\ map\langle\lambda z.\, z\langle\rangle, v\rangle$$

for a fresh variable v. Note how this achieves that peculiar scoping rule: the initial variables in $\exists x_1 \cdots x_n.\, e_1$ are in scope in e_2. Moreover, any effects (like being multi-valued) in e_2 will not affect the choices defined by e_1 since the effects are contained within that lambda. So, for example, $\mathbf{for}(x := (10 \mid 20))\ \mathbf{do}\ (x \mid x + 1)$ will reduce to $\langle 10, 20\rangle \mid \langle 10, 21\rangle \mid \langle 11, 20\rangle \mid \langle 11, 21\rangle$. At this point, it is crucial that the desugaring of $\mathbf{for}$ uses map, not $flatMap$.

Given that tuple indexing expands into choices, we can iterate over tuple indices and elements using $\mathbf{for}$. For example, $\mathbf{for}(\exists i\, x.\, x = t(i))\ \mathbf{do}\ (x + i)$ produces a tuple with the elements of t, each increased by its index within t. Notice the absence of the fencepost-error-prone iteration of i over $(0 .. size\,(t) - 1)$, common in most languages.

---

[7]But similar to C++, Java, Fortress, and Swift, and explained in $\mathcal{VC}$ by the subsequent desugaring into $\mathbf{all}$.

3 REWRITE RULES

How can we give a precise semantics to a programming language? Here are some possibilities:

- *A denotational semantics* is the classical approach, but it is tricky to give a (perspicuous) denotational semantics to a functional logic language because of the logical variables. We have such a denotational semantics under development, which we offer for completeness in Appendix E, but that is the subject of another paper.
- *A big-step operational semantics* typically involves explaining how a (heap, expression) starting point evaluates to a (heap, value) pair; Launchbury's natural semantics for lazy evaluation [Launchbury 1993] is the classic paper. The heap, threaded through the semantics, accounts for updating thunks as they are evaluated. Despite its "operational semantics" title, the big-step approach does not convey accurate operational intuition, because it goes all the way to a value in one step.
- *A small-step operational semantics* addresses this criticism: it describes how an abstract machine state (typically a (heap, expression, stack) triple) evolves, one small step at a time (*e.g.*, [Peyton Jones 1992]). The difficulty is that the description is now so low-level that it is again hard to explain to programmers.
- *A rewrite semantics* steers a middle path between the mathematical abstraction of denotational semantics and the over-concrete details of an operational semantics, whether big or small step. For example, Ariola *et al.*'s "A call-by-need lambda calculus" [Ariola et al. 1995] shows how to give the semantics of a call-by-need language as a set of rewrite rules. The great advantage of this approach is that it is easily explicable to programmers. In fact, teachers almost always explain the execution of Haskell or ML programs as a succession of rewrites of the program, such as: inline this call, simplify this **case** expression, *etc.*

Up to this point, there has been no satisfying rewrite semantics for functional logic languages (see Section 6 for previous work). Our main technical contribution is to fill this gap with a rewrite semantics for $\mathcal{VC}$, one that has the following properties:

- The semantics is expressed as a set of rewrite rules (Fig. 3). These rules apply to the core language of Fig. 1, after all desugaring.
- Any rule can be applied, in either direction, anywhere in the program term (including under lambdas).
- The rules are (mostly) oriented, with the intent that using them left-to-right makes progress.
- Despite this orientation, the rules do not say which rule should be applied where; that is the task of a separate *evaluation strategy* (Section 3.7).
- The rules can be applied by programmers to reason about what their program does, and by compilers to transform (and hopefully optimize) the program.
- There is no "magical rewriting" (Section 6.3): all the free variables on the right-hand side of a rule are bound on the left.

3.1 Functions and Function Application Rules

Looking at Fig. 3, the rule for a primitive operator like addition, APP-ADD, should be familiar: it simply rewrites an application of **add** to integer constants. For example **add**$\langle 3, 4 \rangle \longrightarrow 7$. Rules APP-GT and APP-GT-FAIL are more interesting: **gt**$\langle k_1, k_2 \rangle$ fails if $k_1 \leqslant k_2$ (rather than returning *False* as is more conventional), and returns k_1 otherwise (rather than returning *True*). An amusing consequence is that $(10 > x > 0)$ succeeds iff x is between 10 and 0 (comparison is right-associative).

An application of a primitive operator can rewrite only when its arguments are ground values; an application like **gt**$\langle x, 3 \rangle$ or **add**$\langle 7, x \rangle$ is *stuck* awaiting a value for x, which may arrive later, by substitution (Section 3.2). An ill-typed application, such as **gt**$\langle 3, \lambda x. x \rangle$, is simply stuck forever.

Application:

APP-ADD $\qquad\qquad\qquad$ $\mathbf{add}\langle k_1, k_2\rangle \longrightarrow k_3$ $\qquad\qquad\qquad\qquad$ where $k_3 = k_1 + k_2$

APP-GT $\qquad\qquad\qquad\quad$ $\mathbf{gt}\langle k_1, k_2\rangle \longrightarrow k_1$ $\qquad\qquad\qquad\qquad\;\;$ if $k_1 > k_2$

APP-GT-FAIL $\qquad\qquad$ $\mathbf{gt}\langle k_1, k_2\rangle \longrightarrow \mathbf{fail}$ $\qquad\qquad\qquad\quad$ if $k_1 \leqslant k_2$

APP-BETA$^\alpha$ $\qquad\qquad$ $(\lambda x.\, e)(v) \longrightarrow \exists x.\, x = v;\, e$ $\qquad\qquad\quad$ if $x \notin \mathrm{fvs}(v)$

APP-TUP $\qquad\quad$ $\langle v_0, \cdots, v_n\rangle(v) \longrightarrow \exists x.\, x = v;\, (x = 0;\, v_0)\,|\cdots|\,(x = n;\, v_n)$ $\;$ fresh $x \notin \mathrm{fvs}(v, v_0, \cdots, v_n)$

APP-TUP-0 $\qquad\qquad\qquad$ $\langle\rangle(v) \longrightarrow \mathbf{fail}$

Unification:

U-LIT $\qquad\qquad\qquad\;$ $k_1 = k_2;\, e \longrightarrow e$ $\qquad\qquad\qquad\qquad\;\;$ if $k_1 = k_2$

U-TUP $\quad$ $\langle v_1, \cdots, v_n\rangle = \langle v_1', \cdots, v_n'\rangle;\, e \longrightarrow v_1 = v_1';\, \cdots;\, v_n = v_n';\, e$

U-FAIL $\qquad\qquad\;$ $hnf_1 = hnf_2;\, e \longrightarrow \mathbf{fail}$ $\qquad\quad$ if U-LIT, U-TUP do not match

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\;$ and neither hnf_1 nor hnf_2 is a lambda

U-OCCURS $\qquad\qquad$ $x = V[x];\, e \longrightarrow \mathbf{fail}$ $\qquad\qquad\quad$ if $V \neq \square$

SUBST $\qquad\qquad$ $X[x = v;\, e] \longrightarrow (X\{v/x\})[x = v;\, e\{v/x\}]$ $\;$ if $v \neq V[x]$

HNF-SWAP $\qquad\qquad$ $hnf = v;\, e \longrightarrow v = hnf;\, e$

VAR-SWAP $\qquad\qquad$ $y = x;\, e \longrightarrow x = y;\, e$ $\qquad\qquad\quad$ if $x < y$

SEQ-SWAP $\qquad\qquad$ $eq;\, x = v;\, e \longrightarrow x = v;\, eq;\, e$ $\qquad$ unless (eq is $y = v'$ and $y \leq x$)

Elimination:

VAL-ELIM $\qquad\qquad\qquad$ $v;\, e \longrightarrow e$

EXI-ELIM $\qquad\qquad\qquad$ $\exists x.\, e \longrightarrow e$ $\qquad\qquad\qquad\qquad\quad$ if $x \notin \mathrm{fvs}(e)$

EQN-ELIM $\qquad$ $\exists x.\, X[x = v;\, e] \longrightarrow X[e]$ $\qquad\qquad\;$ if $x \notin \mathrm{fvs}(X[e])$ and $v \neq V[x]$

FAIL-ELIM $\qquad\qquad\quad$ $X[\mathbf{fail}] \longrightarrow \mathbf{fail}$

Normalization:

EXI-FLOAT$^\alpha$ $\qquad\qquad$ $X[\exists x.\, e] \longrightarrow \exists x.\, X[e]$ $\qquad\qquad\quad$ if $x \notin \mathrm{fvs}(X)$

SEQ-ASSOC $\qquad\quad$ $(eq;\, e_1);\, e_2 \longrightarrow eq;\, (e_1;\, e_2)$

EQN-FLOAT $\qquad$ $v = (eq;\, e_1);\, e_2 \longrightarrow eq;\, (v = e_1;\, e_2)$

EXI-SWAP $\qquad\qquad$ $\exists x.\, \exists y.\, e \longrightarrow \exists y.\, \exists x.\, e$

Choice:

ONE-FAIL $\qquad\qquad\qquad$ $\mathbf{one}\{\mathbf{fail}\} \longrightarrow \mathbf{fail}$

ONE-VALUE $\qquad\qquad\qquad\;$ $\mathbf{one}\{v\} \longrightarrow v$

ONE-CHOICE $\qquad\qquad$ $\mathbf{one}\{v\,|\,e\} \longrightarrow v$

ALL-FAIL $\qquad\qquad\qquad\;$ $\mathbf{all}\{\mathbf{fail}\} \longrightarrow \langle\rangle$

ALL-VALUE $\qquad\qquad\qquad$ $\mathbf{all}\{v\} \longrightarrow \langle v\rangle$

ALL-CHOICE $\qquad\;$ $\mathbf{all}\{v_1\,|\cdots|\,v_n\} \longrightarrow \langle v_1, \cdots, v_n\rangle$

CHOOSE-R $\qquad\qquad\qquad$ $\mathbf{fail}\,|\,e \longrightarrow e$

CHOOSE-L $\qquad\qquad\qquad$ $e\,|\,\mathbf{fail} \longrightarrow e$

CHOOSE-ASSOC $\qquad$ $(e_1\,|\,e_2)\,|\,e_3 \longrightarrow e_1\,|\,(e_2\,|\,e_3)$

CHOOSE $\qquad$ $SX[CX[e_1\,|\,e_2]] \longrightarrow SX[CX[e_1]\,|\,CX[e_2]]$

Note: In the rules marked with a superscript α, use α-conversion to satisfy the side condition.

Fig. 3. **The Verse Calculus: Rewrite rules**

$$
\begin{array}{ll}
\text{Execution contexts} & X ::= \Box \mid v = X;\ e \mid X;\ e \mid eq;\ X \\
\text{Value contexts} & V ::= \Box \mid \langle v_1, \cdots, V, \cdots, v_n \rangle \\
\text{Scope contexts} & SX ::= \mathbf{one}\{SC\} \mid \mathbf{all}\{SC\} \\
& SC ::= \Box \mid SC \mathbin{|} e \mid e \mathbin{|} SC \\
\text{Choice contexts} & CX ::= \Box \mid v = CX;\ e \mid CX;\ e \mid ceq;\ CX \mid \exists x.\ CX \\
\text{Choice-free exprs} & ce ::= v \mid ceq;\ ce \mid \mathbf{one}\{e\} \mid \mathbf{all}\{e\} \mid \exists x.\ ce \mid op(v) \\
& ceq ::= ce \mid v = ce
\end{array}
$$

Note: The $\Box$ in X can only be an expression, not an equation.

Fig. 4. The syntax of contexts

Note that equality, written $v = e$, is not a primitive operator; it is a language construct, and enjoys a rich set of built-in rewrite rules (see Section 3.2).

β-reduction is performed quite conventionally by APP-BETA; the only unusual feature is that on the RHS of the rule, we use an $\exists$ to bind x, together with $(x = v)$ to equate x to the argument. The rule may appear to use call-by-value, because the argument is a value v, but remember that values include variables, and a variable may be bound to an as-yet-unevaluated expression. For example:

$$\exists y.\ y = 3 + 4;\ (\lambda x.\ x + 1)(y)\ \longrightarrow\ \exists y.\ y = 3 + 4;\ \exists x.\ x = y;\ x + 1$$

Finally, the side condition $x \notin \mathrm{fvs}(v)$ in APP-BETA ensures that the $\exists x$ does not capture any variables free in v. If x appears free in v, α-conversion may be used on $\lambda x.\ e$ to rename x to $y \notin \mathrm{fvs}(v)$.

In $\mathcal{VC}$, tuples behave like (finite) functions in which application is indexing. Rule APP-TUP describes how tuple application works on non-empty tuples, while APP-TUP-0 deals with empty tuples. Notice that APP-TUP does not require the argument to be evaluated to an integer k; instead the rule works by narrowing. So the expression $\exists x.\ \langle 2, 3, 2, 7, 9 \rangle(x) = 2;\ x$ does not suspend awaiting a value for x; instead it explores all the alternatives (a form of narrowing), returning $(0 \mathbin{|} 2)$. This is a free design decision: a suspending semantics would be equally easy to express.

3.2 Unification Rules

Next we study unification, again in Fig. 3. Rules U-LIT and U-TUP are the standard rules for unification, going back nearly 60 years [Robinson 1965]. Rule U-FAIL makes unification fail (return zero results) on two different head-normal forms (see Fig. 1 for the syntax of *hnf*), except that it gets stuck if you attempt to unify a lambda with any other value, including itself. Why? Because equality of functions is undecidable, so in $\mathcal{VC}$ we simply refuse to run a program that tries to do so, just as we refuse to run $\mathbf{gt}\langle 3, \lambda x.\ x \rangle$ (see Section 3.1). This choice has consequences for confluence: see Section 4.1.1.

The standard "occurs check" is rule U-OCCURS, which makes use of a *context* V, whose syntax is given in Fig. 4. In general, a context [Felleisen and Friedman 1986; Felleisen et al. 1987] is a syntax tree containing a single hole, written $\Box$. The notation $V[v]$ is the term obtained by filling the hole in V with v. For example, U-OCCURS reduces $x = \langle 1, x, 3 \rangle;\ e$ to **fail** using the context $V = \langle 1, \Box, 3 \rangle$.

The key innovation in $\mathcal{VC}$ is the way bindings (that is, just ordinary equalities) of logical variables are propagated. The key rule is:

$$\text{SUBST} \quad X[x = v;\ e] \quad \longrightarrow \quad (X\{v/x\})[x = v;\ e\{v/x\}] \quad \text{if } v \neq V[x]$$

The rule says that if we have an equation $(x = v)$, we can replace the occurrences of x by v within the following expression and also within a surrounding context. This rule uses context X (Fig. 4),

and uses the notation $e\{v/x\}$ to mean "capture-avoiding substitution of v for x in e" (and similarly $X\{v/x\}$ to mean "capture-avoiding substitution of v for x in X"). There are several things to notice:

- SUBST fires only when the right-hand side of the equation is a *value* v, so that the substitution does not risk duplicating either work or choices. This restriction is precisely the same as the LET-v rule of Ariola et al. [1995] and, by not duplicating choices, it neatly implements so-called *call-time choice* [Hanus 2013]. We do not need a heap, or thunks, or updates; the equalities of the program elegantly suffice to express the necessary sharing.
- SUBST replaces all occurrences of x in X and e, but *it leaves the original $(x = v)$ undisturbed*, because X might not be big enough to encompass all occurrences of x. For example, we can rewrite $(y = x + 1;\ (x = 3;\ x + 3))$ to $(y = x + 1;\ (x = 3;\ 3 + 3))$, using $X = \square$, where the $\square$ is $x = 3;\ x + 3$. But that leaves an occurrence of x in $(y = x + 1)$. When there are no remaining occurrences of x, we may eliminate the binding by using rule EQN-ELIM (see Section 3.4).
- The side condition $v \neq V[x]$ in SUBST avoids an overlap with U-OCCURS.

3.3 Swapping and Binding Order

Two rules allow the two sides of an equation to be swapped. Rule HNF-SWAP helps SUBST to fire by putting the variable on the left. Rule VAR-SWAP, which swaps two variables, is needed for a more subtle reason. Consider the following example expression, where a and b are bound within some containing context, perhaps by lambdas. It can be rewritten in two different ways, where each column is a reduction sequence starting from the same initial term:

$$\exists x.\ x = \langle a \rangle;\ x = \langle b \rangle;\ x$$

$\longrightarrow\{\text{SUBST}\}$	$\exists x.\ x = \langle a \rangle;\ \langle a \rangle = \langle b \rangle;\ \langle a \rangle$		$\longrightarrow\{\text{SUBST}\}$	$\exists x.\ \langle b \rangle = \langle a \rangle;\ x = \langle b \rangle;\ \langle b \rangle$
$\longrightarrow\{\text{U-TUP}\}$	$\exists x.\ x = \langle a \rangle;\ a = b;\ \langle a \rangle$		$\longrightarrow\{\text{U-TUP}\}$	$\exists x.\ b = a;\ x = \langle b \rangle;\ \langle b \rangle$
$\longrightarrow\{\text{EQN-ELIM}\}$	$a = b;\ \langle a \rangle$		$\longrightarrow\{\text{EQN-ELIM}\}$	$b = a;\ \langle b \rangle$
$\longrightarrow\{\text{SUBST}\}$	$a = b;\ \langle b \rangle$		$\longrightarrow\{\text{SUBST}\}$	$b = a;\ \langle a \rangle$

The two sequences differ when it comes to which equation for x is chosen for SUBST in the first step of each column. As you can see, they conclude with two terms that are "obviously" the same semantically, but which are syntactically different. Rule VAR-SWAP allows them to be brought together (for example, $a = b;\ \langle b \rangle \longrightarrow\{\text{VAR-SWAP}\} b = a;\ \langle b \rangle \longrightarrow\{\text{SUBST}\} b = a;\ \langle a \rangle$), so that the unification rules can be syntactically confluent.

Rule SEQ-SWAP, which swaps adjacent equations within a sequence under certain conditions, is needed for a similar reason. Consider this example:

$$c = a;\ c = b;\ c$$

$\longrightarrow\{\text{SUBST}\}$	$c = a;\ a = b;\ a$	$\longrightarrow\{\text{SUBST}\} b = a;\ c = b;\ b$
$\longrightarrow\{\text{VAR-SWAP}\}$	$c = a;\ b = a;\ a$	$\longrightarrow\{\text{SUBST}\} b = a;\ c = a;\ a$

Again, the concluding terms of the two columns are "obviously" the same (they differ only in the order of the equations $b = a$ and $c = a$); SEQ-SWAP allows one term to be rewritten to match the other, making explicit our intuition that the order of equations of the form $x = v$ should not matter.

Observe the mysterious side condition $x \prec y$ in rule VAR-SWAP, and a similar one in SEQ-SWAP. In the overall proof of confluence, it turns out to be very helpful if the unification rules are *terminating* (see Section 4.2). To achieve this, VAR-SWAP fires on $y = x$ only if x's binding occurs in the scope of y's binding, written $x \prec y$, so that the innermost-bound variable ends up on the left. Similarly, the side condition on SEQ-SWAP prevents it firing infinitely.

Other rules, notably EXI-SWAP, may change this binding order and thereby re-enable VAR-SWAP or SEQ-SWAP, but the unification rules *considered in isolation* are terminating and confluent, and that is what we need for the proof (but see Section 5.1).

3.4 Elimination and Normalization Rules

Four *elimination* rules allow dead code to be dropped (Fig. 3): val-elim discards a value to the left of a semicolon; exi-elim discards a dead existential; eqn-elim discards an existential $\exists x$ that binds a variable whose only occurrence is a single equation $x = v$; and fail-elim discards the execution context surrounding a **fail**. Note that none of these rules, except fail-elim, discards an unevaluated expression, because that expression might fail and we don't want to "lose" that failure (see Section 3.6). The exception is fail-elim, whose purpose is to propagate a known failure, so losing some other failure that might have been in the discarded execution context does not matter.

Four *normalization* rules help to put the expression in a form that allows other rules to fire (Fig. 3): exi-float allows an existential to float outwards; seq-assoc makes semicolon right-associated; eqn-float moves work out of the right-hand side of an equation $v = e$. For example, we cannot use subst to substitute for x in $(x = (e;\ 3);\ x + 2)$, because the RHS of the x-equation is not a value; but we can instead apply eqn-float to get $e;\ x = 3;\ x + 2$, and now we *can* apply subst.

Rule exi-swap allows you to move an existential inward so that a dead equation can be eliminated by eqn-elim. Rule exi-swap is unusual because it can be infinitely applied; avoiding that eventuality is easily achieved by tweaking the evaluation strategy (Section 3.7).

Note that all these swapping and normalization rules *preserve the left-to-right sequencing of expressions*, which matters because choices are made left-to-right, as we saw in Section 2.3. Moreover, the rules do not float equalities or existentials out of choices: that restriction is the key to localizing unification (Section 2.3) and to the flexible/rigid distinction of Section 2.5. For example, consider the expression $(y = ((x = 3;\ x + 5) \mathbin{\|} (x = 4;\ x + 2));\ \langle x + 1, y \rangle)$. We must not float the binding $(x = 3)$ up to a point where it might interact with the expression $(x + 1)$, because the latter is outside the choice, and a different branch of the choice binds x to 4.

3.5 Rules for Choice

The rules for choice are given in Fig. 3:

- Rules one-fail, one-value, and one-choice describe the semantics of **one**, as in Section 2.5.
- Similarly, all-fail, all-value, and all-choice describe the semantics of **all** (Section 2.6).
- Rules choose-r and choose-l eliminate **fail**, which behaves as an identity for choice.
- Rule choose-assoc associates choice to the right, so that one-choice or all-choice can fire. (The dots on the left of all-choice should be read as a string of right-associated choices.)

The most interesting rule is choose, which, just as described in Section 2.2, "floats the choice outwards," duplicating the surrounding context. But what "surrounding context" precisely? We use two new contexts, SX and CX, both defined in Fig. 4. A *choice context* CX is like an execution context X, *but with no possible choices to the left of the hole*:

$$CX ::= \square \mid v = CX \mid CX;\ e \mid ce;\ CX \mid \exists x.\ CX$$

Here, ce is a guaranteed-choice-free expression (syntax in Fig. 4). This syntactic condition is necessarily conservative; for example, a call $f(x)$ is considered not guaranteed-choice-free because it depends on what function f does. We must guarantee not to have choices to the left so that we preserve the order of choices—see Section 2.3.

The context SX (Fig. 4) in choose ensures that CX is as large as possible. This is a very subtle point: without this restriction we lose confluence. To see this, consider[8]:

$$\exists x.\ (\textbf{if } (x > 0) \textbf{ then } 55 \textbf{ else } (44 \mathbin{\|} 2));\ x = 1;\ (77 \mathbin{\|} 99)$$
$$\longrightarrow \{\text{subst}\} \qquad \exists x.\ (\textbf{if } (1 > 0) \textbf{ then } 55 \textbf{ else } (44 \mathbin{\|} 2));\ x = 1;\ (77 \mathbin{\|} 99)$$
$$\longrightarrow \{\text{simplify } \textbf{if}\} \ \ \exists x.\ 55;\ x = 1;\ (77 \mathbin{\|} 99) \quad \longrightarrow \{\text{val-elim, eqn-elim}\} \ \ 77 \mathbin{\|} 99$$

---

[8]Remember, **if** is syntactic sugar for a use of **one** (see Section 2.5), but using **if** makes the example easier to understand.

But suppose instead we floated the choice out *partway*, like this[9]:

$$\exists x. \,(\textbf{if } (x > 0) \textbf{ then } 55 \textbf{ else } (44 \,\|\, 2)); \; x = 1; \; (77 \,\|\, 99)$$
$$\longrightarrow\{\text{Bogus \textsc{choose}}\} \quad \exists x. \,(\textbf{if } (x > 0) \textbf{ then } 55 \textbf{ else } (44 \,\|\, 2)); \; ((x = 1; \; 77) \,\|\, (x = 1; \; 99))$$

Now the $(x = 1)$ is inside the choice branches, so we cannot use \textsc{subst} to substitute for x in the condition of the **if**. Nor can we use \textsc{choose} again to float the choice further out because the **if** is not guaranteed choice-free (in this example, the **else** branch has a choice). So, alas, we are stuck! Our not-entirely-satisfying solution is to force \textsc{choose} to float the choice all the way to the top. The *SX* context (Fig. 4) formalizes what we mean by "the top": rule \textsc{choose} can float a choice outward only when it becomes part of the choice tree (context *SC*) immediately under a **one** or **all** construct (context *SX*).

Rule \textsc{choose} moves choices around; only \textsc{one-choice} and \textsc{all-choice} *decompose* choices. So choice behaves a bit like a data constructor, or normal form, of the language. This contrasts with other semantic approaches that eliminate choice by non-deterministically picking one branch or the other, which immediately gives up confluence. However, in implementation terms, treating choice a bit like a data constructor is the basis for *pull-tabbing* in logically-complete implementations of non-deterministic functional logic languages like Curry [Antoy 2011].

3.6 The Verse Calculus is Lenient

$\mathcal{VC}$ is *lenient* [Schauser and Goldstein 1995], not lazy (call-by-need), nor strict (call-by-value). Under lenient evaluation, functions can run before their arguments have a ground value (as in a lazy language), but (almost) everything is eventually evaluated (as in a strict language).

$\mathcal{VC}$ cannot be lazy, because of $\mathcal{VC}$'s commitment to determinism and, particularly, its first-class **all** operator. Consider:

$$\textbf{all}\{\exists x. \, x = e; \; 3\}$$

In a lazy language like Curry, the expression $\exists x. \, x = e; \; 3$ would return one result, 3, regardless of e, because x is unused, and so $x = e$ is simply dead code. But in $\mathcal{VC}$ we must evaluate e, and the **all** expression will yield a tuple with one element for each result yielded by e. If e fails (returns zero results), then the result tuple is empty. Since **all** reifies these results into a tuple, the multiplicity of results is visible to the context of the **all** expression, via the length of the returned tuple. Moreover, consider

$$\textbf{all}\{\exists x, y. \, x = ((y = 3; \, 1) \,\|\, (y = 4; \, 2)); \; y\}$$

In $\mathcal{VC}$, the expression to which x is equated must be evaluated, yielding two values, 1 and 2, which are then discarded since x is unused. However, each of these values is accompanied by a distinct binding for y, so the result of the whole expression is the tuple $\langle 3, 4 \rangle$. In short, we must eventually evaluate the expression to which x is equated not only to get the *size* of the tuple, but also the *values of its elements*, via the bindings of y.

So it seems hard to reconcile laziness with a deterministic **all**. But $\mathcal{VC}$ is not strict, either; in $\mathcal{VC}$ a function can be called without its argument having a ground value. For example:

$$\exists f. \, f = (\lambda x. \, x = 3; \; x); \; \exists y. \, f(y)$$

Here we can call f on an uninstantiated existential variable y; we do not have to wait until y gets a ground value from the calling context. Rather, in this example it is the body of f that then constrains y to be 3. This is part of the essence of functional logic programming, and indeed in $\mathcal{VC}$ a variable *is* a value (Fig. 1). When lenience was introduced in the data-flow language Id [Schauser and Goldstein

---

[9]As in the previous example, here and elsewhere we freely rewrite terms that have not been fully desugared, but that is just an expository aid; formally, the rewrite rules of Fig. 3 apply only to programs in the basic "Syntax" language of Fig. 1.

1995], it was a way to get *parallelism*, and it certainly justifies parallelism in an implementation of $\mathcal{VC}$. But for $\mathcal{VC}$ lenience has *semantic* significance, too, as we see in this example.

So in $\mathcal{VC}$ functions can be called on as-yet-unevaluated arguments, but almost all redexes that are not under a lambda get evaluated eventually. You can see this enforced in the rule for **one** in Figure 3, which fires only when the first choice in the body is a value v, and similarly for **all**. Why "almost" all redexes? Because the **one** operator returns the value of its first choice and abandons all other choices. For example **one**$\{1 \mid loop\langle\rangle\}$ returns 1, discarding the infinite $loop\langle\rangle$.

Note that lenience, like laziness, supports *abstraction* in the presence of unification. For example, we can replace an expression $(x = \langle y, 3 \rangle; \ y > 7)$ by

$$\exists f. f = (\lambda\langle p, q \rangle. \ p = \langle q, 3 \rangle; \ q > 7); \ f\langle x, y \rangle$$

Here, we abstract over the free variables of the expression and define a named function f. Calling the function is just the same as writing the original expression. This transformation would not be valid under call-by-value.

3.7 Evaluation Strategy

Any rewrite rule can apply anywhere in the term, at any time. For example, in the term $(x = 3 + 4; \ y = 4 + 2; \ x + y)$ the rewrite rules do not say whether to rewrite $3 + 4 \rightarrow 7$ and then $4 + 2 \rightarrow 6$, or the other way around. The rules do, however, require us to reduce $3 + 4 \rightarrow 7$ before substituting for x in $x + y$, because rule subst fires only when the RHS is a value. The rewrite rules thereby express *semantics*.

For example, in the lambda calculus, by changing the rewrite rule β to βV, we change the language from call-by-name to call-by-value; by adding **let**, plus suitable rewrite rules, we can express call-by-need [Ariola et al. 1995]. In $\mathcal{VC}$, the rewrite rules are carefully crafted in a similar way; for example, subst will substitute $x = v$ only when the equation binds a variable to a *value*, rather like βV in the lambda calculus. Similarly, the elimination rules never discard a term that could fail.

In any term there may of course be many redexes—that is good. An *evaluation strategy* answers the question: given a closed term, which unique redex, out of the many possible redexes, should be rewritten next to make progress toward the result? Let us call an evaluation strategy *normalizing* if it guarantees to terminate if there is *any* terminating sequence of reductions—that is, if any path terminates with a value, then a normalizing evaluation strategy will terminate with that same value[10]. For example, in the pure lambda calculus, *normal-order reduction*, sometimes called *leftmost-outermost reduction*, is a normalizing strategy.

For $\mathcal{VC}$, a first step towards a normalizing strategy is to avoid no-ops. For example, there is no point in applying rule subst if $x \notin \text{fvs}(e)$; nor in applying fail-elim if $X = \square$. Next, a normalizing strategy must be careful with exi-swap, because it can easily apply infinitely; its role is to make it possible to use exi-float. Dealing with these no-ops and flip-flops is not hard.

With that done, we believe that a fair outermost strategy is normalizing for $\mathcal{VC}$: at each step, just select any redex that is not within a lambda or another redex. That sounds easy, but is tricky in practice for two reasons. First, a reduction may "unlock" a redex far to its left. For example, consider $(x + y; \ \langle x, 3 \rangle = \langle 2, y \rangle; \ x)$. The $(x + y)$ is not a redex, but the equation is; we can apply unification to get $(x = 2; \ y = 3)$, and then use substitution to rewrite the $(x + y)$ to $(2 + 3)$; and *now* the $(2 + 3)$ is a redex. Hence, a challenge for an implementation is to find the next redex efficiently. Secondly, as with other functional-logic languages such as Curry, a normalizing strategy must be able to pursue redexes within multiple subterms in parallel (or alternately), without starving any of

---

[10]It would be even better if the strategy could guarantee to find the result in the minimal number of rewrite steps—so-called "optimal reduction" [Asperti and Guerrini 1999; Lamping 1990; Lévy 1978]—but optimal reduction is typically very hard, even in theory, and invariably involves reducing under lambdas, so for practical purposes it is well out of reach.

them, hence "fair". For example ($loop\langle\rangle$; **fail**) and (**fail**; $loop\langle\rangle$) should both fail, so a normalizing strategy must, in both cases, avoid getting stuck in the $loop$.

We have not yet, however, formalized such a strategy for $\mathcal{VC}$ or proved it normalizing. We have several prototype implementations of $\mathcal{VC}$, each involving an abstract machine with a stack, a heap, a bunch of previously blocked computations, and so on. Exploring this design space is, however, beyond the scope of this paper.

3.8 Logical Completeness

Our rewrite rules seek to implement the idea articulated in Section 2.1 that a program executes by solving its equations, specifically by finding values for the logical variables that satisfy the equations. Ideally, we would like to promise to find *all* substitutions for the logical variables that satisfy the equations: this is *logical completeness*. Alas, in practice logical completeness is well out of reach, at least in systems that support arithmetic. For example, consider $\exists x. (x * x * x + 3 * x - 8) = 0; x$ (assuming a version of $\mathcal{VC}$ extended to have a multiplication operator $*$ and subtraction operator $-$). To solve this would require solving a cubic equation, and it is equally easy to express very difficult problems such as Fermat's last theorem. Moreover, remembering that $\mathcal{VC}$ is deterministic, in what order would the (infinite sequence of) results of, say, $\exists x. x + 1$ be returned? So instead we satisfy ourselves with a computable and deterministic approximation to this Platonic ideal, an approximation that is described precisely by our rewrite rules, using unification as the mechanism. Expressions that unification cannot solve, like the one above, are stuck; no rewrite rule applies. All functional logic languages share this challenge to logical completeness.[11]

The boundary between "stuck" and "soluble" can be quite subtle. Consider this tricky term: $\exists x. x = \mathbf{if}\ (x = 0;\ x > 1)\ \mathbf{then}\ 33\ \mathbf{else}\ 55; x$. At first we might think it was stuck—how can we simplify the **if** when its condition mentions x, which is not yet defined? But in fact, rule SUBST allows us to substitute *locally* in any X-context surrounding the equation ($x = 0$) thus:

$$\exists x. x = \mathbf{if}\ (x = 0;\ x > 1)\ \mathbf{then}\ 33\ \mathbf{else}\ 55; x$$
$$\longrightarrow\{\text{SUBST}\} \qquad \exists x. x = \mathbf{if}\ (x = 0;\ 0 > 1)\ \mathbf{then}\ 33\ \mathbf{else}\ 55; x$$
$$\longrightarrow\{\text{APP-GT-FAIL}\}\ \exists x. x = \mathbf{if}\ (x = 0;\ \mathbf{fail})\ \mathbf{then}\ 33\ \mathbf{else}\ 55; x$$
$$\longrightarrow\{\text{FAIL-ELIM}\} \qquad \exists x. x = \mathbf{if}\ \mathbf{fail}\ \mathbf{then}\ 33\ \mathbf{else}\ 55; x$$
$$\longrightarrow\{\text{simplify}\ \mathbf{if}\}\ \exists x. x = 55; x \quad \longrightarrow\{\text{SUBST}\}\ \exists x. x = 55; 55 \quad \longrightarrow\{\text{EQN-ELIM}\}\ 55$$

Minor variants of the same example get stuck instead of reducing. For example, if we replace the $x = 0$ with $x = 100$ then rewriting gets stuck, as the reader may verify (we cannot eliminate the equation $x = 100$); and yet there is a substitution that will satisfy the equations, namely $\{55/x\}$. And if we replace $x = 0$ with $x = 55$, then rewriting again gets stuck; this time there is *no* substitution that will satisfy the equations, and yet the expression gets stuck rather than failing.

3.9 Developing and Debugging Rules

The rules we describe here should both be able to transform a program to its value, and also be confluent. To aid in the development of the rules, we have used several mechanized tools to automate reduction, random test-case generation, and confluence checking. Initially, we used PLT Redex [Felleisen et al. 2009], which is very easy to use but not very efficient. For better efficiency we switched to a Haskell library for term rewriting. The library provides a DSL for writing rules, and provides the infrastructure to apply the rules everywhere, detect cycles, provide traces, *etc.* Some sample rewrite rules can be found in Fig. 5.

---

[11]For countable domains, like the integers, or even the algebraic numbers—which would be adequate to solve our cubic equation example—we could in theory simply try all possible values of x one by one; but that is hardly practical.

$$
\begin{aligned}
\textit{rules lhs} \ = \ & \texttt{"APP-ADD"} \ \text{`name`} \ (\textbf{do} \ Op \ Add : @ : Tup \ [\textit{Int k1}, \textit{Int k2}] \leftarrow [\textit{lhs}] \\
& \qquad\qquad\qquad\qquad\quad \textit{pure} \ (\textit{Int} \ (k1 + k2))) \\
\diamondsuit \ & \texttt{"EXI-SWAP"} \ \text{`name`} \ (\textbf{do} \ EXI \ x \ (EXI \ y \ e) \leftarrow [\textit{lhs}] \\
& \qquad\qquad\qquad\qquad\quad \textit{pure} \ (EXI \ y \ (EXI \ x \ e))) \\
\diamondsuit \ & \texttt{"EQN-ELIM"} \ \text{`name`} \ (\textbf{do} \ EXI \ x \ a \leftarrow [\textit{lhs}] \\
& \qquad\qquad\qquad\qquad\quad (\textit{ctx}, (\textit{Var} \ x' := : \textit{Val} \ v) :> : e) \leftarrow \textit{execX} \ a \\
& \qquad\qquad\qquad\qquad\quad \textit{guard} \ (x = x' \ \wedge \ x \notin \textit{free} \ (\textit{ctx} \ (v :> : e))) \\
& \qquad\qquad\qquad\qquad\quad \textit{pure} \ (\textit{ctx} \ e))
\end{aligned}
$$

Fig. 5. Sample Haskell reduction rules

We used this infrastructure in two ways. First, we have a set of examples with known results, against which we can test a potential rule set. Second, before beginning a proof of confluence, we used QuickCheck [Claessen and Hughes 2000] to generate test cases and check them for confluence. QuickCheck turned out to be invaluable at finding counterexamples to otherwise reasonable-looking rules; it has run on the order of 100 million tests on the current rule set.

4 METATHEORY

The rules of our rewrite semantics can be applied anywhere, in any order; they give meaning to programs without committing to a specific evaluation strategy. But then it had better be true that no matter how the rules are applied, one always obtains the same result! That is, our rules should be *confluent*. In this section, we describe our proof of confluence. Because the rule set is quite big (compared, say, to the pure lambda calculus), this proof turns out to be a substantial undertaking.

Reductions and confluence. A *binary relation* is a set of pairs of related items. A *reduction relation* $\mathcal{R}$ is the *compatible closure*[12] of any binary relation on a set of tree-structured *terms,* such as the terms generated by some BNF grammar. We write $\mathcal{R}^*$ for the reflexive transitive closure of $\mathcal{R}$. We write $e \rightarrow_{\mathcal{R}} e'$ (*e steps to e'*) if $(e, e') \in \mathcal{R}$ and $e \twoheadrightarrow_{\mathcal{R}} e'$ (*e reduces to e'*) if $(e, e') \in \mathcal{R}^*$. A reduction relation $\mathcal{R}$ is *confluent* if whenever $e \twoheadrightarrow_{\mathcal{R}} e_1$ and $e \twoheadrightarrow_{\mathcal{R}} e_2$, there exists an e' such that $e_1 \twoheadrightarrow_{\mathcal{R}} e'$ and $e_2 \twoheadrightarrow_{\mathcal{R}} e'$. Confluence gives us the assurance that we will not get different results when choosing different rules, or get stuck with some rules and not with others.

Normal forms and unicity. A term e is an $\mathcal{R}$-*normal form* if there does not exist any e' such that $e \rightarrow_{\mathcal{R}} e'$. Confluence implies uniqueness of normal forms (unicity): if $e \twoheadrightarrow_{\mathcal{R}} e_1$ and $e \twoheadrightarrow_{\mathcal{R}} e_2$, and e_1 and e_2 are normal forms, then $e_1 = e_2$ [Barendregt 1984, Corollary 3.1.13(ii)].

4.1 Two Challenges to Confluence

Confluence is a purely syntactic property, which leads to two tiresome challenges.

4.1.1 Tiresome Challenge 1: Unifying Lambdas. In $\mathcal{VC}$, an attempt to unify a lambda with another value is stuck (Section 3.2). That choice has two consequences. One is easy to handle, one is tiresome.

First, while u-lit lets us eliminate equalities on the same literal $k = k$, $\mathcal{VC}$ has no analogous u-var rule to drop equalities on the same variable $x = x$. To see why not, suppose we had such a u-var rule, and consider the term $(\exists x. \ x = (\lambda y. \ y); \ x = x; \ 0)$. If we first apply u-var to eliminate the equality $x = x$, then the remainder reduces to 0. However, if we first subst the equality $x = (\lambda y. \ y)$, we get $((\lambda y. \ y) = (\lambda y. \ y); \ 0)$, which is stuck. Therefore, to preserve confluence, $\mathcal{VC}$ has no rule u-var: such equalities can be eliminated only after the value of x is substituted in and checked to not be a lambda.

---

[12]"Compatible closure" means that, for any context E and any two terms M and N, if $(M, N) \in \mathcal{R}$ then $(E[M], E[N]) \in \mathcal{R}$.

Second, the inability to fail when unifying different lambdas leads to non-confluence. Here is an expression that rewrites in two different ways, depending on which equation we subst first:

$$(\lambda p.\,1) = (\lambda q.\,2);\ 1 \quad \twoheadleftarrow \quad \exists x.\ x = (\lambda p.\,1);\ x = (\lambda q.\,2);\ x\ \langle\rangle \quad \twoheadrightarrow \quad (\lambda q.\,2) = (\lambda p.\,1);\ 2$$

These two outcomes cannot be joined. This non-confluence is tiresome because the program is wrong anyway: we should not be attempting to unify lambdas.

4.1.2 Tiresome Challenge 2: Recursion and the Notorious Even/odd Problem. It is well known that adding **letrec** to the lambda calculus makes it non-confluent, in a very tiresome, but hard-to-avoid, way [Ariola and Blom 2002]. In our context, consider the term:

$$\exists x\,y.\ x = \langle 1, y \rangle;\ y = (\lambda z.\,x);\ x \quad \twoheadrightarrow \quad \exists y.\ y = (\lambda z.\,\langle 1, y \rangle);\ \langle 1, y \rangle \qquad \text{(1) substitute for } x \text{ first}$$

$$\exists x\,y.\ x = \langle 1, y \rangle;\ y = (\lambda z.\,x);\ x \quad \twoheadrightarrow \quad \exists x.\ x = \langle 1, \lambda z.\,x \rangle;\ x \qquad \text{(2) substitute for } y \text{ first}$$

The results of (1) and (2) have the same meaning (are indistinguishable by a $\mathcal{VC}$ context) but cannot be joined by our rewrite rules. Nor is this easily fixed by adding new rules, as we did when we added var-swap (Section 3.2) and seq-swap (Section 3.4). Why not? Because the terms are equivalent only under some kind of graph isomorphism.

4.1.3 Resolving the Two Challenges. We explored three different ways to address these challenges. The first is simply to abandon confluence as a goal altogether. Confluence is, after all, purely *syntactic*, and hence much stronger than what we really need, which is only that each of our rules be *semantics*-preserving. But that, of course, requires an independent notion of semantics, a direction we sketch in Appendix E.

Second, we can simply exclude programs that unify lambdas, or that use recursion. To be precise:

- An equation is *obviously problematic* if either
  - it is an equation of form $hnf = \lambda x.\,e$ or $\lambda x.\,e = hnf$, or
  - it is an equation of the form $x = V[\lambda y.\,e]$, where $x \in \text{fvs}(e)$.
- A term e is *obviously problematic* if it contains an obviously problematic equation.
- A term e is *problematic* if there exists an obviously problematic term e' such that $e \twoheadrightarrow e'$.
- A term e is *well-behaved* if it is not problematic.

Then we prove confluence only for well-behaved terms. We take this approach for our main proof in this paper (Section 4.2). Excluding lambda unification is fine; programmers simply cannot expect that to work. Excluding recursion may seem drastic, but no expressiveness is lost thereby: in our untyped setting, one can still write recursive (and non-terminating) programs using one's favorite fixpoint combinator, such as **Y** or **Z**.

But excluding recursion is not entirely satisfying: it is hard to prove that a term has no recursion, and it is clumsy to write recursive programs using only **Y**-combinators. Our third approach is to adopt the idea of *skew confluence* [Ariola and Blom 2002], a clever technique developed specifically to handle the even/odd problem; we give an overview of skew confluence in Section 4.3, and provide details of our approach to a proof of skew confluence in Appendix D, with several new lemmas, but we emphasize that the proof of skew confluence is not yet complete.

4.2 Proof of Confluence

Our main result is that $\mathcal{VC}$'s reduction rules are confluent for well-behaved terms. We sketch the proof here, with full details in Appendix C (and relevant preliminaries in Appendix B).

THEOREM 4.1 (CONFLUENCE). *The reduction relation in Fig. 3 is confluent for well-behaved terms.*

Proof sketch. Our proof strategy is to (1) divide the rules into groups for application, unification, *etc.*, approximately as in Fig. 3, (2) prove confluence for each separately, and then (3) prove that their

combination is confluent via commutativity. Given two reduction relations R and S, we say that R *commutes* with S if for all terms e, e_1, e_2 such that $e \twoheadrightarrow_R e_1$ and $e \twoheadrightarrow_S e_2$ there exists e' such that $e_1 \twoheadrightarrow_S e'$ and $e_2 \twoheadrightarrow_R e'$. We prove each individual sub-relation is confluent and that they pairwise commute. Then confluence of their union follows, using a theorem of Huet [1980]: If R and S are confluent and commute, then $R \cup S$ is confluent. Proving confluence for application, elimination and choice is easy: they all satisfy the *diamond property*—namely, that two different reduction steps can be joined at a common term *by a single step*—which suffices to show the relations are confluent [Barendregt 1984]. The diamond property itself can be verified easily by considering critical pairs of transitions. The rules for unification and normalization, however, pose two problems.

The unification problem. The first problem is that the unification relation does not satisfy the diamond property—it may need multiple steps to join the results of two different one-step reductions. For example, consider the term $(x = \langle 1, y \rangle; \ x = \langle z, 2 \rangle; \ x = \langle 1, 2 \rangle; \ 3)$. The term can be reduced in one step by substituting x in the third equation by either $\langle 1, y \rangle$ or $\langle z, 2 \rangle$. After this, it will take multiple steps to join the two terms.

Following a well-trodden path in proofs of confluence for the λ-calculus (*e.g.*, [Barendregt 1984]), our proof of confluence for the unification rules works in two stages. First, we prove that the reductions are *locally confluent*, meaning if e single-steps to each of e_1 and e_2, then e_1 and e_2 can be joined at some e' by taking *multiple* unification rule steps. Second, we prove that the unification reductions are *terminating*, which relies upon eliminating recursion in tuples via U-OCCURS and in lambdas via the well-behaved-ness condition. Newman's Lemma [Huet 1980, Lemma 2.4] then implies that the locally confluent, terminating unification relation is also confluent.

The normalization problem. The second problem is that the normalization rules do not commute with the unification rules. Recall from Section 3.3 that the unification rules rely upon variable ordering to orient equations between variables in a canonical fashion. The normalization rule EXI-SWAP can *change* the variable order and hence, its behavior is deeply intertwined with unification and cannot be factored out via a commutativity argument. Instead, we prove that the union of unification and normalization is confluent by showing that unification *postpones after* normalization [Hindley 1964]; see Appendix C for the gory details.

4.3 Overview of Skew Confluence

We travel a path very similar to the one blazed by Ariola and her co-authors. Ariola and Klop studied a form of the lambda calculus with an added letrec construct and determined (like us) that their calculus was not confluent; then they added a specific constraint on recursive substitution and proved that the modified calculus is confluent [Ariola and Klop 1994, 1997]. In a later paper, Ariola and Blom proved that their calculus without the constraint, while not confluent, does obey a weaker related property, which they invented, called *skew confluence* [Ariola and Blom 2002]. We believe, and currently are trying to prove, that $\mathcal{VC}$ without the pesky no-recursion side condition of Theorem 4.1 is skew confluent.

$$\text{Confluence:} \quad \forall e, e_1, e_2. \ e \twoheadrightarrow_{\mathcal{R}} e_1 \wedge e \twoheadrightarrow_{\mathcal{R}} e_2 \implies \exists e'. \ e_1 \twoheadrightarrow_{\mathcal{R}} e' \wedge e_2 \twoheadrightarrow_{\mathcal{R}} e'.$$

$$\text{Skew confluence:} \quad \forall e, e_1, e_2. \ e \twoheadrightarrow_{\mathcal{R}} e_1 \wedge e \twoheadrightarrow_{\mathcal{R}} e_2 \implies \exists e'. \ e_1 \twoheadrightarrow_{\mathcal{R}} e' \wedge e_2 \preceq_{\omega_{\mathcal{R}}} e'.$$

These are depicted here as two commutative diagrams, which differ only on the bottom edge:

For each diagram, given e, e_1, e_2 that obey the relationships indicated by all the solid lines, there exists e' such that all relationships indicated by dotted lines are also satisfied.

You can understand skew confluence as follows: if two different reduction paths from e produce terms e_1, e_2, then e_1 can be further reduced to some e' such that all of e_2's permanent structure is present in e', written $e_2 \preceq_{\omega_{\mathcal{R}}} e'$. By "permanent structure" we mean an outer shell of tuples, lambdas, and constants, that will never change no matter how much further reduction takes place. For example, however far we reduce the term $\langle 1, \lambda z.\, e \rangle$, the result will always look like $\langle 1, \lambda z.\, e' \rangle$, where $e \twoheadrightarrow_{\mathcal{R}} e'$. We can formalize the notion of permanent structure by defining an *information content function* $\omega_{\mathcal{R}}(e)$ that replaces all the impermanent bits of e with a new dummy term Ω. Thus $\omega_{\mathcal{R}}(\langle 1, \lambda z.\, x \rangle) = \langle 1, \lambda z.\, \Omega \rangle$. Then $e_2 \preceq_{\omega_{\mathcal{R}}} e'$ if $\omega_{\mathcal{R}}(e_2)$ can be made equal to e' by replacing each occurrence of Ω in $\omega_{\mathcal{R}}(e_2)$ with an (individually-chosen) term.

Consider the even-odd problem discussed in Section 4.1.2.

$$\exists x\, y.\ x = \langle 1, y \rangle;\ y = \lambda z.\, x;\ x$$

$\twoheadrightarrow \exists y.\ y = \lambda z.\, \langle 1, y \rangle;\ \langle 1, y \rangle$	$\twoheadrightarrow \exists x.\ x = \langle 1, \lambda z.\, x \rangle;\ x$
$\rightarrow \exists y.\ y = \lambda z.\, \langle 1, y \rangle;\ \langle 1, \lambda z.\, \langle 1, y \rangle \rangle$	$\rightarrow \exists x.\ x = \langle 1, \lambda z.\, x \rangle;\ \langle 1, \lambda z.\, x \rangle$
$\rightarrow \exists y.\ y = \lambda z.\, \langle 1, y \rangle;\ \langle 1, \lambda z.\, \langle 1, \lambda z.\, \langle 1, y \rangle \rangle \rangle$	$\rightarrow \exists x.\ x = \langle 1, \lambda z.\, x \rangle;\ \langle 1, \lambda z.\, \langle 1, \lambda z.\, x \rangle \rangle$
$\rightarrow \cdots$	$\rightarrow \cdots$

The two columns can never join up, but if you pick any term in either column, there is a term in the other column that has at least as much permanent structure. That in turn means that the terms in the left-hand column are contextually equivalent to those in the right-hand column, because the context can inspect only the permanent structure. This contextual equivalence is the real reason for seeking confluence in the first place.

In $\mathcal{VC}$, the notion of permanent structure must be generalized from an outer shell of tuples, lambdas, and constants to a *sequence* (right-associative tree) of *choices* of such outer shells. For example, however far we reduce the term $\langle 68, \lambda x.\, e_1 \rangle \mid \langle \lambda y.\, e_2, 57, \lambda z.\, e_3 \rangle \mid e_4$, the result will always look like $\langle 68, \lambda x.\, e_1' \rangle \mid \langle \lambda y.\, e_2', 57, \lambda z.\, e_3' \rangle \mid e_4'$, where $e_k \twoheadrightarrow_{\mathcal{R}} e_k'$ for $1 \leqslant k \leqslant 4$; or $\langle 68, \lambda x.\, e_1' \rangle \mid \langle \lambda y.\, e_2', 57, \lambda z.\, e_3' \rangle$ if $e_4 \twoheadrightarrow_{\mathcal{R}} \mathbf{fail}$. In Appendix D we sketch our plan to adapt the proof strategy of Section 4.2 and Appendix C for skew confluence.

5 VARIATIONS AND CHOICES

In a calculus like $\mathcal{VC}$, there is room for many design variations. We discuss some of them here.

5.1 Simplifying the Rules

The rules of Figure 3 are more complicated than we would like, because our proof of confluence requires that the unification rules form a *terminating* rewrite system (Section 4.2). Tantalizingly, we have found simpler versions of four of the unification and elimination rules that we believe are equally expressive and still support (modified versions of) our proofs of termination and confluence:

SUBST$'$	$x = v;\ e \longrightarrow x = v;\ e\{v/x\}$	if $v \neq V[x]$
VAR-SWAP$'$	$y = x;\ e \longrightarrow x = y;\ e$	(unconditionally)
SEQ-SWAP$'$	$eq;\ x = v;\ e \longrightarrow x = v;\ eq;\ e$	(unconditionally)
EQN-ELIM$'$	$\exists x.\ x = v;\ e \longrightarrow e$	if $x \notin \mathrm{fvs}(v, e)$

Each of these rules is simpler than its counterpart in Figure 3. In particular, rules SUBST$'$ and EQN-ELIM$'$ need no context X: instead, they rely on other rules (EQN-FLOAT and SEQ-SWAP$'$) to float the equation of interest upward and to the left. Our testing framework (Section 3.9) has given us strong confidence that the entire set of rules, with these four simplifications, remains confluent; actually *proving* it confluent is work in progress. We speculate that these same simplifications will also support a proof of skew confluence.

5.2 Ordering and Choices

As we discussed in Section 3.5, rule CHOOSE is less than satisfying for two reasons. First, the CX context uses a conservative, syntactic analysis for choice-free expressions; and second, the SX context is needed to force CX to be maximal. A rule like this would be more satisfying:

$$\text{SIMPLER-CHOOSE} \quad CX[\, e_1 \mid e_2 \,] \longrightarrow CX[\, e_1 \,] \mid CX[\, e_2 \,]$$

The trouble with this is that it may change the order of the results (Section 2.3). Another possibility would be to accept that results may come out in the "wrong" order, but have some kind of sorting mechanism to put them back into the "right" order. Something like this:

$$\text{LABELED-CHOOSE} \quad CX[\, e_1 \mid e_2 \,] \longrightarrow CX[\, L \triangleright e_1 \,] \mid CX[\, R \triangleright e_2 \,]$$

Here, the two branches are labeled with L and R. We can add new rules to reorder such labeled expressions, something in the spirit of:

$$\text{SORT} \quad (R \triangleright e_1) \mid (L \triangleright e_2) \longrightarrow (L \triangleright e_2) \mid (R \triangleright e_1)$$

We believe this can be made to work, and it would allow more programs to evaluate, but it adds unwelcome clutter to program terms, and the cure may be worse than the disease. However, the idea inspired our denotational semantics (Appendix E.4), where it seems to work rather beautifully.

5.3 Generalizing one and all

In $\mathcal{VC}$, we introduced **one** and **all** as the primitive choice-consuming operators, and neither is more general than the other, as discussed in Section 2.6. We could have introduced a more general operator **split**[13] as $e ::= \cdots \mid \textbf{split}\{e\}\langle v_1, v_2 \rangle$ and rules:

$$
\begin{array}{llll}
\text{SPLIT-FAIL} & \textbf{split}\{\textbf{fail}\}\langle f, g \rangle & \longrightarrow & f\langle\rangle \\
\text{SPLIT-VALUE} & \textbf{split}\{v\}\langle f, g \rangle & \longrightarrow & g\langle v, \lambda\langle\rangle.\, \textbf{fail}\rangle \\
\text{SPLIT-CHOICE} & \textbf{split}\{v \mid e\}\langle f, g \rangle & \longrightarrow & g\langle v, \lambda\langle\rangle.\, e\rangle
\end{array}
$$

The intuition behind **split** is that it distinguishes a failing computation from one that returns at least one value. If e fails, it calls f; but if e returns at least one value, it passes that to g together with the remaining computation, safely tucked away within a lambda. When adding more effects to $\mathcal{VC}$ (see Appendix F), it is in fact crucial to use **split** to exactly control the order of effects.

Indeed, this is more general, as we can implement **one** and **all** with **split**:

$$
\begin{array}{lll}
\textbf{one}\{e\} \equiv f(x) := \textbf{fail};\ g\langle x, y \rangle := x; & & \textbf{split}\{e\}\langle f, g \rangle \\
\textbf{all}\{e\} \equiv f(x) := \langle\rangle;\quad g\langle x, y \rangle := cons\langle x, \textbf{split}\{y\langle\rangle\}\langle f, g \rangle\rangle;\ \textbf{split}\{e\}\langle f, g \rangle
\end{array}
$$

For this paper, we stuck to the arguably simpler **one** and **all**, to avoid confusing the presentation with these higher-order encodings, but there are no complications using **split** instead.

6 $\mathcal{VC}$ IN CONTEXT: REFLECTIONS AND RELATED WORK

Functional logic programming has a rich literature; excellent starting points are the CACM review article by Antoy and Hanus [2010] and the longer survey by Hanus [2013]. Now that we know what $\mathcal{VC}$ is, we can identify its distinctive features, and compare them to other approaches.

---

[13]The name **split** was inspired by Kiselyov et al. [2005].

6.1 Choice and Non-determinism

A significant difference between our presentation and earlier works is our treatment of choice. Consider an expression like $(3 + (20 \mathbin{|} 30))$. Choice is typically handled by a pair of non-deterministic rewrite rules:

$$e_1 \mathbin{|} e_2 \longrightarrow e_1 \qquad e_1 \mathbin{|} e_2 \longrightarrow e_2$$

So our expression rewrites (non-deterministically) to either $(3 + 20)$ or $(3 + 30)$, and that in turn allows the addition to make progress. Of course, including non-deterministic choice means the rules are non-confluent by construction. Instead, one must generalize to say that a reduction does not change the *set* of results; in the context of lambda calculi, see for example Kutzner and Schmidt-Schauß [1998]; Schmidt-Schauß and Machkasova [2008].

In contrast, our rules never pick one side or the other of a choice. And yet, $(3 + (20 \mathbin{|} 30))$ can still make progress by floating out the choice (rule CHOOSE in Fig. 3), thus $(3 + 20) \mathbin{|} (3 + 30)$. In effect, *choices are laid out in space* (in the syntax of the term), rather than being explored by non-deterministic selection. Rule CHOOSE is not a new idea: it is common in calculi with choice, see *e.g.*, de'Liguoro and Piperno [1995, Section 6.1] and Dal Lago et al. [2020, Section 3], and, more recently, has been used to describe functional logic languages, where it is variously called *bubbling* [Antoy et al. 2007] or *pull-tabbing* [Antoy 2011]. However, laziness imposes significant additional complications, including the need to attach an identifier to each choice.

Recent work by Barenbaum et al. [2020, 2021] describes a functional logic calculus that, like $\mathcal{VC}$, is just an extension of lambda calculus. They deal with lambda, existentials, choice, and equations in a similar way to $\mathcal{VC}$, but lack the **all** and **one** operators which are central to $\mathcal{VC}$. As a result they can always float choice up to the top level of the whole program, and indeed they do so in the very structure of their terms.

6.2 One and all

Logical variables, choice, and equalities are present in many functional logic languages. However, **one** and **all** are distinctive features of $\mathcal{VC}$, with the notable exception of Fresh, a very interesting design introduced in a technical report nearly 40 years ago [Smolka and Panangaden 1985] that also aims to unify functional and logical constructs. Fresh reifies choice into data via *confinement* (corresponding to **one**) and *collection* (corresponding to **all**). However, Fresh differs from $\mathcal{VC}$ in crucial ways. First, it solves equations in a strictly left-to-right fashion, which means that it is not lenient in the sense discussed in Section 3.6. Second, its semantics are presented in an operational fashion with explicit stacks and heaps, in contrast to our focus on developing an equational account of functional logic programming. Finally, Fresh appears not to have been implemented.

Several aspects of **all** and **one** are worth noting. First, **all** *reifies* choice (a control operator) into a tuple (a data structure); for example, **all**$\{1 \mathbin{|} 7 \mathbin{|} 2\}$ returns the tuple $\langle 1, 7, 2 \rangle$. In the other direction, indexing turns a tuple into choice (for example, $\exists i. \langle 1, 7, 2 \rangle(i)$ yields $(1 \mathbin{|} 7 \mathbin{|} 2)$). Other languages can reify choices into a (non-deterministic) list, via an operator called bagof, or a mechanism called *set-functions* in an extension of Curry [Antoy and Hanus 2021, Section 4.2.7], implemented in the Kiel Curry System [Antoy and Hanus 2009; Braßel and Huch 2007, 2009]. But in Curry, this is regarded as a somewhat sophisticated feature, whereas it is part of the foundational fabric of $\mathcal{VC}$. Curry's set-functions need careful explanation about sharing across non-deterministic choices[14], or what is "inside" and "outside" the set function, something that appears as a straightforward consequence of $\mathcal{VC}$'s single rule CHOOSE.

Second, even under the reification of **all**, $\mathcal{VC}$ is *deterministic*. $\mathcal{VC}$ takes pains to maintain order, so that when reifying choice into a tuple, the order of elements in that tuple is completely

---

[14]Again, this complexity in Curry appears to be a consequence of laziness.

determined. This determinism has a price: we have to take care to maintain the left-to-right order of choices (see Section 2.3 and Section 3.5, for example). However, maintaining that order has other payoffs. For example, it is relatively easy to add effects other than choice, including mutable variables and input/output, to $\mathcal{VC}$. To substantiate this claim, Appendix F gives the additional syntax and rewrite rules for mutable variables.

Thirdly, **one** allows us to reify failure: to try something and take different actions depending on whether or not it succeeds. Prolog's "cut" operator [Clocksin and Mellish 2003, Chapter 4] has a similar flavor, and Curry's set-functions allow one to do the same thing.

Finally, **one** and **all** neatly encapsulate the idea of "flexible" *vs.* "rigid" logical variables. As we saw in Section 2.5, logical variables bound outside **one**/**all** cannot be unified inside it; they are "rigid." This notion is nicely captured by the fact that equalities cannot float outside **one** and **all** (Section 3.4).

6.3 The semantics of logical variables

Our logical variables, introduced by ∃, are often called *extra variables* in the literature, because they are typically introduced as variables that appear on the right-hand side of a function definition, but are not bound on the left. For example, in Curry we can write:

```
first x | x =:= (a,b) = a where a,b free
```

Here, a and b are logical variables, not bound on the left; they get their values through unification (written "=:="). In Curry, they are explicitly introduced by the "where a,b free" clause, while in many other papers their introduction is implicit in the top-level rules, simply by not being bound on the left. These extra variables (our logical variables) are at the heart of the "logic" part of functional logic programming.

Constructor-based ReWrite Logic (CRWL) [González-Moreno et al. 1999] is the brand leader for high-level semantics for non-strict, non-deterministic functional logic languages. CRWL is a "big-step" rewrite semantics that rewrites a term to a value in a single step. López-Fraguas et al. [2007] make a powerful case for instead giving the semantics of a functional logic language using "small-step" rewrite rules, more like those of the lambda calculus, that successively rewrite the term, one step at a time, until it reaches a normal form. Their paper does exactly this, and proves equivalence to the CRWL framework. Their key insight (like us, inspired by Ariola et al. [1995]'s formalization of the call-by-need lambda calculus) is to use **let** to make sharing explicit.

However, both CRWL and López-Fraguas *et al.* are in some ways *too* high level: they require something we call *magical rewriting*. A key rewrite rule is this:

$$f(\theta(e_1), \ldots, \theta(e_n)) \longrightarrow \theta(rhs)$$
$$\text{if } (e_1, \ldots, e_n) \longrightarrow rhs \text{ is a top-level function binding, and}$$
$$\theta \text{ is a substitution mapping variables to closed values, s.t. } dom(\theta) = fvs(e_1, \ldots, e_n, rhs)$$

The substitution for the free variables of the left-hand-side can readily be chosen by matching the left-hand-side against the call. But the substitution for the extra variables on the right-hand side must be chosen "magically" [López-Fraguas et al. 2007, Section 7] or clairvoyantly, so as to make the future execution work out. This is admirably high-level because it hides everything about unification, but it is not much help to a programmer trying to understand a program, nor is it directly executable. In a subsequent journal paper, they refine CRWL to avoid magical rewriting by using "let-narrowing" [López-Fraguas et al. 2014, Section 6]; this system looks rather different to ours, especially in its treatment of choice, but is rather close in spirit.

To explain actual execution, the state of the art is described by Albert et al. [2005]. They give both a big-step operational semantics (in the style of Launchbury [1993]), and a small-step operational

semantics. These two approaches both thread a *heap* through the execution, which holds the unification variables and their unification state; the small-step semantics also has a *stack*, to specify the focus of execution. The trouble is that heaps and stacks are difficult to explain to a programmer, and do not make it easy to reason about program equivalence. In addition to this machinery, the model is further complicated with concurrency to account for residuation.

In contrast, our rewrite rules give a complete, executable (*i.e.*, no "magic") account of logical variables and choice, directly as small-step rewrites on the original program, rather than as the evolution of a (heap, control, stack) configuration. Moreover, we have no problem with residuation.

The Escher language [Lloyd 1999] extends a Haskell-like functional base with logical (existentially quantified) variables and constraints to yield an integration of functional and logical constructs, and informally describes a rewriting based evaluation mechanism. However, the language does not have the equivalent of **one** and **all** which reify choice as data, the rules are not precisely formalized, and no confluence result is established for them.

Goffin [Chakravarty et al. 1998] also presents a way to extend a higher-order functional language with existentially quantified logical variables. However, Goffin is rather different from $\mathcal{VC}$ in that it falls in the tradition of Concurrent Constraint Programming [Saraswat and Rinard 1990], where the constraints over logical variables are used to declare data dependencies as a means to co-ordinate efficient concurrent execution of sub-computations, rather than in the tradition of logic programming (and $\mathcal{VC}$) where we seek to declare the properties of solutions. Consequently, in Goffin, the logical variables are never "unified" but instead are "updated" as values are computed in parallel, yielding a completely different model of computation.

6.4 Flat *vs.* Higher Order

When giving the semantics of functional logic languages, a first-order presentation is almost universal. User-defined functions can be defined at top level only, and the names of functions are syntactically distinguished from ordinary variables. As Hanus describes, it is possible to translate a higher-order program into a first-order form using defunctionalization [Hanus 2013, Section 3.3] and a built-in **apply** function. (Hanus does not mention this, but for a language with arbitrarily nested lambdas, one would need to do lambda-lifting [Johnsson 1985] as well; this is perhaps a minor point.) Sadly, this encoding is hardly a natural rendition of the lambda calculus, and it obstructs the goal of using rewrite rules to explain to programmers how their program might behave. In contrast, a strength of our $\mathcal{VC}$ presentation is that it deals natively with the full lambda calculus.

6.5 Intermediate Language

Hanus's *Flat Language* [Albert et al. 2005, Fig 1], FLC, plays the same role as $\mathcal{VC}$: it is a small core language into which a larger surface language can be desugared. There are some common features: variables, literals, constructor applications, and sequencing (written hnf in FLC). However, it seems that $\mathcal{VC}$ has a greater economy of concepts. In particular, FLC has two forms of equality (==) and (=:=), and two forms of case-expression, case and fcase. In each pair, the former suspends if it encounters a logical variable; the latter unifies or narrows respectively. In contrast, $\mathcal{VC}$ has a single equality (=), and the orthogonal **one** construct, to deal with all four concepts.

FLC has let-expressions (let x=e in b) where $\mathcal{VC}$ uses ∃ and (again) unification. FLC also uses the same construct for a different purpose, to bring a logical variable into scope, using the strange binding x=x, thus (let x=x in e). In contrast, ∃x. e seems more direct.

6.6 Comparison with Various Functional Language Extensions to Datalog

Datalog [Ceri et al. 1989] in its purest form is a subset of Prolog [Clocksin and Mellish 2003; Warren et al. 1977] that is carefully limited so that every Datalog program is guaranteed to terminate.

Datalog variants are often used as database query languages and to concisely express fixpoint computations on lattices (for example, static-analysis phases for compilers). Some of these variants support functional programming while preserving the guarantee that every program terminates.

- Datafun [Arntzenius and Krishnaswami 2016] is perhaps closest in spirit to Verse in that it supports higher-order functions as first-class values; it uses a strong type system to prohibit testing functions for equality and to prohibit unbounded recursion.
- Flix [Madsen et al. 2016] makes it easy to work with arbitrary user-defined lattices; programs are guaranteed to terminate as long as (a) all declared lattices are complete and of finite height, and (b) all functions are strict and monotone. Functions are not first-class and their use is syntactically restricted.
- QL [Avgustinov et al. 2016] has an object-oriented, class-based syntax that is then compiled to pure Datalog; this allows the programmer to use inheritance to organize code on related types of database records. In some cases what looks like a "method call" appears to produce multiple values, but this is just a bit of syntactic sugar for expressing relations.
- Formulog [Bembenek et al. 2020] extends Datalog with a first-order functional language (a subset of ML) and an SMT solver. Functions are not first-class; functions, predicates, and variables have separate namespaces.
- Functional IncA [Pacak and Erdweg 2022] is a language that supports first-class higher-order functions over sets and algebraic data types; this language is compiled to pure Datalog. The compilation process eliminates first-class functions through defunctionalization. The paper does not discuss an equality operator or whether it may be applied to functions.

None of these languages supports unbounded computation, a choice operator, or failure.

6.7 Comparison with Icon

There are many obvious similarities between Verse and the Icon programming language [Griswold 1979; Griswold and Griswold 1983, 2002; Griswold et al. 1979, 1981]:

- An expression can (successively) produce any number of values. An expression that produces zero values is said to *fail* [Griswold et al. 1981, §3.1]; an expression that produces at least one value is said to *succeed*.
- The expression $e_1 \mid e_2$ produces all the values of e_1 followed by all the values of e_2.
- There is a way to turn an array (or tuple) a into a sequence of produced values. In Icon, this is written $!a$ [Griswold et al. 1979, §3]; in Verse, a?; in $\mathcal{VC}$, $\exists i. a(i)$.
- Most "scalar" operations (such as addition and comparisons) run through all possible combinations of values of their operand expressions, using a specific left-to-right evaluation order and automatic chronological backtracking.
- Success and failure are used in place of boolean values for control-structure purposes, just as in Section 2.5.
- The "$\mid$" construct is idiomatically used as a *logical or* operation [Griswold et al. 1979, §3].
- There is a control structure that executes a specified expression once for every value produced by another expression. In Icon, this is every e_1 do e_2 and in Verse, it is written for(e_1) do e_2. In each language, the "do e_2" may be omitted to simply evaluate e_1 repeatedly until it has yielded all its values. In Icon, every e may also be written as repeat e.
- It is impossible to name a generator (Icon) or choice (Verse); if e produces multiple values, x := e will provide one value at a time from e to be named by variable x.

But there are also major differences. Icon was designed primarily to use expressions as generators to automatically explore a combinatorial space of possibilities ("goal-directed evaluation"), and secondarily to use success/failure rather than booleans to drive control structure. But in other

respects, *Icon is a fairly conventional imperative language*, relying on side effects (assignments) to process the generated combinations. The designers judged that the interactions of such side effects with completely unrestrained control backtracking would be difficult for programmers to understand [Griswold et al. 1981, §3.1]; therefore, the design of Icon emphasizes limited scopes for control backtracking and tools for controlling the backtracking process [Griswold et al. 1981, §3.3].

In contrast, *Verse is a declarative language* and avoids these difficulties by using a functional logic approach rather than an imperative approach to processing generated combinations:

- While Icon typically processes multiple values from an expression by using *assignment*, Verse typically processes multiple values by using *equations* (which are then *solved*).
- Verse also has a concise way to turn a finite sequence of multiple values into a tuple directly. For example, to make variable a refer to an array containing all values generated by expression e, code such as the following (using a repeat loop containing an assignment) is idiomatic in Icon [Griswold et al. 1979, §8]:

        a := array 0 string; i := 0; repeat a[i+] := e; close(a)

  In Verse, a = for{e} does the job; in $\mathcal{VC}$, $a = \mathbf{all}\{e\}$ is all it takes.
- Backtracking in Icon is "only control backtracking"; side effects, such as assignments, are not undone [Griswold et al. 1981, §3.1].
- Both languages have an implicit "cut" (permanent acceptance of the first produced value) after the predicate part of an **if-then-else**, but Icon furthermore has an implicit cut at each statement end (semicolon or end of line) [Griswold et al. 1981, §3.1], each closing brace "}", and most keywords [Icon PC 1980].

7   LOOKING BACK, LOOKING FORWARD

We believe that this is the first presentation of a functional logic language as a deterministic rewrite system. A rewrite system has the advantage (compared to more denotational, or more operational, methods) that it is sufficiently low-level to capture the *computational model* of the language, and yet sufficiently high-level to be *illuminating* to a programmer or compiler writer. Our focus on rewriting as a way to define the semantics has forced us to focus on confluence, a syntactic property that is stronger (and hence more delicate and harder to prove) than the contextual equivalence that is all we really need. That in turn led us to study the elegant and ingenious notion of skew confluence, which has been barely revisited during the last 20 years, but which we believe deserves a wider audience.

We have much left to do. The full Verse language has statically checked types. In the dynamic semantics, the types can be represented by partial identity functions—identity for the values of the type, and **fail** otherwise. This gives a distinctive new perspective on type systems, one that we intend to develop in future work. The full Verse language also has a statically checked effect system, including both mutable references and input/output. All these effects must be *transactional*; for example, when the condition of an **if** fails, any store effects in the condition must be rolled back. We have preliminary reduction rules for references see Appendix F.

REFERENCES

Elvira Albert, Michael Hanus, Frank Huch, Javier Oliver, and German Vidal. 2005. Operational semantics for declarative multi-paradigm languages. *Journal of Symbolic Computation* 40, 1 (2005), 795–829. https://doi.org/10.1016/j.jsc.2004.01.001 Reduction Strategies in Rewriting and Programming special issue.

Sergio Antoy. 2011. On the Correctness of Pull-Tabbing. *Theory and Practice of Logic Programming* 11, 4-5 (July 2011), 713–730. https://doi.org/10.1017/S1471068411000263

Sergio Antoy, Daniel W. Brown, and Su-Hui Chiang. 2007. Lazy Context Cloning for Non-Deterministic Graph Rewriting. *Electronic Notes in Theoretical Computer Science* 176, 1 (May 2007), 3–23. https://doi.org/10.1016/j.entcs.2006.10.026 Proceedings of the Third International Workshop on Term Graph Rewriting (TERMGRAPH 2006).

Sergio Antoy and Michael Hanus. 2009. Set Functions for Functional Logic Programming. In *Proceedings of the 11th ACM SIGPLAN Conference on Principles and Practice of Declarative Programming* (Coimbra, Portugal) *(PPDP '09)*. Association for Computing Machinery, New York, NY, USA, 73–82. https://doi.org/10.1145/1599410.1599420

Sergio Antoy and Michael Hanus. 2010. Functional Logic Programming. *Commun. ACM* 53, 4 (April 2010), 74–85. https://doi.org/10.1145/1721654.1721675

Sergio Antoy and Michael Hanus. 2021. *Curry: A Tutorial Introduction.* Technical Report. Kiel University (Christian-Albrechts-Universität zu Kiel). https://web.archive.org/web/20220121070135/https://www.informatik.uni-kiel.de/~curry/tutorial/tutorial.pdf

Zena M. Ariola and Stefan Blom. 2002. Skew confluence and the lambda calculus with letrec. *Annals of Pure and Applied Logic* 117, 1 (2002), 95–168. https://doi.org/10.1016/S0168-0072(01)00104-X

Zena M. Ariola and Jan Willem Klop. 1994. Cyclic lambda graph rewriting. In *Proceedings of the Ninth Annual IEEE Symposium on Logic in Computer Science (LICS '94)*. IEEE, 416–425. https://doi.org/10.1109/LICS.1994.316066

Zena M. Ariola and Jan Willem Klop. 1997. Lambda Calculus with Explicit Recursion. *Information and Computation* 139, 2 (Dec. 1997), 154–233. https://doi.org/10.1006/inco.1997.2651

Zena M. Ariola, John Maraist, Martin Odersky, Matthias Felleisen, and Philip Wadler. 1995. A Call-by-Need Lambda Calculus. In *Proceedings of the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Francisco, California, USA) *(POPL '95)*. Association for Computing Machinery, New York, NY, USA, 233246. https://doi.org/10.1145/199448.199507

Michael Arntzenius and Neelakantan R. Krishnaswami. 2016. Datafun: A Functional Datalog. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming* (Nara, Japan) *(ICFP 2016)*. Association for Computing Machinery, New York, NY, USA, 214–227. https://doi.org/10.1145/2951913.2951948

Andrea Asperti and Stefano Guerrini. 1999. *The Optimal Implementation of Functional Programming Languages.* Cambridge University Press.

Pavel Avgustinov, Oege de Moor, Michael Peyton Jones, and Max Schäfer. 2016. QL: Object-oriented Queries on Relational Data. In *30th European Conference on Object-Oriented Programming (ECOOP 2016) (Leibniz International Proceedings in Informatics (LIPIcs))*, Shriram Krishnamurthi and Benjamin S. Lerner (Eds.), Vol. 56. Schloss Dagstuhl—Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2:1–2:25. https://doi.org/10.4230/LIPIcs.ECOOP.2016.2

Pablo Barenbaum, Federico Lochbaum, and Mariana Milicich. 2020. Semantics of a Relational λ-Calculus. In *Theoretical Aspects of Computing (ICTAC 2020)*, Violet Ka I Pun, Volker Stolz, and Adenilso Simao (Eds.). Springer International Publishing, Cham, 242–261. https://doi.org/10.1007/978-3-030-64276-1_13

Pablo Barenbaum, Federico Lochbaum, and Mariana Milicich. 2021. Semantics of a Relational λ-Calculus (Extended Version), version 4. *CoRR* abs/2009.10929 (28 Feb. 2021), 51. arXiv:2009.10929 https://arxiv.org/abs/2009.10929

H. P. (Hendrik Pieter) Barendregt. 1984. *The Lambda Calculus: Its Syntax and Semantics* (revised ed.). Studies in Logic and the Foundations of Mathematics, Vol. 103. North-Holland (Elsevier Science Publishers), Amsterdam.

Aaron Bembenek, Michael Greenberg, and Stephen Chong. 2020. Formulog: Datalog for SMT-Based Static Analysis. *Proc. ACM Programming Languages* 4, OOPSLA, Article 141 (Nov. 2020), 31 pages. https://doi.org/10.1145/3428209

Bernd Braßel, Michael Hanus, and Frank Huch. 2004a. Encapsulating Non-Determinism in Functional Logic Computations [extended journal version]. *Journal of Functional and Logic Programming* 2004, 6 (Dec. 2004), 28. http://danae.uni-muenster.de/lehre/kuchen/JFLP/articles/2004/S04-01/A2004-06/JFLP-A2004-06.pdf Special Issue 1. Archived at https://web.archive.org/web/20060505093657/http://danae.uni-muenster.de/lehre/kuchen/JFLP/articles/2004/S04-01/A2004-06/JFLP-A2004-06.pdf.

Bernd Braßel, Michael Hanus, and Frank Huch. 2004b. Encapsulating Non-Determinism in Functional Logic Computations [workshop version]. In *13th International Workshop on Functional and (Constraint) Logic Programming (WFLP'04)*. 74–90. Full proceedings available at http://www-i2.informatik.rwth-aachen.de/WFLP04/WFLP-proceedings.pdf. RWTH Aachen, Department of Computer Science, Technical Report AIB-2004-05. Archived at https://web.archive.org/web/20040713053113/http://www-i2.informatik.rwth-aachen.de/WFLP04/WFLP-proceedings.pdf.

Bernd Braßel and Frank Huch. 2007. On a Tighter Integration of Functional and Logic Programming. In *5th Asian Symposium on Programming Languages and Systems (APLAS 2007) (LNCS 4807)*, Zhong Shao (Ed.). Springer-Verlag, Berlin, Heidelberg,

122–138. https://doi.org/10.1007/978-3-540-76637-7_9

Bernd Braßel and Frank Huch. 2009. The Kiel Curry System KiCS. In *Applications of Declarative Programming and Knowledge Management (LNAI 5437)*, Dietmar Seipel, Michael Hanus, and Armin Wolf (Eds.). Springer-Verlag, Berlin, Heidelberg, 195–205. https://doi.org/10.1007/978-3-642-00675-3_13

Stefano Ceri, Georg Gottlob, and Letizia Tanca. 1989. What You Always Wanted to Know about Datalog (and Never Dared to Ask). *IEEE Transactions on Knowledge and Data Engineering* 1, 1 (March 1989), 146–166. https://doi.org/10.1109/69.43410 A comprehensive survey with an extensive bibliography.

Manuel M.T. Chakravarty, Yike Guo, Martin Köhler, and Hendrik C.R. Lock. 1998. GOFFIN: Higher-order Functions Meet Concurrent Constraints. *Science of Computer Programming* 30, 1 (June 1998), 157–199. https://doi.org/10.1016/S0167-6423(97)00010-5

Jan Christiansen, Daniel Seidel, and Janis Voigtländer. 2011. An Adequate, Denotational, Functional-Style Semantics for Typed FlatCurry. In *Functional and Constraint Logic Programming: 19th International Workshop, WFLP 2010*, Julio Mariño (Ed.). Springer-Verlag, Berlin, Heidelberg, 119–136. https://doi.org/10.1007/978-3-642-20775-4_7

Koen Claessen and John Hughes. 2000. QuickCheck: A Lightweight Tool for Random Testing of Haskell Programs. In *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming* (Montreal, Canada) *(ICFP '00)*. Association for Computing Machinery, New York, NY, USA, 268–279. https://doi.org/10.1145/351240.351266

William F. Clocksin and Christopher S. Mellish. 2003. *Programming in Prolog (Using the ISO Standard)* (fifth ed.). Springer-Verlag New York Berlin Heidelberg.

Ugo Dal Lago, Giulio Guerrieri, and Willem Heijltjes. 2020. Decomposing Probabilistic Lambda-Calculi. In *Foundations of Software Science and Computation Structures: 23rd International Conference (FoSSaCS'20) (LNCS 12077)*, Jean Goubault-Larrecq and Barbara König (Eds.). Springer International, 136–156. https://doi.org/10.1007/978-3-030-45231-5_8

Ugo de'Liguoro and Adolfo Piperno. 1995. Nondeterministic Extensions of Untyped λ-Calculus. *Information and Computation* 122, 2 (1995), 149–177. https://doi.org/10.1006/inco.1995.1145

Matthias Felleisen, Robert Bruce Findler, and Matthew Flatt. 2009. *Semantics Engineering with PLT Redex*. MIT Press, Cambridge, Massachusetts, USA. https://mitpress.mit.edu/9780262062756/semantics-engineering-with-plt-redex/

Matthias Felleisen and Daniel P. Friedman. 1986. Control Operators, the SECD Machine, and the λ-Calculus. In *Formal Description of Programming Concepts III: Proceedings of the IFIP TC 2/WG 2.2 Working Conference* (Ebberup, Denmark). Elsevier Science Publishers (North-Holland), 193–217. https://web.archive.org/web/20220709064643/https://www.cs.tufts.edu/~nr/cs257/archive/matthias-felleisen/cesk.pdf

Matthias Felleisen, Daniel P. Friedman, Eugene Kohlbecker, and Bruce Duba. 1987. A Syntactic Theory of Sequential Control. *Theoretical Computer Science* 52, 3 (1987), 205–237. https://doi.org/10.1016/0304-3975(87)90109-5

J.C. González-Moreno, M.T. Hortalá-González, F.J. López-Fraguas, and M. Rodríguez-Artalejo. 1999. An approach to declarative programming based on a rewriting logic. *The Journal of Logic Programming* 40, 1 (July 1999), 47–87. https://doi.org/10.1016/S0743-1066(98)10029-8

Ralph E. Griswold. 1979. *User's Manual for the Icon Programming Language*. Technical Report TR 78-14. Department of Computer Science, University of Arizona. https://www2.cs.arizona.edu/icon/ftp/doc/tr78_14.pdf

Ralph E. Griswold and Madge T. Griswold. 1983. *The Icon Programming Language*. Prentice-Hall, Englewood Cliffs, New Jersey.

Ralph E. Griswold and Madge T. Griswold. 2002. *The Icon Programming Language* (third ed.). Peer-to-Peer Communications. https://web.archive.org/web/20040723085807/https://www2.cs.arizona.edu/icon/ftp/doc/lb1up.pdf

Ralph E. Griswold, David R. Hanson, and John T. Korb. 1979. The Icon Programming Language: An Overview. *SIGPLAN Notices* 14, 4 (April 1979), 18–31. https://doi.org/10.1145/988078.988082

Ralph E. Griswold, David R. Hanson, and John T. Korb. 1981. Generators in Icon. *ACM Trans. Programming Languages and Systems* 3, 2 (April 1981), 144–161. https://doi.org/10.1145/357133.357136

Michael Hanus. 2013. Functional Logic Programming: From Theory to Curry. In *Programming Logics: Essays in Memory of Harald Ganzinger*, Andrei Voronkov and Christoph Weidenbach (Eds.). LNCS, Vol. 7797. Springer, Berlin, Heidelberg, 123–168. https://doi.org/10.1007/978-3-642-37651-1_6

Michael Hanus, Sergio Antoy, Bernd Braßel, Herbert Kuchen, Francisco J. López-Fraguas, Wolfgang Lux, Juan José Moreno Navarro, Björn Peemöller, and Frank Steiner. 2016. *Curry: An Integrated Functional Logic Language (Version 0.9.0)*. Technical Report. University of Kiel. https://web.archive.org/web/20161020144634/https://www-ps.informatik.uni-kiel.de/currywiki/_media/documentation/report.pdf

J. Roger Hindley. 1964. *The Church-Rosser Property and a Result in Combinatory Logic*. Ph.D. Dissertation. University of Newcastle-upon-Tyne, United Kingdom.

Gérard Huet. 1980. Confluent Reductions: Abstract Properties and Applications to Term Rewriting Systems. *J. ACM* 27, 4 (Oct. 1980), 797–821. https://doi.org/10.1145/322217.322230

Icon PC 1980. Programming Corner from Icon Newsletter 4. https://www2.cs.arizona.edu/icon/progcorn/pc_inl04.htm

Thomas Johnsson. 1985. Lambda lifting: Transforming programs to recursive equations. In *Functional Programming Languages and Computer Architecture*, Jean-Pierre Jouannaud (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 190–203.

Oleg Kiselyov, Chung-chieh Shan, Daniel P. Friedman, and Amr Sabry. 2005. Backtracking, interleaving, and terminating monad transformers: (functional pearl). In *Proceedings of the 10th ACM SIGPLAN International Conference on Functional Programming, ICFP 2005, Tallinn, Estonia, September 26-28, 2005*, Olivier Danvy and Benjamin C. Pierce (Eds.). ACM, 192–203. https://doi.org/10.1145/1086365.1086390

Arne Kutzner and Manfred Schmidt-Schauß. 1998. A Non-Deterministic Call-by-Need Lambda Calculus. In *Proceedings of the Third ACM SIGPLAN International Conference on Functional Programming* (Baltimore, Maryland, USA) *(ICFP '98)*. Association for Computing Machinery, New York, NY, USA, 324–335. https://doi.org/10.1145/289423.289462

John Lamping. 1990. An Algorithm for Optimal Lambda Calculus Reduction. In *Proceedings of the 17th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Francisco, California, USA) *(POPL '90)*. Association for Computing Machinery, New York, NY, USA, 16–30. https://doi.org/10.1145/96709.96711

John Launchbury. 1993. A Natural Semantics for Lazy Evaluation. In *Proceedings of the 20th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Charleston, South Carolina, USA) *(POPL '93)*. Association for Computing Machinery, New York, NY, USA, 144–154. https://doi.org/10.1145/158511.158618

Jean-Jacques Lévy. 1976. An Algebraic Interpretation of the $\lambda\beta$K-Calculus; and an Application of a Labelled λ-Calculus. *Theoretical Computer Science* 2, 1 (June 1976), 97–114. https://doi.org/10.1016/0304-3975(76)90009-8

Jean-Jacques Lévy. 1978. *Réductions Correctes et Optimales dans le Lambda-calcul*. Ph.D. Dissertation. Université Paris VII. https://web.archive.org/web/20051016053439/http://pauillac.inria.fr/~levy/pubs/78phd.pdf

John W. Lloyd. 1999. Programming in an Integrated Functional and Logic Language. *Journal of Functional and Logic Programming* 1999, 3 (March 1999). http://danae.uni-muenster.de/lehre/kuchen/JFLP//articles/1999/A99-03/JFLP-A99-03.pdf Archived at https://web.archive.org/web/20040627000956/http://danae.uni-muenster.de/lehre/kuchen/JFLP/articles/1999/A99-03/JFLP-A99-03.pdf. Also available at https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=31d9a554f61ced1304ecd0ded9fca6fee2173b99.

Francisco Javier López-Fraguas, Enrique Martin-Martin, Juan Rodríguez-Hortalá, and Jaime Sánchez-Hernández. 2014. Rewriting and narrowing for constructor systems with call-time choice semantics. *Theory and Practice of Logic programming* 14, 2 (March 2014), 165–213. https://doi.org/doi:10.1017/S1471068412000373 Published online on 30 October 2012.

Francisco J. López-Fraguas, Juan Rodríguez-Hortalá, and Jaime Sánchez-Hernández. 2007. A Simple Rewrite Notion for Call-Time Choice Semantics. In *Proceedings of the 9th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming* (Wroclaw, Poland) *(PPDP '07)*. Association for Computing Machinery, New York, NY, USA, 197–208. https://doi.org/10.1145/1273920.1273947

Magnus Madsen, Ming-Ho Yee, and Ondřej Lhoták. 2016. From Datalog to Flix: A Declarative Language for Fixed Points on Lattices. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Santa Barbara, California, USA) *(PLDI '16)*. Association for Computing Machinery, New York, NY, USA, 194–208. https://doi.org/10.1145/2908080.2908096

André Pacak and Sebastian Erdweg. 2022. Functional Programming with Datalog. In *36th European Conference on Object-Oriented Programming (ECOOP 2022) (Leibniz International Proceedings in Informatics (LIPIcs))*, Karim Ali and Jan Vitek (Eds.), Vol. 222. Schloss Dagstuhl—Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 7:1–7:28. https://doi.org/10.4230/LIPIcs.ECOOP.2022.7

Simon L. Peyton Jones. 1992. Implementing Lazy Functional Languages on Stock Hardware: The Spineless Tagless G-machine. *Journal of Functional Programming* 2, 2 (April 1992), 127–202. https://doi.org/10.1017/S0956796800000319 Also available at https://www.microsoft.com/en-us/research/publication/implementing-lazy-functional-languages-on-stock-hardware-the-spineless-tagless-g-machine/.

John C. Reynolds. 1972. Definitional Interpreters for Higher-Order Programming Languages. In *Proceedings of the ACM Annual Conference—Volume 2* (Boston, Massachusetts, USA) *(ACM '72)*. Association for Computing Machinery, New York, NY, USA, 717–740. https://doi.org/10.1145/800194.805852

J. A. Robinson. 1965. A Machine-Oriented Logic Based on the Resolution Principle. *J. ACM* 12, 1 (Jan. 1965), 23–41. https://doi.org/10.1145/321250.321253

Amr Sabry and Matthias Felleisen. 1992. Reasoning about Programs in Continuation-Passing Style. In *Proceedings of the 1992 ACM Conference on LISP and Functional Programming, LFP '92*. ACM.

Vijay A. Saraswat and Martin C. Rinard. 1990. Concurrent Constraint Programming. In *Proceedings of the 17th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Francisco, California, USA) *(POPL '90)*, Frances E. Allen (Ed.). ACM Press, 232–245. https://doi.org/10.1145/96709.96733

Klaus E. Schauser and Seth C. Goldstein. 1995. How Much Non-strictness Do Lenient Programs Require?. In *Proceedings of the Seventh International Conference on Functional Programming Languages and Computer Architecture* (La Jolla, California,

USA) *(FPCA '95)*. Association for Computing Machinery, New York, NY, USA, 216–225. https://doi.org/10.1145/224164.224208

Manfred Schmidt-Schauß and Elena Machkasova. 2008. A Finite Simulation Method in a Non-deterministic Call-by-Need Lambda-Calculus with Letrec, Constructors, and Case. In *19th International Conference on Rewriting Techniques and Applications (RTA '08) (LNCS 5117)*. Springer, Berlin, Heidelberg, 321–335. https://doi.org/10.1007/978-3-540-70590-1_22

Gert Smolka and Prakash Panangaden. 1985. *FRESH: A Higher-Order Language with Unification and Multiple Results*. Technical Report TR 85-685. Cornell University, Ithaca, New York, USA. https://hdl.handle.net/1813/6525

Guy Lewis Steele Jr. 1978. *Rabbit: A Compiler for Scheme*. Technical Report 474. Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA. https://web.archive.org/web/20211108071621/https://dspace.mit.edu/bitstream/handle/1721.1/6913/AITR-474.pdf Master's Dissertation.

David H D Warren, Luis M. Pereira, and Fernando Pereira. 1977. Prolog - The Language and Its Implementation Compared with Lisp. In *Proceedings of the 1977 Symposium on Artificial Intelligence and Programming Languages* (Rochester, New York, USA). Association for Computing Machinery, New York, NY, USA, 109–115. https://doi.org/10.1145/800228.806939

A   EXAMPLE

A complete reduction sequence for a small example can be found in Fig. 6. This example shows how constraining the output of a function call can constrain the argument. While most of the reductions are administrative in nature, these are the highlights: At ① the *swap* function is inlined so that at ② a β-reduction can happen. Step ③ inlines the argument, and ④ does the matching of the tuple. At ⑤ and ⑥ the actual numbers are inlined.

	$swap\langle x, y\rangle := \langle y, x\rangle;\ \exists p.\ swap(p) = \langle 2, 3\rangle;\ p$
$\longrightarrow\{\text{DESUGAR}\}$	$\exists swap.\ swap = (\lambda xy.\ \exists x\,y.\ \langle x, y\rangle = xy;\ \langle y, x\rangle);\ \exists p\,t.\ t = swap(p);\ t = \langle 2, 3\rangle;\ p$
① $\longrightarrow\{\text{SUBST,EQN-ELIM}\}$	$\exists p\,t.\ t = (\lambda xy.\ \exists x\,y.\ \langle x, y\rangle = xy;\ \langle y, x\rangle)(p);\ t = \langle 2, 3\rangle;\ p$
$\longrightarrow\{\text{SUBST,EQN-ELIM}\}$	$\exists p.\ (2, 3) = (\lambda xy.\ \exists x\,y.\ \langle x, y\rangle = xy;\ \langle y, x\rangle)(p);\ p$
② $\longrightarrow\{\text{APP-BETA}\}$	$\exists p.\ (2, 3) = (\exists xy.\ xy = p;\ \exists x\,y.\ \langle x, y\rangle = xy;\ \langle y, x\rangle);\ p$
$\longrightarrow\{\text{EXI-FLOAT}\}$	$\exists p\,xy.\ (2, 3) = ((xy = p;\ \exists x\,y.\ \langle x, y\rangle = xy;\ \langle y, x\rangle));\ p$
③ $\longrightarrow\{\text{SUBST,EQN-ELIM}\}$	$\exists p.\ (2, 3) = (\exists x\,y.\ \langle x, y\rangle = p;\ \langle y, x\rangle);\ p$
$\longrightarrow\{\text{EXI-FLOAT,EXI-FLOAT}\}$	$\exists p\,x\,y.\ (2, 3) = ((\langle x, y\rangle = p;\ \langle y, x\rangle));\ p$
$\longrightarrow\{\text{EQN-FLOAT,SEQ-ASSOC}\}$	$\exists p\,x\,y.\ \langle x, y\rangle = p;\ (2, 3) = \langle y, x\rangle;\ p$
$\longrightarrow\{\text{HNF-SWAP}\}$	$\exists p\,x\,y.\ p = \langle x, y\rangle;\ (2, 3) = \langle y, x\rangle;\ p$
$\longrightarrow\{\text{SUBST,EQN-ELIM}\}$	$\exists x\,y.\ (2, 3) = \langle y, x\rangle;\ \langle x, y\rangle$
④ $\longrightarrow\{\text{U-TUP,SEQ-ASSOC}\}$	$\exists x\,y.\ 2 = y;\ 3 = x;\ \langle x, y\rangle$
$\longrightarrow\{\text{HNF-SWAP}\}$	$\exists x\,y.\ y = 2;\ 3 = x;\ \langle x, y\rangle$
⑤ $\longrightarrow\{\text{SUBST,EQN-ELIM}\}$	$\exists x.\ 3 = x;\ \langle x, 2\rangle$
$\longrightarrow\{\text{HNF-SWAP}\}$	$\exists x.\ x = 3;\ \langle x, 2\rangle$
⑥ $\longrightarrow\{\text{SUBST,EQN-ELIM}\}$	$\langle 3, 2\rangle$

Fig. 6. A sample reduction sequence

B CONFLUENCE: PRELIMINARIES

B.1 Reduction relations

Definition B.1 (Binary relations). A *binary relation* is a set of pairs of related items; if R is a relation, then we may write $a\,R\,b$ to mean $(a, b) \in R$.

Definition B.2 (Prototype reduction relations and rewrite rules). Let $\widehat{\mathcal{R}}$ be any binary relation on a set of tree-structured *terms*, such as the terms generated by some BNF grammar; we sometimes refer to $\widehat{\mathcal{R}}$ as a *prototype reduction relation*.

Often a prototype reduction relation is specified by a *rewrite rule* of the form $\alpha \to \beta$, which indicates that for any substitution σ that consistently instantiates all the metavariables (BNF nonterminals) in α and β, $(\sigma(\alpha), \sigma(\beta))$ is a member of the prototype reduction relation. A prototype reduction relation may also be specified by a set of rewrite rules, in which case the prototype reduction relation is the union of the prototype reduction relations specified by the individual rewrite rules.

Definition B.3 (Reduction relations). A *reduction relation* $\mathcal{R}$ is the *compatible closure* of some prototype reduction relation $\widehat{\mathcal{R}}$; compatibility means that, for any context E and any two terms M and N, if $(M, N) \in \mathcal{R}$ then $(E[M], E[N]) \in \mathcal{R}$. Because most of the relations we consider here are compatible, we find it more convenient to use a hat over relation symbol to indicate that it may *not* be compatible, rather than using some special mark to indicate that a relation *is* compatible or to indicate the taking of a compatible closure.

Definition B.4 (Derived relations). For any relation—but typically for a reduction relation, so we will call it $\mathcal{R}$ here—we write $\mathcal{R}^k$ for the composition of k copies of $\mathcal{R}$ and $\mathcal{R}^*$ for the reflexive and transitive closure of $\mathcal{R}$, *i.e.* $\mathcal{R}^* \equiv \cup_{0 \leqslant k} \mathcal{R}^k$. We write

- $a \to_{\mathcal{R}} b$ (*a steps to b*) if $(a, b) \in \mathcal{R}$,
- $a \xrightarrow{\epsilon}_{\mathcal{R}} b$ (*a skips to b*) if $a \equiv b$ or $(a, b) \in \mathcal{R}$,
- $a \twoheadrightarrow_{\mathcal{R}} b$ (*a reduces to b*) if $(x, y) \in \mathcal{R}^*$.
- $a \xrightarrow{k}_{\mathcal{R}} b$ (*a k-steps to b*) if $(a, b) \in \mathcal{R}^k$, and

Sometimes we use this same notation and terminology with a prototype reduction relation $\widehat{\mathcal{R}}$, thus for example $a \to_{\widehat{\mathcal{R}}} b$. In such a case, the arrow indicates rewriting of the entire term a (at the root), and not of some subterm of a.

Definition B.5 (Size). The *size* of a reduction $a \twoheadrightarrow b$ is the smallest i such that $a \xrightarrow{i} b$.

Definition B.6 (Normal Forms). A term a is an $\mathcal{R}$-*Normal Form* if there does not exist any b such that $a \to_{\mathcal{R}} b$.

For clarity, we will omit the subscript $\mathcal{R}$ when it is clear from the context.

B.2 Confluence

Definition B.7 (Diamond Property). A reduction relation satisfies the *diamond* property if whenever $a \to b$ and $a \to c$, there is a d such that $b \to d$ and $c \to d$.

Definition B.8 (Confluence). Two terms b, c can be $\mathcal{R}$-*joined* written $b \downarrow_{\mathcal{R}} c$, if there is a d such that $b \twoheadrightarrow_{\mathcal{R}} d$ and $c \twoheadrightarrow_{\mathcal{R}} d$. A reduction relation $\mathcal{R}$ is *confluent* if whenever $a \twoheadrightarrow_{\mathcal{R}} b$ and $a \twoheadrightarrow_{\mathcal{R}} c$, we have $b \downarrow_{\mathcal{R}} c$.

Definition B.9 (Local Confluence). A reduction relation $\mathcal{R}$ is *locally confluent* if whenever $a \to_{\mathcal{R}} b$ and $a \to_{\mathcal{R}} c$, we have $b \downarrow_{\mathcal{R}} c$.

Fig. 7. **Diamond Property (L), Local Confluence (M), and Confluence (R)**

Fig. 8. **Strong Confluence**

LEMMA B.10 (DIAMOND [BARENDREGT 1984]). *If $\mathcal{R}$ satisfies the diamond property then $\mathcal{R}$ is confluent.*

LEMMA B.11 (UNICITY [BARENDREGT 1984]). *If $\mathcal{R}$ is confluent then every term reduces to at most one normal form.*

LEMMA B.12 (CLOSURE [BARENDREGT 1984]). *If $\mathcal{R}$ is confluent then $\mathcal{R}^*$ is confluent.*

Definition B.13 (Noetherian Reduction). A reduction relation $\mathcal{R}$ is *Noetherian* if there is no infinite sequence $a_0 \rightarrow_\mathcal{R} a_1 \rightarrow_\mathcal{R} \ldots \rightarrow_\mathcal{R} a_n \rightarrow_\mathcal{R} \ldots$.

The following result is known as Newman's Lemma [Barendregt 1984; Huet 1980].

LEMMA B.14 (NEWMAN'S LEMMA). *If $\mathcal{R}$ is locally confluent and Noetherian then $\mathcal{R}$ is confluent.*

Definition B.15 (Strong Confluence). A reduction relation is *strongly confluent* if whenever $a \rightarrow b$ and $a \rightarrow c$, either $b \twoheadrightarrow c$ or there is a d such that $b \twoheadrightarrow d$ and $c \rightarrow d$, as shown in Fig. 8, where the ϵ label indicates 0 or 1 step.

LEMMA B.16 ([HUET 1980, LEMMA 2.5]). *If $\mathcal{R}$ is strongly confluent then $\mathcal{R}$ is confluent.*

B.3 Commutativity

Definition B.17 (Commutativity). A reduction relation R *commutes* with S if for all terms a, b, c such that $a \twoheadrightarrow_R b$ and $a \twoheadrightarrow_S c$ there exists d such that $b \twoheadrightarrow_S d$ and $c \twoheadrightarrow_R d$, as illustrated on the left in Fig. 9.

Definition B.18 (Strong commutativity). A reduction relation R *strongly commutes* with S if for all terms a, b, c such that $a \rightarrow_R b$ and $a \rightarrow_R c$ there exists d such that $b \twoheadrightarrow_S d$ and $c \rightarrow_R d$, as illustrated in the middle in Fig. 9.

Note that if R strongly commutes with itself then, by Definition B.7, R has the diamond property.

LEMMA B.19 (STRONG-COMMUTATIVITY). *If R strongly commutes with S then R commutes with S.*

Fig. 9. **Commutativity (L), Strong Commutativity (C), ∗-Commutativity (R)**

PROOF. Via the following "chase" diagram (probably well known?)

□

LEMMA B.20 (UNION). *If R and S_1 commute and R and S_2 commute then R and $S_1 \cup S_2$ commute.*

PROOF. Via the following chase diagram (probably well known?)

□

Definition B.21 (Postpones). A reduction relation R *strongly postpones* after S if $e \to_R \cdot \to_S e'$ implies $e \twoheadrightarrow_S \cdot \to_R e'$.

LEMMA B.22 ([HINDLEY 1964]). *If R strongly postpones after S then if $e \twoheadrightarrow_{R \cup S} e'$ then $e \twoheadrightarrow_S \cdot \twoheadrightarrow_R e'$.*

Definition B.23 (Hops). A reduction relation R *hops* after S if $e \to_R \cdot \to_S e'$ implies there is an e'' such that $e' \twoheadrightarrow_R e''$ and $e \to_S \cdot \twoheadrightarrow_R e''$.

LEMMA B.24. *If R is confluent and hops after S then*

Proof. By induction on size of $\twoheadrightarrow_R$.

Base case By definition of *hops after*.

Inductive case Assume the lemma for reductions of size upto k and complete the proof via the following diagram where (1) is from the induction hypothesis, (2) is from the definition of hops over, and (3) is from the assumption that R is confluent.

□

LEMMA B.25. *If R is confluent and hops after S then*

Proof. By induction on the number of S steps in $\twoheadrightarrow_{R \cup S}$, via the following diagram.

□

Definition B.26 (half-commutes). A reduction relation R *half-commutes* with S if whenever $e \rightarrow_R e_1$ and $e \rightarrow_S e_2$ there exists e' such that $e_2 \twoheadrightarrow_R e'$ and $e_1 \rightarrow_{S\epsilon} \cdot \twoheadrightarrow_R e'$.

LEMMA B.27. *If R is confluent and R half-commutes with S then R commutes with $R \cup S$.*

Proof. Via the following diagram, where: (1) is R is confluent, and (2) is Lemma B.28.

□

LEMMA B.28. *If R is confluent and R half-commutes with S then if $e \twoheadrightarrow_R e_1$ and $e \twoheadrightarrow_S e_2$ then exists e' such that $e_1 \twoheadrightarrow_{R \cup S} e'$ and $e_2 \twoheadrightarrow_R e'$.*

$$
\begin{array}{ccc}
e & \xrightarrow{\;\;S^k\;\;} & e_2 \\
{\scriptstyle R}\downarrow & & \downarrow{\scriptstyle R} \\
e_1 & \cdots\!\cdots\!\cdots\!\twoheadrightarrow_{(R^* \cup S)^k} & e'
\end{array}
$$

PROOF. By repeatedly tiling (1) Lemma B.30 as follows

$$
\begin{array}{ccccccccc}
\bullet & \xrightarrow{S_1} & \bullet & \xrightarrow{S_2} & \bullet & \dashrightarrow^{S} & \bullet & \xrightarrow{S_k} & \bullet \\
{\scriptstyle R}\downarrow & & {\scriptstyle R}\downarrow & & {\scriptstyle R}\downarrow & & {\scriptstyle R}\downarrow & & \downarrow{\scriptstyle R} \\
\bullet & \twoheadrightarrow_{(R^* \cup S)_1} & (1) & \twoheadrightarrow_{(R^* \cup S)_2} & (1) & \twoheadrightarrow_{(R^* \cup S)} & (1) & \twoheadrightarrow_{(R^* \cup S)_k} & (1)
\end{array}
$$

□

LEMMA B.29. *If R is confluent and R half-commutes with S then if $e \twoheadrightarrow_R e_1$ and $e \rightarrow_{(R^* \cup S)^k} e_2$ then exists e' such that $e_1 \twoheadrightarrow_{(R^* \cup S)^k} e'$ and $e_2 \twoheadrightarrow_R e'$.*

$$
\begin{array}{ccc}
e & \xrightarrow{\;(R^* \cup S)^k\;} & e_2 \\
{\scriptstyle R}\downarrow & & \downarrow{\scriptstyle R} \\
e_1 & \cdots\!\cdots\!\cdots\!\twoheadrightarrow_{(R^* \cup S)^k} & e'
\end{array}
$$

PROOF. Similar to Lemma B.28, by repeatedly "tiling" (1) Lemma B.30 and using (2) R is confluent to match the R^* reductions.

$$
\begin{array}{ccccccccc}
\bullet & \xrightarrow{R} & \bullet & \xrightarrow{S_1} & \bullet & \dashrightarrow & \bullet & \xrightarrow{S_k} & \bullet \\
{\scriptstyle R}\downarrow & & {\scriptstyle R}\downarrow & & {\scriptstyle R}\downarrow & & {\scriptstyle R}\downarrow & & \downarrow{\scriptstyle R} \\
\bullet & \twoheadrightarrow_{R} & (2) & \twoheadrightarrow_{(R^* \cup S)_1} & (1) & \dashrightarrow & (1) & \twoheadrightarrow_{(R^* \cup S)_k} & (1)
\end{array}
$$

□

LEMMA B.30. *If R is confluent and R half-commutes with S then if $e \twoheadrightarrow_R e_1$ and $e \rightarrow_S e_2$ then exists e' such that $e_1 \xrightarrow{\epsilon}_S \cdot \twoheadrightarrow_R e'$ and $e_2 \twoheadrightarrow_R e'$.*

$$
\begin{array}{ccc}
e & \xrightarrow{\;\;S\;\;} & e_2 \\
{\scriptstyle R}\downarrow & & \downarrow{\scriptstyle R} \\
e_1 & \cdots\!\cdots_{S^\epsilon}\rightarrow \bullet \cdots\!\cdots_R\twoheadrightarrow & e'
\end{array}
$$

PROOF. By induction on the size of $e \twoheadrightarrow_R e_1$.

Base Case Immediate from the definition of half-commutes.

Inductive Case Assume the lemma holds for reductions of size k, complete the proof via the following diagram where (1) is due to the induction hypothesis, (2) is from the definition of

R half-commutes with S and, (3) follows from the fact that R is confluent.

$\square$

B.4 ∗-Commutativity

Definition B.31 (∗-Commutativity). A reduction relation R ∗-commutes with S if for all terms a, b, c such that $a \to_R b$ and $a \to_S c$ there exists d such that $b \twoheadrightarrow_S d$ and $c \xrightarrow{\epsilon}_R d$ (right in Fig. 9.)

LEMMA B.32. *If R ∗-commutes with S then for all a, b, c if $a \to_R b$ and $a \twoheadrightarrow_S c$ then there exists d such that $b \twoheadrightarrow_S d$ and $c \xrightarrow{\epsilon}_R d$.*

PROOF. By induction on the size of the reduction $a \twoheadrightarrow_S c$.

(Base case) Here c is the same as a, so just pick $d = b$.

(Ind. case) Assume the lemma for reductions of size less than or equal to n. Suppose that $a \xrightarrow{n+1}_S c$. Then there exists c' such that $a \xrightarrow{n}_S c'$ and $c' \to_S c$. The proof is completed by the diagram:

$\square$

LEMMA B.33. *If R ∗-commutes with S then for all a, b, c if $a \twoheadrightarrow_R b$ and $a \to_S c$ then there exists d such that $b \twoheadrightarrow_S d$ and $c \twoheadrightarrow_R d$.*

PROOF. By induction on the size of the reduction $a \twoheadrightarrow_R b$.

(Base case) Here b is the same as a, so just pick $d = c$.

(Ind. case) Assume the lemma for reductions of size less than or equal to n. Suppose that $a \xrightarrow{n+1}_R b$. Then there exists some b' such that $a \xrightarrow{n}_R b'$ and $b' \to_R b$. The proof is completed by the diagram below.

LEMMA B.34 (∗-COMMUTATIVITY). *If R ∗-commutes with S then R commutes with S.*

PROOF. By induction on the size of the reduction $a \twoheadrightarrow_S c$.

(Base case) Here c is the same as a, so just pick $d = b$.

(Ind. case) Assume the lemma for reductions of size less than or equal to n. Suppose that $a \xrightarrow{n+1}_S c$. Then there exists some c' such that $a \xrightarrow{n}_S c'$ and $c' \rightarrow_S c$. The proof is completed by the diagram below.

□

B.5 Commutativity and Confluence

LEMMA B.35 (COMMUTATIVITY). *If R and S are confluent and commute, then $R \cup S$ is confluent.*

LEMMA B.36 (N-COMMUTATIVITY). *If (i) $\forall 0 \leqslant i \leqslant n$, the reduction relation R_i is confluent, and (ii) $\forall 0 \leqslant i < j \leqslant n$, the reduction relations R_i and R_j commute then $\cup_{i=0}^{n} R_i$ is confluent.*

PROOF. By induction on n using Lemma B.35 and Lemma B.20.   □

B.6 Confluent Kernels

Definition B.37 (Kernel). A reduction relation S is a *kernel* of R, written $S \preceq R$ if (1) $S \subseteq R$ and (2) If $a \rightarrow_R b$ there exists c such that $a, b \twoheadrightarrow_S c$.

LEMMA B.38 (KERNEL-STEPS). *If $S \preceq R$ and S is confluent and $a \twoheadrightarrow_R b$ then $\exists c. \ a, b \twoheadrightarrow_S c$.*

PROOF. By induction on $a \twoheadrightarrow_R b$.

Base Case: $a \equiv b$ so trivially $a, b \twoheadrightarrow_S a$.

Fig. 10. S **is a kernel of** R **written** $S \leq R$

Inductive Case: Assume theorem for $a \xrightarrow{n}_R b$. Suppose that $a \xrightarrow{n}_R b'$ via $a \xrightarrow{n}_R b$ and $b \to_R b'$. The proof follows from the diagram below: c is from the IH, c' from $S \leq R$ and c'' from the confluence of S.

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ □

THEOREM B.39. *Kernel Confluence If $S \leq R$ and S is confluent, then R is confluent.*

PROOF. Suppose that $a \twoheadrightarrow_R b_1$ and $a \twoheadrightarrow_R b_2$. The following diagram shows how to construct c such that $b_1 \twoheadrightarrow_R c$ and $b_2 \twoheadrightarrow_R c$. c_1 (resp. c_2) follows from Lemma B.38 using a and b_1 (resp. b_2). Recall that $S \leq R$ implies every S reduction is also an R reduction.

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ □

C CONFLUENCE OF $\mathcal{VC}$: PROOF

Definition C.1 (Reductions). Let $\mathcal{R}$ be the reduction relation defined as the union $\mathcal{U} \cup \mathcal{N} \cup \mathcal{A} \cup \mathcal{G} \cup C$ of five distinct reduction relations, each of which is defined as the compatible closure of a prototype reduction relation that is in turn defined by rewrite rules in Fig. 3, as follows:

- $\mathcal{U}$ (Unification) is the compatible closure of $\widehat{\mathcal{U}}$, which is the union of the prototype reduction relations specified by rules U-LIT, U-TUP, U-FAIL, U-OCCURS, SUBST, HNF-SWAP, VAR-SWAP, CHOOSE, SEQ-ASSOC, EQN-FLOAT, and SEQ-SWAP.
- $\mathcal{N}$ (Normalization) is the compatible closure of $\widehat{\mathcal{N}}$, which is the union of the prototype reduction relations specified by rules EXI-SWAP, EXI-FLOAT, SUBST (restricted to $x = y$), and VAR-SWAP.
- $\mathcal{A}$ (Application) is the compatible closure of $\widehat{\mathcal{A}}$, which is the union of the prototype reduction relations specified by rules APP-ADD, APP-GT, APP-GT-FAIL, APP-BETA, APP-TUP, and APP-TUP-0.
- $\mathcal{G}$ (Garbage Collection) is the compatible closure of $\widehat{\mathcal{G}}$, which is the union of the prototype reduction relations specified by rules FAIL-ELIM, VAL-ELIM, EXI-ELIM, and EQN-ELIM.

$\mathcal{U}$ and $\widehat{\mathcal{U}}$	$\mathcal{N}$ and $\widehat{\mathcal{N}}$	$\mathcal{A}$ and $\widehat{\mathcal{A}}$	$\mathcal{G}$ and $\widehat{\mathcal{G}}$	C and $\widehat{C}$
U-LIT	EXI-SWAP	APP-ADD	FAIL-ELIM	ONE-FAIL
U-TUP	EXI-FLOAT	APP-GT	VAL-ELIM	ONE-VALUE
U-FAIL	SUBST (restricted to $x = y$)	APP-GT-FAIL	EXI-ELIM	ONE-CHOICE
U-OCCURS	VAR-SWAP	APP-BETA	EQN-ELIM	ALL-FAIL
SUBST		APP-TUP		ALL-VALUE
HNF-SWAP		APP-TUP-0		ALL-CHOICE
VAR-SWAP				CHOOSE-L
CHOOSE				CHOOSE-R
SEQ-ASSOC				
EQN-FLOAT				
SEQ-SWAP				

Fig. 11. Division of the rewrite rules shown in Fig. 3 into groups

		$\mathcal{U}$	$\mathcal{N}$	$\mathcal{A}$	$\mathcal{G}$	C
Unification	$\mathcal{U}$	C.19	C.42	C.49	C.50	C.51
Normalization	$\mathcal{N}$		C.31	C.52	C.53	C.54
Application	$\mathcal{A}$			C.55	C.56	C.57
Garbage Collection	$\mathcal{G}$				C.58	C.59
Choice	C					C.60

Fig. 12. Summary of the confluence and commutativity of the reductions in Definition C.1. The lemmas on the diagonal (resp. non-diagonal) entries establish confluence (resp. commutativity) for the respective relation (resp. pairs of relations).

- C (Choice) is the compatible closure of $\widehat{C}$, which is the union of the prototype reduction relations specified by rules ONE-FAIL, ONE-VALUE, ONE-CHOICE, ALL-FAIL, ALL-VALUE, ALL-CHOICE, CHOOSE-L, and CHOOSE-R.

Let $\widehat{\mathcal{R}} = \widehat{\mathcal{U}} \cup \widehat{\mathcal{N}} \cup \widehat{\mathcal{A}} \cup \widehat{\mathcal{G}} \cup \widehat{C}$; then $\mathcal{R}$ may also be described as the compatible closure of $\widehat{\mathcal{R}}$ (because the operation of taking a compatible closure distributes over $\cup$).

These groups correspond approximately to the sub-headings in Fig. 3, *but not precisely*. In particular, some rewrite rules appear in more than one group: VAR-SWAP is used in both $\mathcal{U}$ and $\mathcal{N}$, and SUBST is used in both $\mathcal{U}$ and (in restricted form) $\mathcal{N}$. Moreover, CHOOSE is used in $\mathcal{U}$ but not in C, although it is listed under "Choice" in Fig. 3.

For convenient reference, the five lists of rules are also displayed in tabular form in Fig. 11.

Definition C.2 (Well-behaved Terms). An equation is *obviously problematic* if

- It is an equation of the form $x = V[\lambda y.\, e]$, where $x \in \mathsf{fvs}(e)$, or
- It is an equation of form $hnf = \lambda x.\, e$ or $\lambda x.\, e = hnf$.

A term e is *obviously problematic* if it contains an obviously problematic equation. A term e is *problematic* if there exists an obviously-problematic term e' such that $e \twoheadrightarrow e'$. A term e is *well-behaved* if it is not problematic.

Our main confluence theorem is as follows:

THEOREM C.3 (CONFLUENCE). *If e is well-behaved and $e \twoheadrightarrow_{\mathcal{R}} e_1$ and $e \twoheadrightarrow_{\mathcal{R}} e_2$ then $e_1 \downarrow_{\mathcal{R}} e_2$.*

Notation

a, b, c, d, e	Expressions (syntax in Fig. 1)
Δ	An expression e that has a redex at the root
$e_1 \subset e_2$	The expression e_1 is a strict sub-term of e_2
$e_1 \subseteq e_2$	The expression e_1 is a sub-term of e_2, including e_2 itself
$a \rightarrow_{\widehat{\mathcal{R}}} b$	a reduces to b via one root-level step of $\mathcal{R}$
$a \rightarrow_{\mathcal{R}} b$	a reduces to b in one step of $\mathcal{R}$
$a \xrightarrow{\epsilon}_{\mathcal{R}} b$	a reduces to b in zero or one step of $\mathcal{R}$
$a \twoheadrightarrow_{\mathcal{R}} b$	a reduces to b in zero or more steps of $\mathcal{R}$
$a \xrightarrow{k}_{\mathcal{R}} b$	a reduces to b in k steps of $\mathcal{R}$

Expression contexts

$$E \quad ::= \quad \square \mid E; e \mid v = E; e \mid \exists x.\, E \mid E \,\textvisiblespace\, e \mid e \,\textvisiblespace\, E \mid \mathbf{one}\{E\} \mid \mathbf{all}\{E\}$$
$$\mid \quad E(v) \mid v(E) \mid \langle v_1, \cdots, E, \cdots, v_n \rangle \mid \lambda x.\, E$$

Note: $e_1 \subseteq e_2$ is equivalent to $\exists E.\, E[e_1] \equiv e_2$.

Fig. 13. Summary of notation

PROOF. First, we partition $\mathcal{R}$ into the relations $\mathcal{U} \cup \mathcal{N}$, $\mathcal{A}$, $\mathcal{G}$ and C. Next, we show that each of these relations is confluent and pairwise commutative (Fig. 12). Finally, we use Lemma B.36 to prove their union $\mathcal{R}$ is confluent. □

The no-recursion condition is only needed to prove $\mathcal{U}$ is confluent, but we assume it globally for clarity.

C.1 Disjointness, Reduction under, and the Diamond property

In talking about confluence we often speak of two different reduction steps with a common starting point, thus $e \rightarrow_{\mathcal{R}} e_1$ and $e \rightarrow_{\mathcal{R}} e_2$. In the first of these there is a sub-term of e, say Δ_1, that is the actual redex; the root of Δ_1 matches some rule in $\mathcal{R}$. Δ_1 is just an ordinary expression, but we use the notation "Δ" to stress that it is the root of a redex (see Fig. 13). Δ_1 is a sub-term of e (or possibly $\Delta = e$), which we write $\Delta_1 \subseteq e$ (again in Fig. 13). Note that $e_1 \subseteq e_2$ is equivalent to saying that there exists some expression context E such that $E[e_1] \equiv e_2$, *i.e.* that e_2 can be decomposed into a context E whose hole is filled by e_1.

Similarly we may identify Δ_2, the redex that is reduced by $e \rightarrow_{\mathcal{R}} e_2$. Now there are two cases to consider:

(1) Δ_1 is disjoint from Δ_2 in e; or
(2) $\Delta_1 \subseteq \Delta_2$, or $\Delta_2 \subseteq \Delta_1$.

One might wonder if Δ_1 can *overlap* Δ_2, but that is not possible: we are discussing syntax trees, not graphs, and so for distinct Δ_1 and Δ_2, either the root of Δ_1 is a child of the root of Δ_2, or vice versa, or neither.

In the first case (a) we have the diamond property immediately:

LEMMA C.4 (DISJOINT). *Let $e \equiv \ldots \Delta_1 \ldots \Delta_2 \ldots$ be an expression with two* disjoint *redexes Δ_1 and Δ_2. If $e \rightarrow \ldots \Delta_1' \ldots \Delta_2 \ldots \equiv e_1$ and $e \rightarrow \ldots \Delta_1 \ldots \Delta_2' \ldots \equiv e_2$ then there exists e' such that $e_1 \rightarrow e'$ and $e_2 \rightarrow e'$.*

PROOF. Trivial: $e' = \ldots \Delta_1' \ldots \Delta_2' \ldots$. □

C.2 Lemmas for Reductions-Under

So to prove the diamond property for a relation $\mathcal{R}$, we should focus attention only on case (b) where the redexes are not disjoint, *i.e.* one occurs under the other. To this end, it suffices to consider the case where one of the reductions is *at the root*, written $e \to_{\widehat{\mathcal{R}}} e_1$ (see Fig. 13 and Appendix B.1), and the *other* occurs under e *i.e.* is of the form $E[\Delta] \to_{\mathcal{R}} e_2$ where $e_2 \equiv E[\Delta']$, and $\Delta \to_{\widehat{\mathcal{R}}} \Delta'$.

Next, we prove a set of "reductions-under" R lemmas that say that if a term e can be (1) reduced using two different rules R and S as $e \to_R e_R$ and $e_S \to_S$, such that (2) the redex for the S reduction *occurs under* the redex for the R reduction, then there exists some e' such that e_R (resp. e_S) can be reduced to e' using some number of S (resp. R) reductions.

The lemmas will be used in two ways. First, to show that two *different* relations commute. Second, that a relation (strongly) commutes with *itself*, *i.e.* has the diamond property, and hence is confluent. In each case, we will split cases on which relation is the "outer" reduction and which is the "inner" and then applying the appropriate "reduction-under" lemma for the outer relation, and using Lemma C.4 for the case where the redexes are disjoint.

C.2.1 Application. The following lemma says that if a term $\Delta_{\mathcal{A}}$ is the root of an $\mathcal{A}$ reduction $\Delta_{\mathcal{A}} \to_{\mathcal{A}} \Delta'_{\mathcal{A}}$ and the $\Delta_{\mathcal{A}}$ additionally contains under it a *subterm* Δ that is the root of some $\mathcal{R}$ reduction $\Delta \to_{\mathcal{R}} \Delta'$ then it is possible to join the result of the $\mathcal{R}$ and $\mathcal{A}$ reduction at a common term $\Delta''_{\mathcal{A}}$ by executing a single step of the *other* reduction, *i.e.* $\mathcal{A}$ and $\mathcal{R}$ respectively. (Recall that $E[e'] \equiv e$ means that $e' \subseteq e$ *i.e.* e' occurs under or is a sub-term of e).

LEMMA C.5 (UNDER-$\mathcal{A}$). *If $\Delta_{\mathcal{A}} \to_{\widehat{\mathcal{A}}} \Delta'_{\mathcal{A}}$ and $\Delta_{\mathcal{A}} \equiv E[\Delta]$ and $\Delta \to_{\widehat{\mathcal{R}}} \Delta'$ then there exists $\Delta''_{\mathcal{A}}$ such that $\Delta'_{\mathcal{A}} \to_{\mathcal{R}} \Delta''_{\mathcal{A}}$ and $E[\Delta'] \to_{\widehat{\mathcal{A}}} \Delta''_{\mathcal{A}}$.*

$$
\begin{array}{ccc}
\Delta_{\mathcal{A}} \equiv E[\Delta] & \xrightarrow{\;\mathcal{R}\;} & E[\Delta'] \\
{\scriptstyle\widehat{\mathcal{A}}}\Big\downarrow & & \Big\downarrow{\scriptstyle\widehat{\mathcal{A}}} \\
\Delta'_{\mathcal{A}} & \cdots\cdots\!\!\!\xrightarrow[\mathcal{R}]{}\!\!\! & \Delta''_{\mathcal{A}}
\end{array}
$$

PROOF. Split cases on the rule used in $\widehat{\mathcal{A}}$.

Case: APP-BETA *i.e.* $\Delta_{\mathcal{A}} \to_{\mathcal{A}} \Delta'_{\mathcal{A}} \equiv (\lambda x.\, e)v \longrightarrow \exists x.\, x = v;\ e$. If $\Delta \subseteq e$, *i.e.* $\mathcal{R} : e \longrightarrow e'$, then join at $\Delta''_{\mathcal{A}} \equiv \exists x.\, x = v;\ e'$. If $\Delta \subseteq v$, *i.e.* $\mathcal{R} : v \longrightarrow v'$, then join at $\Delta''_{\mathcal{A}} \equiv \exists x.\, x = v';\ e$.

Case: APP-TUP *i.e.* $\Delta_{\mathcal{A}} \to_{\mathcal{A}} \Delta'_{\mathcal{A}} \equiv \langle v_0\ \ldots\ v_n\rangle v \longrightarrow \exists x.\, x = v;\ (x = 0;\ v_0\ |\ldots|\ x = n;\ v_n)$. If $\Delta \subseteq v_i$, *i.e.* $\mathcal{R} : v_i \longrightarrow v'_i$, then join at $\exists x.\, x = v;\ (x = 0;\ v_0\ |\ldots|\ x = i;\ v'_i\ \ldots|\ x = n;\ v_n)$. If $\Delta \subseteq v$, *i.e.* $\mathcal{R} : v \longrightarrow v'$, then join at $\exists x.\, x = v';\ (x = 0;\ v_0\ |\ldots|\ x = n;\ v_n)$.

Case: APP-TUP0 *i.e.* $\Delta_{\mathcal{A}} \to_{\mathcal{A}} \Delta'_{\mathcal{A}} \equiv \langle\rangle v \longrightarrow$ **fail**. Here, $\Delta \subseteq v$, *i.e.* $\mathcal{R} : v \longrightarrow v'$, then join at $\Delta''_{\mathcal{A}} \equiv$ **fail**.

Case: APP-ADD, APP-GT-* In any of the primitive application rules, $\Delta \not\subseteq \Delta_{\mathcal{A}}$.

$\square$

C.2.2 Unification.

LEMMA C.6 (UNDER-$\mathcal{U}$). *Let $\mathcal{R}' \equiv \mathcal{R} - \textsc{subst} - \textsc{var-swap}$. If $\Delta_{\mathcal{U}} \to_{\widehat{\mathcal{U}}} \Delta'_{\mathcal{U}}$ and $\Delta_{\mathcal{U}} \equiv E[\Delta]$ and $\Delta \to_{\widehat{\mathcal{R}'}} \Delta'$ then there exists $\Delta''_{\mathcal{U}}$ such that $\Delta'_{\mathcal{U}} \twoheadrightarrow_{\mathcal{R}'} \Delta''_{\mathcal{U}}$ and $E[\Delta'] \to_{\widehat{\mathcal{U}}} \Delta''_{\mathcal{U}}$.*

$$
\begin{array}{ccc}
\Delta_{\mathcal{U}} \equiv E[\Delta] & \xrightarrow{\;\mathcal{R}'\;} & E[\Delta'] \\
{\scriptstyle\widehat{\mathcal{U}}}\Big\downarrow & & \Big\downarrow{\scriptstyle\widehat{\mathcal{U}}} \\
\Delta'_{\mathcal{U}} & \cdots\cdots\!\!\!\twoheadrightarrow\!\!\! & \Delta''_{\mathcal{U}} \\
& {\scriptstyle\mathcal{R}'} &
\end{array}
$$

PROOF. Split cases on the rule used in $\widehat{\mathcal{U}}$.

Case SUBST : Here, $\Delta_{\mathcal{U}} \equiv X[\,x = v\,]$. Split cases on the occurrence of Δ.

   Case $\Delta \subseteq X$, *i.e.* $X \equiv X'[\Delta]$.

$$X'[\Delta][x = v] \xrightarrow{\quad \mathcal{U} \quad} X'\{v/x\}[\Delta\{v/x\}][x = v]$$
$$\mathcal{R}' \downarrow \qquad\qquad\qquad\qquad \vdots\; \mathcal{R}' \text{ via } Lemma\; C.11$$
$$X'[\Delta'][x = v] \xrightarrow{\quad \mathcal{U} \quad} X'\{v/x\}[\Delta'\{v/x\}][x = v]$$

   Case $\Delta \subseteq v$, *i.e.* $v \rightarrow_{\mathcal{R}'} v'$

$$X[x = v] \xrightarrow{\quad \mathcal{U} \quad} X\{v/x\}[x = v]$$
$$\mathcal{R}' \downarrow \qquad\qquad\qquad\qquad \mathcal{R}' \text{ (repeat at each } v)$$
$$X[x = v'] \xrightarrow{\quad \mathcal{U} \quad} X\{v'/x\}[x = v']$$

Case HNF-SWAP : $\Delta_{\mathcal{U}} \equiv hnf = x$ and $h \rightarrow_{\mathcal{R}'} h'$, so join at $\Delta''_{\mathcal{U}} \equiv x = hnf'$.

$$hnf = x \xrightarrow{\quad \mathcal{U} \quad} x = hnf$$
$$\mathcal{R}' \downarrow \qquad\qquad\qquad\qquad \mathcal{R}' \downarrow$$
$$hnf' = x \xrightarrow{\quad \mathcal{U} \quad} x = hnf'$$

Case U-OCCURS : $\Delta_{\mathcal{U}} \equiv x = V[x]$ and $V[x] \rightarrow_{\mathcal{R}'} V'[x]$, so join at $\Delta''_{\mathcal{U}} \equiv$ **fail**.

$$x = V[x] \xrightarrow{\quad \mathcal{U} \quad} \textbf{fail}$$
$$\mathcal{R}' \downarrow \qquad\qquad\qquad \mathcal{U}$$
$$x = V'[x]$$

Case VAR-SWAP : Impossible, no $\Delta \subseteq \Delta_{\mathcal{U}}$
Case U-LIT : Impossible, no $\Delta \subseteq \Delta_{\mathcal{U}}$
Case U-FAIL : Join at $\Delta''_{\mathcal{U}} \equiv$ **fail**.
Case U-TUP : $\Delta_{\mathcal{U}} \equiv (u_1 \ldots u_n) == (v_1 \ldots v_n)$.
   Case $\Delta \subseteq u_i$ i.e. $u_i \rightarrow_{\mathcal{R}'} u'_i$ Join at $\Delta''_{\mathcal{U}} \equiv u_1 = v_1; \ldots u'_i = v_i; \ldots u_n = v_n$.
   Case $\Delta \subseteq v_j$ i.e. $v_j \rightarrow_{\mathcal{R}'} v'_j$ Join at $\Delta''_{\mathcal{U}} \equiv u_1 = v_1; \ldots u_j = v'_j; \ldots u_n = v_n$.
Case SEQ-ASSOC : $\Delta_{\mathcal{U}} \equiv (eq; e_1); e_2 \rightarrow eq; (e_1; e_2) \equiv \Delta'_{\mathcal{U}}$. Split cases on where Δ occurs, which
  as we're precluding SUBST is either in eq or in e_1 or in e_2.
   Case $\Delta \subseteq eq$ i.e. $eq \rightarrow_{\mathcal{R}'} eu'$ Join at $\Delta''_{\mathcal{U}} \equiv eu'; (e_1; e_2)$.
   Case $\Delta \subseteq e_1$ i.e. $e_1 \rightarrow_{\mathcal{R}'} e'_1$ Join at $\Delta''_{\mathcal{U}} \equiv eq; (e'_1; e_2)$.
   Case $\Delta \subseteq e_2$ i.e. $e_2 \rightarrow_{\mathcal{R}'} e'_2$ Join at $\Delta''_{\mathcal{U}} \equiv eq; (e_1; e'_2)$.
   Case Δ spans $(eq; e_1)$ or $(eq; e_1); e_2$ via FAIL-ELIM. Join at **fail**.
Case EQN-FLOAT : $\Delta_{\mathcal{U}} \equiv v = (eq; e_1); e_2 \rightarrow eq; (v = e_1; e_2) \equiv \Delta'_{\mathcal{U}}$. Split cases on where Δ occurs,
  which as we're precluding SUBST is either in v, eq, e_1 or in e_2.
   Case $\Delta \subseteq v$ i.e. $v \rightarrow_{\mathcal{R}'} v'$ Join at $\Delta''_{\mathcal{U}} \equiv eq; (v' = e_1; e_2)$.
   Case $\Delta \subseteq eq$ i.e. $eq \rightarrow_{\mathcal{R}'} eu'$ Join at $\Delta''_{\mathcal{U}} \equiv eu'; (v = e_1; e_2)$.
   Case $\Delta \subseteq e_1$ i.e. $e_1 \rightarrow_{\mathcal{R}'} e'_1$ Join at $\Delta''_{\mathcal{U}} \equiv eq; (v = e'_1; e_2)$.
   Case $\Delta \subseteq e_2$ i.e. $e_2 \rightarrow_{\mathcal{R}'} e'_2$ Join at $\Delta''_{\mathcal{U}} \equiv eq; (v = e_1; e'_2)$.
   Case Δ spans $(eq; e_1)$ or $v = (eq; e_1); e_2$ via FAIL-ELIM. Join at **fail**.
Case CHOOSE : via Lemma C.7.

$\square$

LEMMA C.7 (UNDER-CHOOSE). *If $\Delta_{ch} \rightarrow_{\overline{CHOOSE}} \Delta'_{ch}$ and $\Delta_{ch} \equiv E[\Delta]$ and $\Delta \rightarrow_{\widehat{\mathcal{R}}} \Delta'$ then there exists Δ''_{ch} such that $\Delta'_{ch} \rightarrow\!\!\!\rightarrow_{\mathcal{R}} \Delta''_{ch}$ and $E[\Delta'] \rightarrow_{\overline{CHOOSE}} \Delta''_{ch}$.*

$$
\begin{array}{ccc}
\Delta_{ch} \equiv E[\Delta] & \xrightarrow{\ \mathcal{R}\ } & E[\Delta'] \\
{\scriptstyle \overline{CHOOSE}}\Big\downarrow & & \Big\downarrow{\scriptstyle \overline{CHOOSE}} \\
\Delta'_{ch} & \cdots\cdots\cdots\!\!\!\xrightarrow[\mathcal{R}]{}\!\!\!\twoheadrightarrow & \Delta''_{ch}
\end{array}
$$

PROOF. By the definition of CHOOSE we have

$$
\Delta_{ch} \equiv SX[\,CX[\,e_1 \mid e_2\,]\,] \longrightarrow SX[\,CX[\,e_1\,] \mid CX[\,e_2\,]\,] \equiv \Delta'_{ch}
$$

Split cases on where Δ occurs

Case $\Delta \subseteq e_1$ *i.e.* $e_1 \rightarrow_{\mathcal{R}} e'_1$, so join at $SX[\,CX[\,e'_1\,] \mid CX[\,e_2\,]\,]$.
Case $\Delta \subseteq e_2$ *i.e.* $e_2 \rightarrow_{\mathcal{R}} e'_2$, so join at $SX[\,CX[\,e_1\,] \mid CX[\,e'_2\,]\,]$.
Case $\Delta \subseteq e_1 \mid e_2$ *i.e.* $e_1 \mid e_2 \rightarrow_{\mathcal{R}} e_i$ where e_{3-i} is **fail** so join at $SX[\,CX[\,e_i\,]\,]$.
Case $\Delta \subseteq CX$ *i.e.* $CX \rightarrow_{\mathcal{R}} CX'$ so join (via *two* $\mathcal{R}$ steps) at $SX[\,CX'[\,e_1\,] \mid CX'[\,e_2\,]\,]$.
Case $\Delta \subseteq CX[\,e_1 \mid e_2\,]$ *i.e.* $CX[\,e_1 \mid e_2\,] \rightarrow_{\mathcal{R}} CX'[\,e'_1 \mid e'_2\,]$, so join at $SX[\,CX'[\,e'_1\,] \mid CX'[\,e'_2\,]\,]$.

□

C.2.3 Normalization.

LEMMA C.8 (UNDER-$\mathcal{N}$). *Let $\mathcal{R}' = \mathcal{R} - \mathcal{N} - \mathcal{U}$. If $\Delta_{\mathcal{N}} \rightarrow_{\widehat{\mathcal{N}}} \Delta'_{\mathcal{N}}$ and $\Delta_{\mathcal{N}} \equiv E[\Delta]$ and $\Delta \rightarrow_{\widehat{\mathcal{R}'}} \Delta'$ then exists $\Delta''_{\mathcal{N}}$ such that $\Delta'_{\mathcal{N}} \rightarrow_{\mathcal{R}'} \Delta''_{\mathcal{N}}$ and $E[\Delta'] \rightarrow_{\widehat{\mathcal{N}}} \Delta''_{\mathcal{N}}$.*

PROOF. Split cases on the reduction rule used in $\Delta_{\mathcal{N}} \rightarrow_{\widehat{\mathcal{N}}} \Delta'_{\mathcal{N}}$

Case EXI-SWAP : *i.e.* $\mathcal{N} : \exists x. \exists y. e \longrightarrow \exists y. \exists x. e$. Split cases on the position of Δ.
  Case $\Delta \subseteq e$: *i.e.* $e \rightarrow_{\mathcal{R}'} e'$; join at $\exists x. \exists y. e'$.
  Case $\Delta \subseteq (\exists y. e)$: *i.e.* y eliminated via an EXI-ELIM or EQN-ELIM $\exists y. e \rightarrow_{\mathcal{R}'} e'$; join at $\exists x. e'$.
  Case $\Delta \subseteq (\exists x. \exists y. e)$: *i.e.* x eliminated via an EXI-ELIM or EQN-ELIM $\exists x. \exists y. e \rightarrow_{\mathcal{R}'} \exists y. e'$; join at $\exists y. e'$.
Case EXI-FLOAT : *i.e.* $\mathcal{N} : X[\exists x. e] \longrightarrow \exists x. X[\,e\,]$. Split cases on the position of Δ.
  Case $\Delta \subseteq e$: *i.e.* $e \rightarrow_{\mathcal{R}'} e'$; join at $\exists x. X[\,e'\,]$.
  Case $\Delta \subseteq (\exists x. e)$: *i.e.* x eliminated via an EXI-ELIM or EQN-ELIM $\exists x. e \rightarrow_{\mathcal{R}'} e'$; join at $X[\,e'\,]$.
  Case $\Delta \subseteq X$: *i.e.* $X[\exists x. e] \rightarrow_{\mathcal{R}'} X'[\exists x. e']$; join at $\exists x. X'[\,e'\,]$.
Case SUBST-VAR : *i.e.* $\mathcal{N} : X[\,x = y\,] \longrightarrow (X\{y/x\})[\,x = y\,]$. The only possible position of Δ is $\Delta \subseteq X$ *i.e.* $X^{[}x = y] \rightarrow_{\mathcal{R}'} X'^{[}x = y]$; join at $(X'\{y/x\})[\,x = y\,]$.
Case VAR-SWAP : *i.e.* $\mathcal{N} : x = y \longrightarrow y = x$. Impossible to have $\Delta \subseteq x = y$.

□

C.2.4 Garbage Collection.

LEMMA C.9 (UNDER-$\mathcal{G}$). *If $\Delta_{\mathcal{G}} \rightarrow_{\widehat{\mathcal{G}}} \Delta'_{\mathcal{G}}$ and $\Delta_{\mathcal{G}} \equiv E[\Delta]$ and $\Delta \rightarrow_{\widehat{\mathcal{R}}} \Delta'$ then there exists $\Delta''_{\mathcal{G}}$ such that $\Delta'_{\mathcal{G}} \xrightarrow{\epsilon}_{\mathcal{R}} \Delta''_{\mathcal{G}}$ and $E[\Delta'] \rightarrow_{\widehat{\mathcal{G}}} \Delta''_{\mathcal{G}}$.*

PROOF. Let $\Delta_{\mathcal{G}} \rightarrow_{\widehat{\mathcal{G}}} \Delta'_{\mathcal{G}}$ be the $\mathcal{G}$ redex and split cases on the reduction rule used in the step.

Case VAL-ELIM : *i.e.* $\mathcal{G} : v; e \longrightarrow e$. Split cases on position of Δ
  Case $\Delta \subseteq v$: Join at e.
  Case $\Delta \subseteq e$: *i.e.* $e \rightarrow_{\mathcal{R}} e'$; join at e'.
  Case $\Delta \subseteq v; e$: *i.e.* $v; e \rightarrow_{\text{FAIL-ELIM}}$ **fail** as $e \equiv X[\textbf{fail}]$; join at **fail**.
Case FAIL-ELIM : *i.e.* $\mathcal{G} : X[\textbf{fail}] \longrightarrow \textbf{fail}$. Then $X[\textbf{fail}] \rightarrow_{\mathcal{R}} X'[\textbf{fail}]$ hence join at **fail**.

Case EXI-ELIM : *i.e.* $\mathcal{G} : \exists \overline{y}, x, \overline{z}. e \longrightarrow \exists \overline{y}, \overline{z}. e$; (We can generalize EXI-ELIM to first use a sequence of EXI-SWAP to bring the x binder to the end before applying EXI-ELIM as this does not change the order of the *remaining* binders.) Split cases on position of Δ

Case $\Delta \subseteq e$: *i.e.* $e \rightarrow_\mathcal{R} e'$; join at $\exists \overline{y}, \overline{z}. e'$.
Case $\Delta \subseteq \exists \overline{y}, x, \overline{z}. e$: *i.e.* via EXI-SWAP; join at $\exists \overline{y}, \overline{z}. e$.
Case EQN-ELIM : *i.e.* $\mathcal{G} : \exists x. X[x = v; e] \longrightarrow X[e]$ where $x \notin \text{fvs}(X[v; e])$. (We can generalize EXI-ELIM to first use a sequence of EXI-SWAP to bring the x binder to the end before applying EQN-ELIM as this does not change the order of the *remaining* binders.) Split cases on position of Δ

Case $\Delta \subseteq v$: *i.e.* $v \rightarrow_\mathcal{R} v'$; join at $X[e]$.
Case $\Delta \subseteq e$: *i.e.* $e \rightarrow_\mathcal{R} e'$ (where $\text{fvs}(e') = \text{fvs}(e)$); join at $X[e']$.
Case $\Delta \subseteq X$: *i.e.* $X[x = v; e] \rightarrow_\mathcal{R} X'[x = v; e]$ (where $\text{fvs}(X') = \text{fvs}(X)$); join at $X'[e]$.

□

C.2.5 Choice.

LEMMA C.10 (UNDER-C). *If* $\Delta_C \rightarrow_{\widehat{C}} \Delta'_C$ *and* $\Delta_C \equiv E[\Delta]$ *and* $\Delta \rightarrow_\mathcal{R} \Delta'$ *then there exists* Δ''_C *such that* $\Delta'_C \xrightarrow{\epsilon}_\mathcal{R} \Delta''_C$ *and* $E[\Delta'] \rightarrow_{\widehat{C}} \Delta''_C$.

PROOF. Split cases on the rule used in $\Delta_C \rightarrow_{\widehat{C}} \Delta'_C$.

Case ONE-FAIL (symmetric ALL-FAIL) Impossible as $\Delta \nsubseteq \Delta_C$.
Case ONE-VALUE : Here $\Delta_C \rightarrow_C \Delta'_C \equiv \mathbf{one}\{v\} \longrightarrow v$. Hence $\Delta \subseteq v$ *i.e.* $\mathcal{R} : v \longrightarrow v'$, so join at v'.
Case ALL-VALUE : Here $\Delta_C \rightarrow_C \Delta'_C \equiv \mathbf{all}\{v\} \longrightarrow \langle v \rangle$. Hence $\Delta \subseteq v$ *i.e.* $\mathcal{R} : v \longrightarrow v'$, so join at (v').
Case ONE-CHOICE : Here $\Delta_C \rightarrow_C \Delta'_C \equiv \mathbf{one}\{v \mathbin{|} e\} \longrightarrow v$. If $\Delta \subseteq v$, *i.e.* $\mathcal{R} : v \longrightarrow v'$ then join at v'. If $\Delta \subseteq e$, *i.e.* $\mathcal{R} : e \longrightarrow e'$ then join at v.
Case ALL-CHOICE : Here $\Delta_C \rightarrow_C \Delta'_C \equiv \mathbf{all}\{v_1 \mathbin{|} \ldots \mathbin{|} v_n\} \longrightarrow \langle v_1, \ldots, v_n \rangle$. If $\Delta \subseteq v_i$ ie $\mathcal{R} : v_i \longrightarrow v'_i$ then join at $\langle v_1, \ldots, v'_i, \ldots, v_n \rangle$.
Case CHOOSE-L : (symmetric CHOOSE-R) Here $\Delta_C \rightarrow_C \Delta'_C \equiv \mathbf{fail} \mathbin{|} e \longrightarrow e$. Here, $\Delta \subseteq e$, *i.e.* $\mathcal{R} : e \longrightarrow e'$ so join at e'.
Case CHOOSE-ASSOC : *i.e.* $\Delta_C \rightarrow_C \Delta'_C \equiv (e_1 \mathbin{|} e_2) \mathbin{|} e_3 \longrightarrow e_1 \mathbin{|} (e_2 \mathbin{|} e_3)$. Split cases on where Δ occurs.

Case $\Delta \subseteq e_1$, *i.e.* $\mathcal{R} : e_1 \longrightarrow e'_1$ so join at $e'_1 \mathbin{|} (e_2 \mathbin{|} e_3)$.
Case $\Delta \subseteq e_2$, *i.e.* $\mathcal{R} : e_2 \longrightarrow e'_2$ so join at $e_1 \mathbin{|} (e'_2 \mathbin{|} e_3)$.
Case $\Delta \subseteq e_3$, *i.e.* $\mathcal{R} : e_3 \longrightarrow e'_3$ so join at $e_1 \mathbin{|} (e_2 \mathbin{|} e'_3)$.
Case $\Delta \equiv e_1 \mathbin{|} \mathbf{fail}$, *i.e.* $\mathcal{R} : e_1 \mathbin{|} \mathbf{fail} \longrightarrow e_1$ so join at $e_1 \mathbin{|} e_3$.
Case $\Delta \equiv \mathbf{fail} \mathbin{|} e_2$, *i.e.* $\mathcal{R} : \mathbf{fail} \mathbin{|} e_2 \longrightarrow e_2$ so join at $e_2 \mathbin{|} e_3$.

□

C.3 Lemmas for Substitution and Unification

LEMMA C.11 (SUBSTITUTION). *Let* $\mathcal{R}' \equiv \mathcal{R} - \mathcal{U}$. *If* $\Delta \rightarrow_{\widehat{\mathcal{R}'}} \Delta'$ *then* $\Delta\{v/x\} \rightarrow_{\widehat{\mathcal{R}'}} \Delta'\{v/x\}$.

PROOF. By induction on the structure of Δ, splitting cases on the reduction rule used and using the fact that e, v, X, CX, SX are all closed under value substitution.   □

LEMMA C.12 (SUBST-SWAP). *If* $e \rightarrow_{SUBST} e_1$ *and* $e \rightarrow_{SWAP} e_2$ *then exists* e' *such that* $e_1, e_2 \twoheadrightarrow_\mathcal{U} e'$.

PROOF. Let $\Delta_1 \rightarrow_{SUBST} \Delta'_1$ and $\Delta_2 \rightarrow_{SWAP} \Delta'_2$ be the respective reducts. Via Lemma C.4 it suffices to consider two cases:

Case swap under subst : *i.e.* $\Delta_2 \subseteq \Delta_1$ Let $\Delta_1 \equiv X[x = v]$; split cases on Δ_2 position.

Case $\Delta_2 \equiv x = v$: via rule VAR-SWAP:

$$
\begin{array}{ccc}
X[x=y] & \xrightarrow{\ \ \mathcal{U}\ \ } & X\{y/x\}[x=y] \\[2pt]
& & \Big\downarrow \mathcal{R}' \\[2pt]
\mathcal{R}' \Big\downarrow & & X\{y/x\}[y=x] \\[2pt]
& & \Big\downarrow \mathcal{U} \\[2pt]
X[y=x] & \dashrightarrow[\ \mathcal{U}\ ]{} & X\{x/y\}[y=x]
\end{array}
$$

Case $\Delta_2 \subseteq v$: i.e. $v \rightarrow_{\text{SWAP}} v'$.

$$
\begin{array}{ccc}
X[x=v] & \xrightarrow{\ \ \text{SUBST}\ \ } & X\{v/x\}[x=v] \\[2pt]
\text{SWAP}\Big\downarrow & & \Big\downarrow \text{SWAP (repeat at each } v) \\[2pt]
X[x=v'] & \dashrightarrow[\text{SUBST}]{} & X\{v'/x\}[x=v']
\end{array}
$$

Case $\Delta_2 \subseteq X$: i.e. $X \equiv X'[\ldots \Delta_2 \ldots]$. Let $u' \equiv u\{v/x\}$, split cases on SWAP RHS.
   Case same variable : $\Delta_2 \equiv u = x$ where u is HNF or variable.

$$
\begin{array}{ccc}
X'[\ldots u = x \ldots][x=v] & \xrightarrow{\ \ \text{SUBST}\ \ } & X'\{v/x\}[\ldots u' = v \ldots][x=v] \\[2pt]
& & \Big\Downarrow \\[2pt]
{}^{*}\text{-SWAP}\Big\downarrow & \textit{Lemma C.18}\quad\bullet & \\[2pt]
& & \Big\Uparrow \\[2pt]
X'[\ldots x = u \ldots][x=v] & \dashrightarrow[\text{SUBST}]{} & X'\{v/x\}[\ldots v = u' \ldots][x=v]
\end{array}
$$

   Case different variable : $\Delta_2 \equiv u = y$ where u is HNF or variable.

$$
\begin{array}{ccc}
X'[\ldots u = y \ldots][x=v] & \xrightarrow{\ \ \text{SUBST}\ \ } & X'\{v/x\}[\ldots u' = y \ldots][x=v] \\[2pt]
\text{SWAP}\Big\downarrow & & \Big\downarrow \text{SWAP} \\[2pt]
X'[\ldots y = u \ldots][x=v] & \dashrightarrow[\text{SUBST}]{} & X'\{v/x\}[\ldots y = u' \ldots][x=v]
\end{array}
$$

Case SUBST under SWAP : i.e. $\Delta_1 \subseteq \Delta_2$ Let $\Delta_2 \equiv hnf = x$, so $\Delta_1 \subseteq hnf$, i.e. $hnf \rightarrow_{\text{SUBST}} hnf'$, so join at $x = hnf'$.

$$
\begin{array}{ccc}
hnf = x & \xrightarrow{\ \ \text{SUBST}\ \ } & hnf' = x \\[2pt]
\text{SWAP}\Big\downarrow & & \Big\downarrow \text{SWAP} \\[2pt]
x = hnf & \dashrightarrow[\text{SUBST}]{} & x = hnf'
\end{array}
$$

<div align="right">□</div>

Definition C.13 (Levels). Let $eq_1 \equiv x_1 = v_1$ and $eq_2 \equiv x_2 = v_2$ be two equations in a term e. We say eq_2 *is under* eq_1 if $eq_2 \subseteq X$ and $X[eq_1] \subseteq e$.

LEMMA C.14 (SUBST-SUBST). *If* $e \rightarrow_{\text{SUBST}} e_1$ *and* $e \rightarrow_{\text{SUBST}} e_2$ *then* $e_1 \downarrow_{\mathcal{U}} e_2$.

PROOF. Suppose that the redex $e \rightarrow e_i$ is using the equation $eq_i \equiv x_i = v_i$. Split cases on

   Case eq_1 is under eq_2 and eq_2 is under eq_1: Lemma C.15 completes the proof.
   **Case** eq_1 is under eq_2 and eq_2 is not under eq_1: Lemma C.16 completes the proof.
   Case eq_1 is not under eq_2 and eq_2 is under eq_1: Lemma C.16 completes the proof.

Case eq_1 is not under eq_2 and eq_2 is not under eq_1: The substitutions are disjoint, so Lemma C.4 completes the proof.

$\square$

LEMMA C.15 (SUBST-SAME). *If* $e \rightarrow_{SUBST} e_1$ *using* eq_1 *and* $e \rightarrow_{SUBST} e_2$ *using* eq_2 *such that* eq_1 *is under* eq_2 *and* eq_2 *is under* eq_1, *then* $e_1 \downarrow_{\mathcal{U}} e_2$.

PROOF. Let $eq_1 \equiv x = u$ and $eq_2 \equiv y = v$. As eq_1 is under eq_2 and eq_2 is under eq_1, we have $e \equiv X[x = u; \ b; \ y = w]$ where $b \equiv x_1 = w_1; \ldots ; x_k = w_k$. Let us split cases on whether $x \equiv y$

Case $x \equiv y$: Now, we can further assume that we can use SEQ-ASSOC and SEQ-SWAP to ensure that each $x_i \equiv x$ (by swapping the other equations to the left of eq_1 or right of eq_2 if $x_i < x$ or $x_i > x$ respectively). By the well-behaved assumption, neither u or v are λ-terms, as otherwise, substituting one equation for the other would yield a problematic equation. Hence, we can join e_1 and e_2 using Lemma C.18 via the context $X' \equiv X\{z/x\}[x = z; \square]$ and $b' \equiv b\{z/x\}$ where z is a fresh variable.

$$X[x = u; b; x = v]$$

eq_1 eq_2

$$X\{u/x\}[x = u; \ b\{u/x\}; \ u = v] \qquad\qquad X\{v/x\}[v = u; \ b\{v/x\}; \ x = v]$$

$\equiv$ swap$x=v$over $v=w_i$

$$X'\{u/z\}[b'\{u/z\}; \ u = v] \qquad\qquad X'\{v/z\}[v = u; \ b'\{v/z\}]$$

Lemma C.18

$\gg \bullet \ll$

Case $x \not\equiv y$: Let us split cases on whether x, y appear in $\mathrm{fvs}(u), \mathrm{fvs}(v)$ respectively.
   Case $x \notin \mathrm{fvs}(v), y \notin \mathrm{fvs}(u)$:

$$X[x = u; \ b; \ y = v] \xrightarrow{\quad eq_1 \quad} X\{u/x\}[x = u; \ b\{u/x\}; \ y = v]$$

$eq_2 \downarrow$ $\downarrow eq_2$

$$X\{v/y\}[x = u; \ b\{v/y\}; \ y = v] \dashrightarrow^{\ eu_1\ } X\{u/x, v/y\}[x = u; \ b\{u/x, v/y\}; \ y = v]$$

   Case $x \notin \mathrm{fvs}(v), y \in \mathrm{fvs}(u)$:

$$X[x = u; \ b; \ y = v] \xrightarrow{\qquad\qquad eq_1 \qquad\qquad} X\{u/x\}[x = u; \ b\{u/x\}; \ y = v]$$

$eq_2 \downarrow$ $\downarrow eq_2$

$$X\{v/y\}[x = u\{v/y\}; \ b\{v/y\}; \ y = v] \dashrightarrow^{\ eq_1\ } X\{u\{v/y\}/x, v/y\}[x = u\{v/y\}; \ b\{u\{v/y\}/x, v/y\}; \ y = v]$$

   Case $x \in \mathrm{fvs}(v), y \notin \mathrm{fvs}(u)$: Symmetric to previous case.
   Case $x \in \mathrm{fvs}(v), y \in \mathrm{fvs}(u)$: Split cases on whether u and v are HNF values or not.
      Case *either* u **or** v **are HNF** Join at **fail** if U-OCCURS, else impossible due to well behaved.

Case *neither u* nor *v* are HNF In this case, we have $x \equiv v$ and $y \equiv u$. WLOG assume that $x \prec y$ (the other case is symmetric) so we can join e_1 and e_2 using the following diagram.

$$X[x = y;\ b;\ y = x] \xrightarrow{\quad eq_2(y=x) \quad} X\{x/y\}[x = x;\ b\{x/y\};\ y = x]$$

$$\left\downarrow eq_1(x=y) \right. \qquad\qquad\qquad\qquad\qquad\qquad \vdots \text{ VAR-SWAP}$$

$$X\{x/y\}[x = x;\ b\{x/y\};\ x = y]$$

$$\vdots \text{ SUBST}(x=y)$$

$$X\{y/x\}[x = y;\ b\{y/x\};\ y = y] \qquad\qquad\qquad X\{y/x\}[y = y;\ b\{y/x\};\ x = y]$$

$$\text{SEQ-SWAP} \qquad\qquad\qquad\qquad\qquad\qquad \text{SEQ-SWAP}$$

$\square$

LEMMA C.16 (SUBST-DIFF). *If $e \to_{SUBST} e_1$ using eq_1 and $e \to_{SUBST} e_2$ using eq_2 such that eq_1 is not under eq_2 and eq_2 is under eq_1, then $e_1 \downarrow_{\mathcal{U}} e_2$.*

PROOF. Here, we have $e \equiv X_1[...X_2[x_2 = v_2]...][x_1 = v_1]$ where the substitution with $x_2 = v_2$ does not affect X_1, x_1, v_1. Split cases on whether $x_1 \equiv x_2$.

Case $x_1 \equiv x_2 \equiv x$: By well-behaved we have $x \notin \text{fvs}(v_1), x \notin \text{fvs}(v_2)$, and neither v_1 nor v_2 are λ-terms, as otherwise, substituting one equation for the other would yield a problematic equation. Hence, we can join e_1 and e_2 using Lemma C.17 on the sub-terms $X_2\{v_1/x\}[v_1 = v_2]$ and $X_2\{v_2/x\}[v_1 = v_2]$.

$$X_1[...X_2[x = v_2]...][x = v_1] \xrightarrow{\quad eq_2 \quad} X_1[...X_2\{v_2/x\}[x = v_2]...][x = v_1]$$

$$\left\downarrow eq_1 \right. \qquad\qquad\qquad Lemma\ C.17 \qquad\qquad \vdots\ eq_1$$

$$X_1\{v_1/x\}[...X_2\{v_1/x\}[v_1 = v_2]...][x = v_1] \quad\cdots\!\!\gg \bullet \ll\!\!\cdots\quad X_1\{v_1/x\}[...X_2\{v_2/x\}[v_1 = v_2]...][x = v_1]$$

Case $x_1 \not\equiv x_2$: Let $v_1' \equiv v_1\{v_2/x_2\}$ and $v_2' \equiv v_2\{v_1/x_1\}$. Split cases on whether $x_i \in \text{fvs}(v_{3-i})$.

Case $x_2 \notin \text{fvs}(v_1)$

$$X_1[...X_2[x_2 = v_2]...][x_1 = v_1] \xrightarrow{\quad eq_2 \quad} X_1[...X_2\{v_2/x_2\}[x_2 = v_2]...][x_1 = v_1]$$

$$\left\downarrow eq_1 \right. \qquad\qquad\qquad\qquad\qquad\qquad \vdots\ eq_1$$

$$X_1\{v_1/x_1\}[...X_2\{v_1/x\}[x_2 = v_2']...][x_1 = v_1] \cdots_{\overrightarrow{eq_2}} X_1\{v_1/x\}[...X_2\{v_2'/x_2, v_1/x_1\}[x_2 = v_2']...][x_1 = v_1]$$

Case $x_2 \in \text{fvs}(v_1), x_1 \notin \text{fvs}(v_2)$

$$X_1[...X_2[x_2 = v_2]...][x_1 = v_1] \xrightarrow{\quad eq_2 \quad} X_1[...X_2\{v_2/x_2\}[x_2 = v_2]...][x_1 = v_1]$$

$$\left\downarrow eq_1 \right. \qquad\qquad\qquad\qquad\qquad\qquad \vdots\ eq_1$$

$$X_1\{v_1/x_1\}[...X_2\{v_1/x_1\}[x_2 = v_2]...][x_1 = v_1] \qquad X_1\{v_1/x_1\}[...X_2\{v_1/x_1, v_2/x_2\}[x_2 = v_2]...][x_1 = v_1]$$

$$\cdots\cdots_{\overrightarrow{eq_2}}\cdots\cdots \qquad\qquad \vdots\ eq_2$$

$$X_1\{v_1/x_1\}[...X_2\{v_1'/x_1, v_2/x_2\}[x_2 = v_2]...][x_1 = v_1]$$

Case $x_2 \in \text{fvs}(v_1), x_1 \in \text{fvs}(v_2)$ Split cases on whether v_1, v_2 are HNF values or not.

Case *either* v_1 or v_2 are HNF In this case, we get the below diagram where, since $x_2 \in$ fvs(v_2'), the term $x_2 = v_2'$ either steps to **fail** (and so we can join at **fail**) or the term violates the well-behaved assumption.

$$X_1[\dots X_2[x_2 = v_2]\dots][x_1 = v_1] \xrightarrow{\ eq_2\ } X_1[\dots X_2\{v_2/x_2\}[x_2 = v_2]\dots][x_1 = v_1]$$

$eq_1 \downarrow \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \vdots\, eq_1$

$$X_1\{v_1/x_1\}[\dots X_2\{v_1/x_1\}[x_2 = v_2']\dots][x_1 = v_1] \qquad X_1\{v_1/x_1\}[\dots X_2\{v_1/x_1, v_2'/x_2\}[x_2 = v_2']\dots][x_1 = v_1]$$

Case *neither* v_1 nor v_2 are HNF In this case $v_1 \equiv x_2$ and $v_2 \equiv x_1$, and we can join e_1 and e_2 as shown in the below diagram.

$$X_1[\dots X_2[x_2 = x_1]\dots][x_1 = x_2] \xrightarrow{\ eu_2(x_2=x_1)\ } X_1[\dots X_2\{x_1/x_2\}[x_2 = x_1]\dots][x_1 = x_2]$$

$eu_1(x_1=x_2) \downarrow \qquad\qquad\qquad\qquad\qquad\qquad eu_1(x_1=x_2)$

$$X_1\{x_2/x_1\}[\dots X_2\{x_2/x_1\}[x_2 = x_2]\dots][x_1 = x_2]$$

$\square$

Unification Lemmas The next two *unification* lemmas state that our rewrite rules encode classical unification algorithms. A *value-equation* is an equation of the form $v_1 = v_2$ *i.e.* where both sides are values. A *block* is a sequence of value-equations.

LEMMA C.17 (UNIFY). *If* $\overline{z} \cap (\text{fvs}(\overline{u}) \cup \text{fvs}(\overline{v})) = \emptyset$, $\overline{u}$ *and* $\overline{v}$ *are not* λ-*terms, and* e, e' *are blocks then*

$$X\{\overline{u}/\overline{z}\}[e\{\overline{u}/\overline{z}\}; e'; \overline{u = v}] \downarrow_{\mathcal{U}} X\{\overline{v}/\overline{z}\}[e'; \overline{u = v}; e\{\overline{v}/\overline{z}\}]$$

PROOF. Let $X_{\overline{w}} \equiv X\{\overline{w}/\overline{z}\}$ and $e_{\overline{w}} \equiv e\{\overline{w}/\overline{z}\}$. The proof follows by induction on the triple ($\sharp$free, $\sharp$size, $\sharp n$) where

$$\sharp\text{free} \doteq \sharp\text{fvs}(\overline{u}) + \sharp\text{fvs}(\overline{v})$$
$$\sharp\text{size} \doteq \Sigma_{i=1}^n size(u_i) + size(v_i)$$
$$\sharp n \doteq \text{the cardinality of } \overline{u}, \overline{v}$$

In the base case, the sequences $\overline{u}, \overline{v}$ are empty, and so we trivially have $X[e'; e] \downarrow_{\mathcal{U}} X[e'; e]$ by repeated applications of SEQ-SWAP as e, e' are blocks. In the inductive case, assume that the sequences are non-empty, and split cases on the *first* equation $u_1 = v_1$.

Case $hnf_1 = hnf_2$ with *incompatible* values: Here,

$$X_{\overline{u}}[e_{\overline{u}}; e'; hnf_1 = hnf_2; \overline{u' = v'}] \rightarrow_{\text{U-FAIL}} X_{\overline{u}}[\dots; \text{fail}; \dots]$$

and

$$X_{\overline{v}}[e'; hnf_1 = hnf_2; \overline{u' = v'}; e_{\overline{v}}] \rightarrow_{\text{U-FAIL}} X_{\overline{v}}[\dots; \text{fail}; \dots]$$

after which we can join at **fail** via FAIL-ELIM.

Case $\langle u_1, \dots, u_k \rangle = \langle v_1, \dots, v_k \rangle$ with tuples of the same arity k: use U-TUP to get equations per component and join using the induction hypothesis, which is well-founded as the $\sharp$size is

strictly smaller:

$$X_{\overline{u}}[e_{\overline{u}}; e'; \langle u_1, ..., u_k \rangle = \langle v_1, ..., v_k \rangle; \overline{u' = v'}] \xrightarrow{\text{U-TUP}} X_{\overline{u}}[e_{\overline{u}}; e'; u_1 = v_1, ..., u_k = v_k; \overline{u' = v'}]$$

$$IH \quad \bullet$$

$$X_{\overline{v}}[e'; \langle u_1, ..., u_k \rangle = \langle v_1, ..., v_k \rangle; \overline{u' = v'}; e_{\overline{v}}] \xdashrightarrow{\text{U-TUP}} X_{\overline{v}}[e'; u_1 = v_1, ..., u_k = v_k; \overline{u' = v'}; e_{\overline{v}}]$$

Case $x = y$ (where wlog $x < y$ as else use VAR-SWAP to make it so) use SUBST to replace all occurrences of x with y, and then apply the IH on the *remaining* $n - 1$ equations $\overline{u' = v'}$. Note that the induction is well-founded as in this case $\sharp$free and $\sharp$size are unchanged but the number of equations decreases by one. In the below, $\overline{u''} \equiv \overline{u'}\{y/x\}$ and $\overline{v''} \equiv \overline{v'}\{y/x\}$.

$$X_{\overline{u}}[e_{\overline{u}}; e'; x = y; \overline{u' = v'}] \xrightarrow{\text{SUBST}} X_{\overline{u''}}[e_{\overline{u''}}; \ e'\{y/x\}; \ x = y; \ \overline{u'' = v''}]$$

$$IH \quad \bullet$$

$$X_{\overline{v}}[e'; \ x = y; \ \overline{u' = v'}; \ e_{\overline{v}}] \xdashrightarrow{\text{SUBST}} X_{\overline{v''}}[e'\{y/x\}; \ x = y; \ \overline{u'' = v''}; \ e_{\overline{v''}}]$$

Case $x = h$ where h is an HNF value and $x \notin \text{fvs}(h)$: use SUBST to replace all occurrences of x with h, and then apply the IH on the remaining $n - 1$ equations $\overline{u' = v'}$. Note that the induction is well-founded in this case as $\sharp$free decreases since x is removed from the free variables of $\overline{u'}$ and $\overline{v'}$ and $X_{\overline{u}}$ and $X_{\overline{v}}$ even though the $\sharp$size may increase due to the substitution. In the below, $\overline{u''} \equiv \overline{u'}\{h/x\}$ and $\overline{v''} \equiv \overline{v'}\{h/x\}$.

$$X_{\overline{u}}[e_{\overline{u}}; e'; x = h; \overline{u' = v'}] \xrightarrow{\text{SUBST}} X_{\overline{u''}}[e_{\overline{u''}}; \ e'\{h/x\}; \ x = h; \ \overline{u'' = v''}]$$

$$IH \quad \bullet$$

$$X_{\overline{v}}[e'; \ x = h; \ \overline{u' = v'}; \ e_{\overline{v}}] \xdashrightarrow{\text{SUBST}} X_{\overline{v''}}[e'\{h/x\}; \ x = h; \ \overline{u'' = v''}; \ e_{\overline{v''}}]$$

Case $h = x$ where h is an HNF value and $x \notin \text{fvs}(h)$: we use HNF-SWAP on both sides after which the proof follows as in the previous case.

Case $x = v$ where $x \in \text{fvs}(v)$: either join at **fail** via U-OCCURS or violates the well-behaved assumption.

$\square$

LEMMA C.18 (UNIFY-FLIP). *If* $\overline{z} \cap (\text{fvs}(\overline{u}) \cup \text{fvs}(\overline{v})) = \emptyset$, $\overline{u}$ *and* $\overline{v}$ *are not* λ*-terms, and* e, e' *are blocks then*

$$X\{\overline{u}/\overline{z}\}[e\{\overline{u}/\overline{z}\}; e'; \overline{u = v}] \downarrow_{\mathcal{U}} X\{\overline{v}/\overline{z}\}[e'; \overline{v = u}; e\{\overline{v}/\overline{z}\}]$$

PROOF. Same as Lemma C.17 except using HNF-SWAP and VAR-SWAP to make the equations the same on both sides. $\square$

$$
\begin{array}{lll}
\text{Block} & b ::= \{v = r;\ b\}_\ell \mid \{b; b\}_\ell \mid t \\
\text{RHS} & r ::= b \mid t \\
\text{Tail} & t ::= v \mid v_1\ v_2 \mid \exists x.\ e \mid e_1 \mathbin{\textbf{|}} e_2 \mid \textbf{one}\{e\} \mid \textbf{all}\{e\} \mid \textbf{fail}
\end{array}
$$

Fig. 14. Labeled Blocks

C.4 Unification is Confluent

LEMMA C.19 ($\mathcal{U}$-CONFLUENT). $\mathcal{U}$ is confluent.

PROOF. We prove that $\mathcal{U}$ is confluent via the following strategy inspired by *labeled reductions* [Lévy 1976]. Let $\mathcal{U}_k$ which is a subset of $\mathcal{U}$ that only applies reductions to terms that are *under less than k λs*.

(1) First, we show that $\mathcal{U}_k$ is *locally confluent* for all k (Lemma C.21).
(2) Second, we show that $\mathcal{U}_k$ is *terminating* for all k (Lemma C.23).
(3) Third, consequently, by Lemma B.14 we obtain that $\mathcal{U}_k$ is confluent for all k.
(4) Finally, we show that $\mathcal{U}$ is confluent by using the largest k in two traces, to join two arbitrary sequences of $\mathcal{U}$ reductions Lemma C.22.

□

Definition C.20 (k-Unification). A k-labeled term is a term where each subterm occurring under at most k λ's is marked by a special label ℓ. Let $\mathcal{U}_k$ be defined as the set of all $\mathcal{U}$ reductions where: (1) the $\mathcal{U}$-redex is a ℓ-labeled or occurs under $\leqslant k$ λs, and (2) the SUBST preserves labels.

LEMMA C.21. $\mathcal{U}_k$ is locally confluent.

PROOF. For simplicity, we directly prove that $\mathcal{U}$ is locally confluent (Lemma C.28). The proof carries over to $\mathcal{U}_k$ as the only required $\mathcal{U}$-reductions under $> k$ λs are on *labeled* subterms.   □

We can now prove that any two $\mathcal{U}_k$ reductions (and hence $\mathcal{U}$ reductions) can be joined.

LEMMA C.22 ($\mathcal{U}_k$-JOIN). If $e \twoheadrightarrow_{\mathcal{U}_i} e_i$ and $e \twoheadrightarrow_{\mathcal{U}_j} e_j$ then there exists e' such that $e_i, e_j \twoheadrightarrow_{\mathcal{U}} e'$.

PROOF. Let $k = \max(i, j)$. As $\mathcal{U}_i, \mathcal{U}_j \subseteq \mathcal{U}_k$ we have $e \twoheadrightarrow_{\mathcal{U}_k} e_i$ and $e \twoheadrightarrow_{\mathcal{U}_k} e_j$. By Lemma C.23 and Lemma C.21 and Lemma B.14, $\mathcal{U}_k$ is confluent, hence there exists e' such that $e_i, e_j \twoheadrightarrow_{\mathcal{U}_k} e'$, after which $\mathcal{U}_k \subseteq \mathcal{U}$ completes the proof.   □

LEMMA C.23. $\mathcal{U}_k$ is Noetherian.

PROOF. By induction on k.

Base case ($k \equiv 0$) via Lemma C.27.

Inductive case Assume the induction hypothesis that $\mathcal{U}_k$ is Noetherian and prove $\mathcal{U}_{k+1}$ is Noetherian. Let σ be a $\mathcal{U}_{k+1}$ reduction sequence $e \longrightarrow \dots$. We will prove that σ is finite. By the IH there is some *finite prefix* of the trace $e \twoheadrightarrow_{\mathcal{U}_{k+1}} e'$ after which there are no more $\mathcal{U}$ steps at level $\leqslant k$. Note that e' is finite and of the form $\dots (\lambda x_1.\ e_1) \dots (\lambda x_n.\ e_n) \dots$ comprising n disjoint λ terms. Every $\mathcal{U}_{k+1}$ reduction from e' is a $\mathcal{U}_k$ reduction from some e_i, that occur "in parallel" *i.e.* without influencing each other, and which can be sequenced to get a $\mathcal{U}_{k+1}$ reduction sequence. Again, by the induction hypothesis, each of these reduction sequences (for each e_i) is finite, and hence their sequencing is finite, hence σ must be finite.

□

Labeled Blocks We prove the base case of Lemma C.23 by stratifying expressions into *labeled blocks*, *tails*, *rhs* and *expressions* as shown in Fig. 14. A *tail* is a term that is "inert" for the purposes of $\mathcal{U}_0$ reduction: namely a value, application, existential, one, all or choice. An *rhs* is either a block or a tail (which includes a value). A *labeled block* is a sequence of equations $v = r$ of a value and an RHS r followed by a tail t. We assume each block carries a unique "ghost" label ℓ (that will be used to prove termination). In any block b, for any two labels ℓ_1 and ℓ_2 we write $\ell_1 \prec_b \ell_2$ if the block labelled by ℓ_2 occurs *inside* (under) the block labelled by ℓ_1 in b. We will use b_ℓ to denote the (unique) sub-block of b labeled by ℓ. For rewrites like CHOOSE, SEQ-ASSOC, SEQ-SWAP and EQN-FLOAT, we assume that the rewritten term is given a fresh set of distinct block labels. For rewrites with U-TUP, we assume fresh labels are given to the new (inner) blocks created by tuple matching equations. All other $\mathcal{U}$ rewrites preserve blocks or delete them, so we assume that the same labels carry over to the rewritten terms.

LEMMA C.24. *SEQ-SWAP strongly postpones after* $\mathcal{U}$.

PROOF. Split cases on each reduction of $\mathcal{U}$; the diamond is completed as the rules are non-overlapping. □

Definition C.25 (Elimination). We say a reduction *eliminates* a variable x from a block b if the reduction is (1) a SUBST reduction spanning b or an enclosing block (2) using an equation $x = v$ A reduction sequence *eliminates* a variable x from a block b if there is some reduction in the sequence that eliminates x from b, and the sequence contains no subsequent SUBST reductions spanning any block strictly enclosing b.

LEMMA C.26. $\mathcal{U}_0$ *is Noetherian for all* blocks b.

PROOF. We prove that for any term b that it is only possible to take finitely many $\mathcal{U}_0$ steps from b. Let $\sigma \doteq b \longrightarrow b_1 \longrightarrow b_2 \longrightarrow \dots$ be a $\mathcal{U}_0$ reduction sequence starting at b. Write σ_i for the prefix $b \longrightarrow \dots \longrightarrow b_i$. We will show that σ must be finite. Let $\mathcal{U}_0' \doteq \mathcal{U}_0 - $ SEQ-SWAP. As SEQ-SWAP strongly postpones after $\mathcal{U}$ Lemma C.24, any infinite σ can be translated to a either: (a) A sequence with a *finite* prefix of $\mathcal{U}_0'$ reductions followed by infinitely many SEQ-SWAP, or (b) An infinite sequence of $\mathcal{U}_0'$ reductions. Next, we show neither case is possible.

Case (a) This case is ruled out by the ordering restriction on SEQ-SWAP which ensures that after the finite prefix of $\mathcal{U}_0'$ reductions, we can only keep swapping equations till they reach a canonical linear order after which no further swaps are possible.

Case (b) Next, we (ignore SEQ-SWAP to) show there is no infinite sequence of $\mathcal{U}_0'$ reductions. To do so, suppose that σ is such a reduction sequence. For each prefix (σ_i, e_i) we define the following lexicographic termination metric

$$\sharp(\sigma_i, b_i) \doteq (\sharp\text{choose}(b_i), \sharp\text{semi}(b_i), \text{cands}(\sigma_i, b_i), \text{size}(b_i), \sharp\text{swaps}(b_i))$$

where

$$
\begin{aligned}
\text{choose}(b_i) &\doteq \text{CHOOSE redexes in } b_i \\
\text{semi}(b_i) &\doteq \text{SEQ-ASSOC or EQN-FLOAT redexes in } b_i \\
\text{cands}(\sigma_i, e_i) &\doteq [\dots \ell \mapsto \sharp\text{cand}(\sigma_i, b_i, \ell) \dots \mid \ell \in b_i] \\
&\qquad \text{where labels are ordered by } \prec_{b_i} \\
\text{size}(b_i) &\doteq \text{size of the block } b_i \\
\text{swaps}(b_i) &\doteq \text{VAR-SWAP redexes in } b_i
\end{aligned}
$$

and where, for a finite reduction (prefix) σ' block b and label ℓ

$$\text{cand}(\sigma', b, \ell) \doteq \text{fvs}(b_\ell) - elim(\sigma', b, \ell)$$

$$\text{elim}(\sigma', b, \ell) \doteq \{x \mid \sigma' \text{ eliminates } x \text{ from } b_\ell\}$$

The unification reductions preserve the following invariant: once a variable x has been eliminated from a block, it appears at most once in the block as an LHS of an equation $x = v$ or var-swapped as $v = x$, and in either case that equation can *never* again be used to perform a substitution in that block *unless* new occurrences of x are injected into the block by a substitution performed in an *enclosing* block, in which case, the block metric for the outer block will be *strictly reduced*. Specifically, each application of

- choose strictly reduces $\sharp$choose;
- seq-assoc or eqn-float strictly reduces $\sharp$semi (leaving $\sharp$choose unchanged);
- subst strictly reduces $\sharp$cands (leaving $\sharp$semi, $\sharp$choose unchanged), as it *eliminates* a variable from the block ℓ that the substitution spans, leaving enclosing blocks unchanged;
- u-tup strictly reduces size (leaving cands, $\sharp$semi, $\sharp$choose unchanged), as it preserves elim and hence cand, but reduces the size of ℓ;
- u-lit, u-fail, u-occurs strictly reduces size (leaving cands, $\sharp$semi, $\sharp$choose unchanged);
- var-swap strictly reduces swaps leaving the other components unchanged.

Thus, as $\sharp(\sigma_i, b_i)$ is a strictly decreasing well-founded metric, the sequence $(\sigma_1, b_1), \ldots,$ is finite, and so any sequence of $\mathcal{U}_0$' steps is guaranteed to terminate. □

LEMMA C.27. $\mathcal{U}_0$ is Noetherian for all tails t, rhs r and expressions e.

PROOF. By induction on the structure of t, r and e, using Lemma C.26 for the base case. □

LEMMA C.28. $\mathcal{U}$ is locally confluent.

PROOF. Let $\Delta_1 \to_1 \Delta_1'$ and $\Delta_2 \to_2 \Delta_2'$ denote the two $\mathcal{U}$ reducts. If the reducts are disjoint, then the terms can be joined trivially in a single step via Lemma C.4. By symmetry it suffices to consider the case where Δ_1 occurs under Δ_2. Let us split cases on the rule used for Δ_1.

Case Δ_1 via $\mathcal{U}$ − subst − var-swap join using Lemma C.6.
Case Δ_1 via var-swap join using Lemma C.29.
Case Δ_1 via subst join using Lemma C.30.

□

LEMMA C.29 (VAR-SWAP UNDER). *If* $\Delta_{\mathcal{U}} \to_{\mathcal{U}} \Delta_{\mathcal{U}}'$ *and* $\Delta_{\mathcal{U}} \equiv E[\Delta]$ *and* $\Delta \to_{SWAP} \Delta'$ *then there exists* $\Delta_{\mathcal{U}}''$ *such that* $\Delta_{\mathcal{U}}' \twoheadrightarrow_{SWAP} \Delta_{\mathcal{U}}''$ *and* $E[\Delta'] \to_{\mathcal{U}} \Delta_{\mathcal{R}}''$.

$$
\begin{array}{ccc}
\Delta_{\mathcal{U}} \equiv E[\Delta] & \xrightarrow{\;VAR\text{-}SWAP\;} & E[\Delta'] \\
\mathcal{U} \downarrow & & \vdots\, \mathcal{U} \\
\Delta_{\mathcal{U}}' & \cdots\!\!\cdots\!\!\cdots\!\!\cdots\!\!\cdots\!\!\twoheadrightarrow & \Delta_{\mathcal{U}}'' \\
 & VAR\text{-}SWAP &
\end{array}
$$

PROOF. Split cases on the rule used in $\mathcal{U}$.

Case u-lit or var-swap : impossible as no var-swap redex under $k_1 = k_2$ or $x = y$.
Case u-tup : Here, $\Delta_{\mathcal{U}} \equiv \langle u_1, \ldots, u_n \rangle = \langle v_1, \ldots, v_n \rangle$ and wlog the var-swap redex is $u_1' \to_{u_1}$ so join at $u\_1' = v_1; \ldots; u_n = v_n$.
Case u-fail : Here, $\Delta_{\mathcal{U}} \equiv hnf_i \to hnf_i'$ so join at **fail**

Case u-occurs : Here, $\Delta_{\mathcal{U}} \equiv x = V[x]$ and the var-swap redex is under $V[x]$, i.e. $V^{[}x] \to_{\text{SUBST}}$ $V^{[}x]'$ as the free variables are preserved by var-swap hence we can join at **fail**.

Case hnf-swap : Here, $\Delta_{\mathcal{U}} \equiv hnf = x$ and the var-swap redex is under hnf i.e. $hnf \to_{\text{SUBST}}$ hnf', hence join at $x = hnf'$.

Case subst : via Lemma C.12.

Case choose : via Lemma C.7.

Case seq-assoc : Here, $\Delta_{\mathcal{U}} \equiv (eq; e_1); e_2 \to eq; (e_1; e_2) \equiv \Delta'_{\mathcal{U}}$. Split cases on where Δ occurs.

  Case $\Delta \subseteq eq$ i.e. $eq \to_{\text{VAR-SWAP}} eu'$ Join at $\Delta''_{\mathcal{U}} \equiv eu'; (e_1; e_2)$.
  Case $\Delta \subseteq e_1$ i.e. $e_1 \to_{\text{VAR-SWAP}} e'_1$ Join at $\Delta''_{\mathcal{U}} \equiv eq; (e'_1; e_2)$.
  Case $\Delta \subseteq e_2$ i.e. $e_2 \to_{\text{VAR-SWAP}} e'_2$ Join at $\Delta''_{\mathcal{U}} \equiv eq; (e_1; e'_2)$.

Case eqn-float : $\Delta_{\mathcal{U}} \equiv v = (eq; e_1); e_2 \to eq; (v = e_1; e_2) \equiv \Delta'_{\mathcal{U}}$. Split cases on where Δ occurs.

  Case $\Delta \subseteq v$ i.e. $v \to_{\text{VAR-SWAP}} v'$ Join at $\Delta''_{\mathcal{U}} \equiv eq; (v' = e_1; e_2)$.
  Case $\Delta \subseteq eq$ i.e. $eq \to_{\text{VAR-SWAP}} eu'$ Join at $\Delta''_{\mathcal{U}} \equiv eu'; (v = e_1; e_2)$.
  Case $\Delta \subseteq e_1$ i.e. $e_1 \to_{\text{VAR-SWAP}} e'_1$ Join at $\Delta''_{\mathcal{U}} \equiv eq; (v = e'_1; e_2)$.
  Case $\Delta \subseteq e_2$ i.e. $e_2 \to_{\text{VAR-SWAP}} e'_2$ Join at $\Delta''_{\mathcal{U}} \equiv eq; (v = e_1; e'_2)$.

$\square$

LEMMA C.30 (subst-under). *If* $\Delta_{\mathcal{U}} \to_{\mathcal{U}} \Delta'_{\mathcal{U}}$ *and* $\Delta_{\mathcal{U}} \equiv E[\Delta]$ *and* $\Delta \to_{\text{SUBST}} \Delta'$ *then there exists* $\Delta''_{\mathcal{U}}$ *such that* $\Delta'_{\mathcal{U}} \twoheadrightarrow_{\text{SUBST}} \Delta''_{\mathcal{U}}$ *and* $E[\Delta'] \to_{\mathcal{U}} \Delta''_{\mathcal{U}}$.

$$\Delta_{\mathcal{U}} \equiv E[\Delta] \xrightarrow{\;\;\text{SUBST}\;\;} E[\Delta']$$
$$\mathcal{U} \downarrow \qquad\qquad\qquad \downarrow \mathcal{U}$$
$$\Delta'_{\mathcal{U}} \cdots\cdots\cdots\cdots\cdots\cdots\twoheadrightarrow_{\text{SUBST}} \Delta''_{\mathcal{U}}$$

PROOF. Split cases on the rule used in $\mathcal{U}$.

  Case u-lit or var-swap : impossible as no subst redex under $k_1 = k_2$ or $x = y$.

  Case u-tup : Here, $\Delta_{\mathcal{U}} \equiv \langle u_1, \ldots, u_n \rangle = \langle v_1, \ldots, v_n \rangle$ and wlog the subst redex is $u'_1 \to_{u_1}$ so join at $u\_1' = v_1; \ldots; u_n = v_n$.

  Case u-fail : Here, $\Delta_{\mathcal{U}} \equiv hnf_i \to hnf'_i$ so join at **fail**

  Case u-occurs : Here, $\Delta_{\mathcal{U}} \equiv x = V[x]$ and the subst redex is under $V[x]$, i.e. $V^{[}x] \to_{\text{SUBST}}$ $V^{[}x]'$ as the free variables are preserved by subst hence we can join at **fail**.

  Case hnf-swap : Here, $\Delta_{\mathcal{U}} \equiv hnf = x$ and the subst redex is under hnf i.e. $hnf \to_{\text{SUBST}} hnf'$, hence join at $x = hnf'$.

  Case subst : via Lemma C.14.

  Case choose : via Lemma C.7.

  Case seq-assoc : $\Delta_{\mathcal{U}} \equiv (eq; e_1); e_2 \to eq; (e_1; e_2) \equiv \Delta'_{\mathcal{U}}$. Split cases on where Δ occurs.

    Case $\Delta \subseteq eq$ i.e. $eq \to_{\mathcal{R}'} eu'$ Join at $\Delta''_{\mathcal{U}} \equiv eu'; (e_1; e_2)$.
    Case $\Delta \subseteq e_1$ i.e. $e_1 \to_{\mathcal{R}'} e'_1$ Join at $\Delta''_{\mathcal{U}} \equiv eq; (e'_1; e_2)$.
    Case $\Delta \subseteq e_2$ i.e. $e_2 \to_{\mathcal{R}'} e'_2$ Join at $\Delta''_{\mathcal{U}} \equiv eq; (e_1; e'_2)$.
    Case $\Delta \subseteq (eq; e_1)$ i.e. subst : $(eq; e_1) \longrightarrow (eu'; e'_1)$ Join at $\Delta''_{\mathcal{U}} \equiv eu'; (e'_1; e_2)$.
    Case $\Delta \subseteq ((eq; e_1); e_2)$ i.e. subst : $(eq; e_1); e_2 \longrightarrow (eu'; e'_1); e'_2$ Join at $\Delta''_{\mathcal{U}} \equiv eu'; (e'_1; e'_2)$.

  Case eqn-float : $\Delta_{\mathcal{U}} \equiv v = (eq; e_1); e_2 \to eq; (v = e_1; e_2) \equiv \Delta'_{\mathcal{U}}$. Split cases on where Δ occurs.

    Case $\Delta \subseteq v$ i.e. $v \to_{\mathcal{R}'} v'$ Join at $\Delta''_{\mathcal{U}} \equiv eq; (v' = e_1; e_2)$.
    Case $\Delta \subseteq eq$ i.e. $eq \to_{\mathcal{R}'} eu'$ Join at $\Delta''_{\mathcal{U}} \equiv eu'; (v = e_1; e_2)$.
    Case $\Delta \subseteq e_1$ i.e. $e_1 \to_{\mathcal{R}'} e'_1$ Join at $\Delta''_{\mathcal{U}} \equiv eq; (v = e'_1; e_2)$.
    Case $\Delta \subseteq e_2$ i.e. $e_2 \to_{\mathcal{R}'} e'_2$ Join at $\Delta''_{\mathcal{U}} \equiv eq; (v = e_1; e'_2)$.
    Case $\Delta \subseteq (eq; e_1)$ i.e. subst : $v = (eq; e_1); e_2 \longrightarrow v = (eu'; e'_1); e_2$. Join at $\Delta''_{\mathcal{U}} \equiv eu'; (v = e'_1; e_2)$.

Case $\Delta \subseteq v = (eq; e_1); e_2$ *i.e.* SUBST $: (v = eq; e_1); e_2 \longrightarrow (v' = eu'; e_1'); e_2'$ Join at $\Delta''_{\mathcal{U}} \equiv eu'; (v' = e_1'; e_2')$.

$\square$

C.5 Normalization is Confluent

Recall that $\mathcal{N} \equiv$ EXI-SWAP + EXI-FLOAT + VAR-SWAP + SUBST-VAR where

$$\text{SUBST-VAR} \quad X[x=y;\ e] \longrightarrow (X\{y/x\})[x=y;\ e\{y/x\}]$$

It will be convenient to *factor out* EXI-FLOAT so let

$$\mathcal{SS} \doteq \text{SUBST-VAR} + \text{VAR-SWAP}$$
$$\mathcal{N'} \doteq \mathcal{SS} + \text{EXI-SWAP}$$
$$\mathcal{N} \doteq \mathcal{N'} + \text{EXI-FLOAT}$$

LEMMA C.31 ($\mathcal{N}$-CONFLUENT). *$\mathcal{N}$ is confluent.*

PROOF. The above result follows in two steps. First we show that $\mathcal{N'}$ – *i.e.* normalization-without-EXI-FLOAT – is confluent in Lemma C.35. Second we show that $\mathcal{N'}$ *strongly postpones* after EXI-FLOAT Lemma C.34. Consequently, each $\twoheadrightarrow_{\mathcal{N}}$ can be rewritten as the composition of $\twoheadrightarrow_{\text{EXI-FLOAT}}$ followed by $\twoheadrightarrow_{\mathcal{N'}}$ after which the following diagram completes the proof, where (1) Lemma C.32 (2) Lemma C.33 (3) Lemma C.35. (4) Lemma C.35

$\square$

LEMMA C.32. *If $e \twoheadrightarrow_{\text{EXI-FLOAT}} e_1$ and $e \twoheadrightarrow_{\text{EXI-FLOAT}} e_2$ then exists $e_1 \twoheadrightarrow_{\text{EXI-FLOAT}} e_1', e_2 \twoheadrightarrow_{\text{EXI-FLOAT}} e_2'$, such that $e_1' \downarrow_{\text{EXI-SWAP}} e_2'$.*

PROOF. On each side add the (missing) EXI-FLOAT steps on the other side, and then use (multiple) EXI-SWAP to join. $\square$

LEMMA C.33. *EXI-FLOAT strongly commutes with $\mathcal{N'}$.*

PROOF. Split cases on each possible case of $\mathcal{N}'$, the diamond is completed trivially as the rules are non-overlapping.   □

LEMMA C.34. $\mathcal{N}'$ *strongly postpones after* EXI-FLOAT, *so* $\mathcal{N}^* \equiv \textsc{exi-float}^* \cdot \mathcal{N}'^*$.

PROOF. Split cases on each possible case of $\mathcal{N}'$; the diamond is completed trivially as the rules are non-overlapping.   □

LEMMA C.35. $\mathcal{N}'$ *is confluent.*

PROOF. Via the following diagram, where: (1) is Lemma C.36; (2) is Lemma C.40; (3) is Lemma C.39; (4) is Lemma C.38.

                                                                                                                              □

LEMMA C.36. *If* $e \twoheadrightarrow_{\mathcal{N}'} e'$ *there exists* e'' *such that* $e' \twoheadrightarrow_{SS} e''$ *and* $e \twoheadrightarrow_{EXI\text{-}SWAP} \cdot \twoheadrightarrow_{SS} e''$.

PROOF. By using Lemma B.25 with the facts that SS is confluent (Lemma C.39) and SS *hops after* EXI-SWAP (Lemma C.37).   □

LEMMA C.37. SS *(resp.* $\mathcal{U}$*) hops after* EXI-SWAP.

PROOF. By splitting cases on the SS(resp. $\mathcal{U}$) reduction that precedes the EXI-SWAP.

Case VAR-SWAP Let the $\Delta_{\text{SWAP}} \equiv X[\,x = y\,]$. If the EXI-SWAP preserves the order of x and y then the result follows trivially (as the reductions are non-overlapping.) If the EXI-SWAP toggles the order then the result follows via the diagram

$$\exists x, y. \ldots X[\,x = y\,] \xrightarrow{\text{VAR-SWAP}} \exists x, y. \ldots X[\,y = x\,]$$
$$\text{EXI-SWAP} \downarrow \qquad\qquad\qquad\qquad \downarrow \text{EXI-SWAP}$$
$$\exists y, x. \ldots X[\,x = y\,] \xleftarrow[\text{VAR-SWAP}]{} \exists y, x. \ldots X[\,y = x\,]$$

Case *non*-VAR-SWAP An $\mathcal{U}$ reduction *other than* VAR-SWAP is variable-order independent, so the sequence of $\mathcal{U}$-step followed by EXI-SWAP is equivalent to first doing the EXI-SWAP and then the $\mathcal{U}$step.

                                                                                                                              □

LEMMA C.38. *EXI-SWAP is confluent.*

PROOF. Trivial, via the diamond property. □

LEMMA C.39. $\mathcal{SS} = \textsc{subst-var} + \textsc{swap}$ *is confluent.*

PROOF. Note that $\mathcal{SS}$ is a subset of $\mathcal{U}$; the proof follows similar to the proof of Lemma C.19 (ignoring the bits about U-TUP and U-LIT and U-FAIL and substituting HNF values.) □

LEMMA C.40. $\mathcal{SS}$ *commutes with* $\mathcal{N}'$.

PROOF. Recall that $\mathcal{N}' \equiv \mathcal{SS} + \textsc{exi-swap}$. The proof follows by observing that $\mathcal{SS}$ half-commutes with EXI-SWAP Lemma C.41, recalling that $\mathcal{SS}$ is confluent Lemma C.39, after which Lemma B.27 yields the conclusion $\mathcal{SS}$ commutes with $\mathcal{SS} + \textsc{exi-swap} \equiv \mathcal{N}'$. □

LEMMA C.41. $\mathcal{SS}$ *half-commutes with* EXI-SWAP.

PROOF. Recall that $\mathcal{SS} \equiv \textsc{subst-var} + \textsc{var-swap}$. Split cases and show each reduction half-commutes with EXI-SWAP.

Case SUBST-VAR If the EXI-SWAP occurs under SUBST-VAR then they trivially commute as the variable order is unaffected by the EXI-SWAP. If the SUBST-VAR occurs under EXI-SWAP the proof is completed by the following diagram.

$$\exists y, x. \ldots X[x = y] \xrightarrow{\hspace{2.5cm} \textsc{EXI-SWAP} \hspace{2.5cm}} \exists x, y. \ldots X[x = y]$$

$$\textsc{subst-var} \downarrow \qquad\qquad\qquad\qquad\qquad \downarrow \textsc{var-swap}$$

$$\qquad\qquad\qquad\qquad\qquad\qquad \exists x, y. \ldots X[y = x]$$

$$\qquad\qquad\qquad\qquad\qquad\qquad \vdots\, \textsc{subst-var}$$

$$\exists y, x. \ldots X\{y/x\}[x = y] \xdashrightarrow{\textsc{EXI-SWAP}} \exists x, y. \ldots X\{y/x\}[x = y] \xdashrightarrow{\textsc{var-swap+subst-var}} \exists x, y. \ldots X\{x/y\}[y = x]$$

Case VAR-SWAP (under EXI-SWAP) The non-trivial cases are where the *same* variables x, y are being swapped by both rules (otherwise, the reductions half-commutes trivially via the diamond property). For the variables to be the same, the VAR-SWAP *must* occur under the EXI-SWAP (as otherwise the same variables are not in scope.) Hence, the proof is completed by the following diagram.

$$\exists x, y. \ldots X[x = y] \xrightarrow{\textsc{EXI-SWAP}} \exists y, x. \ldots X[x = y]$$

$$\textsc{var-swap} \downarrow \qquad\qquad\qquad\qquad \uparrow \textsc{var-swap}$$

$$\exists x, y. \ldots X[y = x] \xdashrightarrow{\textsc{EXI-SWAP}} \exists y, x. \ldots X[y = x]$$

□

## C.6 Unification + Normalization is Confluent

Recall that

$$\mathcal{N} \doteq \textsc{exi-float} + \textsc{exi-swap} + \mathcal{SS}$$

and define

$$\mathcal{U}' \doteq \mathcal{U} + \textsc{exi-swap}$$

LEMMA C.42. $\mathcal{U} \cup \mathcal{N}$ *is confluent.*

PROOF. We prove $\mathcal{U} \cup \mathcal{N}$ is confluent by a generalization of the proof of Lemma C.31 where we use the full $\mathcal{U}$ relation (instead of the subset $\mathcal{SS}$). First we show that $\mathcal{U}$' – *i.e.* $\mathcal{U} \cup \mathcal{N}$ without-EXI-FLOAT – is confluent Lemma C.45. Second we show that $\mathcal{U}$' *strongly postpones* after EXI-FLOAT Lemma C.44. Consequently, each $\twoheadrightarrow_{\mathcal{U} \cup \mathcal{N}}$ can be rewritten as the composition of $\twoheadrightarrow_{\text{EXI-FLOAT}}$ followed by $\twoheadrightarrow_{\mathcal{U}'}$ after which the following diagram completes the proof.

□
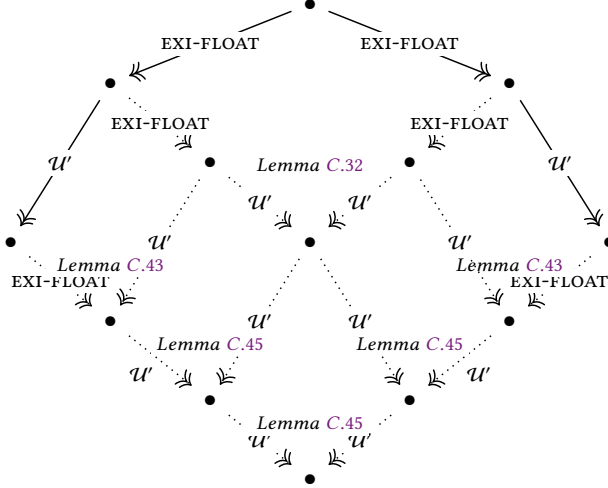
LEMMA C.43. *EXI-FLOAT strongly commutes with $\mathcal{U}$'.*

PROOF. Split cases on each possible case of $\mathcal{U}$'; the diamond is completed trivially as the rules are non-overlapping.   □

LEMMA C.44. *Let $\mathcal{U}' \doteq \mathcal{U} + \text{EXI-SWAP}$. $\mathcal{U}$ strongly postpones after EXI-FLOAT, so $\mathcal{U}'^{*} \equiv \text{EXI-FLOAT}^{*} \cdot \mathcal{U}^{*}$.*

PROOF. Same as Lemma C.34.   □

LEMMA C.45. *Let $\mathcal{U}' \doteq \mathcal{U} + \text{EXI-SWAP}$. $\mathcal{U}'$ is confluent.*

PROOF. Via the following diagram, where: (1) is Lemma C.46; (2) is Lemma C.47; (3) is Lemma C.19; (4) is Lemma C.38.

□

LEMMA C.46. *Let* $\mathcal{U}' = \mathcal{U} + \text{\textsc{exi-swap}}$. *If* $e \twoheadrightarrow_{\mathcal{U}'} e'$ *there exists* e'' *such that* $e' \twoheadrightarrow_{\mathcal{U}} e''$ *and* $e \twoheadrightarrow_{\text{\textsc{exi-swap}}} \cdot \twoheadrightarrow_{\mathcal{U}} e''$.

PROOF. (Similar to Lemma C.36), By using Lemma B.25 with the facts that $\mathcal{U}$ is confluent (Lemma C.19) and $\mathcal{U}$ *hops after* EXI-SWAP (Lemma C.37). □

LEMMA C.47. *Let* $\mathcal{U}' = \mathcal{U} + \text{\textsc{exi-swap}}$. $\mathcal{U}'$ *commutes with* $\mathcal{U}$.

PROOF. The proof follows by observing that $\mathcal{U}$ *half-commutes* with EXI-SWAP Lemma C.48, recalling that $\mathcal{U}$ is confluent Lemma C.19, after which Lemma B.27 yields the conclusion $\mathcal{U}$ commutes with $\mathcal{U} + \text{\textsc{exi-swap}} \equiv \mathcal{U}'$. □

LEMMA C.48. $\mathcal{U}$ *half-commutes with* EXI-SWAP.

PROOF. Same as Lemma C.41; the rules in $\mathcal{U}$ *other than those* in the subset $\mathcal{SS}$ trivially half-commutes as they do not overlap with EXI-SWAP. □

C.7 $\mathcal{U} \cup \mathcal{N}$ Commute With $\mathcal{A} \cup \mathcal{G} \cup C$

LEMMA C.49 ($\mathcal{U}$-$\mathcal{A}$-COMM). $\mathcal{U}$ *and* $\mathcal{A}$ *commute.*

PROOF. We show that $\mathcal{U}$ $*$-commutes with $\mathcal{A}$ and hence commutes via Lemma B.34. Let $\Delta_{\mathcal{U}} \rightarrow_{\mathcal{U}} \Delta'_{\mathcal{U}}$ and $\Delta_{\mathcal{A}} \rightarrow_{\mathcal{A}} \Delta'_{\mathcal{A}}$ denote the reducts for $\mathcal{U}$ and $\mathcal{A}$ respectively.

Case: $\Delta_{\mathcal{U}}$ **and** $\Delta_{\mathcal{A}}$ **disjoint** via Lemma C.4.
Case: $\Delta_{\mathcal{U}} \subseteq \Delta_{\mathcal{A}}$ via Lemma C.5.
Case: $\Delta_{\mathcal{A}} \subseteq \Delta_{\mathcal{U}}$ via Lemma C.6.

□

LEMMA C.50 ($\mathcal{U} - \mathcal{G}$-COMM). $\mathcal{U}$ *and* $\mathcal{G}$ *commute.*

PROOF. We show that $\mathcal{U}$ $*$-commutes commutes wth $\mathcal{G}$, and hence by Lemma B.34, $\mathcal{U}$ commutes wth $\mathcal{G}$. Let $\Delta_{\mathcal{U}} \rightarrow_{\mathcal{U}} \Delta'_{\mathcal{U}}$ and $\Delta_{\mathcal{G}} \rightarrow_{\mathcal{G}} \Delta'_{\mathcal{G}}$ denote the reducts for $\mathcal{U}$ and $\mathcal{G}$ respectively. If the reducts are disjoint then terms can be trivially joined. Let us split cases on whether $\Delta_{\mathcal{U}}$ occurs under $\Delta_{\mathcal{G}}$ or vice versa.

Case $\Delta_{\mathcal{U}} \subseteq \Delta_{\mathcal{G}}$: via Lemma C.9.
Case $\Delta_{\mathcal{G}} \subseteq \Delta_{\mathcal{U}}$: via Lemma C.6.

□

LEMMA C.51 ($\mathcal{U} - C$-COMM). $\mathcal{U}$ *and* C *commute.*

PROOF. We show that $\mathcal{U}$ $*$-commutes with C, and hence by Lemma B.34, $\mathcal{U}$ commutes wth C. Let $\Delta_{\mathcal{U}} \rightarrow_{\mathcal{U}} \Delta'_{\mathcal{U}}$ and $\Delta_C \rightarrow_C \Delta'_C$ denote the reducts for $\mathcal{U}$ and C respectively. If the reducts are disjoint then terms can be trivially joined. Let us split cases on whether $\Delta_{\mathcal{U}}$ occurs under Δ_C or vice versa.

Case $\Delta_{\mathcal{U}} \subseteq \Delta_C$ via Lemma C.10.
Case $\Delta_C \subseteq \Delta_{\mathcal{U}}$ via Lemma C.6.

□

LEMMA C.52. $\mathcal{N}$ *and* $\mathcal{A}$ *commute.*

PROOF. We show that $\mathcal{N}$ strongly commutes with $\mathcal{A}$, hence commutes via Lemma B.19. Let $\Delta_\mathcal{A} \to_\mathcal{A} \Delta'_\mathcal{A}$ and $\Delta_\mathcal{N} \to_\mathcal{N} \Delta'_\mathcal{N}$ denote the reducts for $\mathcal{A}$ and $\mathcal{N}$ respectively. If the reducts are disjoint then terms can be trivially joined in a single step. Let us split cases on whether $\Delta_\mathcal{A}$ occurs under $\Delta_\mathcal{N}$ or vice versa.

Case $\Delta_\mathcal{A} \subseteq \Delta_\mathcal{N}$ via Lemma C.8.
Case $\Delta_\mathcal{N} \subseteq \Delta_\mathcal{A}$ via Lemma C.5.

$\square$

LEMMA C.53. $\mathcal{N}$ and $\mathcal{G}$ commute.

PROOF. We show that $\mathcal{N}$ strongly commutes with $\mathcal{G}$, hence commutes via Lemma B.19. Let $\Delta_\mathcal{G} \to_\mathcal{G} \Delta'_\mathcal{G}$ and $\Delta_\mathcal{N} \to_\mathcal{N} \Delta'_\mathcal{N}$ denote the reducts for $\mathcal{G}$ and $\mathcal{N}$ respectively. If the reducts are disjoint then terms can be trivially joined in a single step. Let us split cases on whether $\Delta_\mathcal{G}$ occurs under $\Delta_\mathcal{N}$ or vice versa.

Case $\Delta_\mathcal{G} \subseteq \Delta_\mathcal{N}$ via Lemma C.8.
Case $\Delta_\mathcal{N} \subseteq \Delta_\mathcal{G}$ via Lemma C.9.

$\square$

LEMMA C.54. $\mathcal{N}$ and C commute.

PROOF. We show that $\mathcal{N}$ strongly commutes with C, hence commutes via Lemma B.19. Let $\Delta_C \to_C \Delta'_C$ and $\Delta_\mathcal{N} \to_\mathcal{N} \Delta'_\mathcal{N}$ denote the reducts for C and $\mathcal{N}$ respectively. If the reducts are disjoint then terms can be trivially joined in a single step. Split cases on whether Δ_C occurs under $\Delta_\mathcal{N}$ or vice versa.

Case $\Delta_C \subseteq \Delta_\mathcal{N}$ via Lemma C.8.
Case $\Delta_\mathcal{N} \subseteq \Delta_C$ via Lemma C.10.

$\square$

C.8 Application

LEMMA C.55. $\mathcal{A}$ is confluent.

PROOF. We show that $\mathcal{A}$ satisfies the diamond property and hence, is confluent by Lemma B.10. Suppose that $e \to_\mathcal{A} e_1$ via the redux $\Delta_1 \to_\mathcal{A} \Delta'_1$, and $e \to_\mathcal{A} e_2$ via the redux $\Delta_2 \to_\mathcal{A} \Delta'_2$. If Δ_1 and Δ_2 are disjoint in e, the terms e_1 and e_2 can be trivially joined in a single step. If $\Delta_1 \subseteq \Delta_2$ (or $\Delta_2 \subseteq \Delta_1$) then Lemma C.5 completes the proof. $\square$

LEMMA C.56. $\mathcal{A}$ and $\mathcal{G}$ commute.

PROOF. We show that $\mathcal{A}$ strongly commutes with $\mathcal{G}$, hence commutes via Lemma B.19. Let $\Delta_\mathcal{A} \to_\mathcal{A} \Delta'_\mathcal{A}$ and $\Delta_\mathcal{G} \to_\mathcal{G} \Delta'_\mathcal{G}$ denote the reducts for $\mathcal{A}$ and $\mathcal{G}$ respectively. If the reducts are disjoint then terms can be trivially joined in a single step. Let us split cases on whether $\Delta_\mathcal{A}$ occurs under $\Delta_\mathcal{N}$ or vice versa.

Case $\Delta_\mathcal{A} \subseteq \Delta_\mathcal{G}$ via Lemma C.9.
Case $\Delta_\mathcal{G} \subseteq \Delta_\mathcal{A}$ via Lemma C.5.

$\square$

LEMMA C.57. $\mathcal{A}$ and C commute.

PROOF. We show that $\mathcal{A}$ strongly commutes with C, hence commutes via Lemma B.19. Let $\Delta_\mathcal{A} \to_\mathcal{A} \Delta'_\mathcal{A}$ and $\Delta_C \to_C \Delta'_C$ denote the reducts for $\mathcal{A}$ and C respectively. If the reducts are disjoint

then terms can be trivially joined in a single step. Let us split cases on whether $\Delta_{\mathcal{A}}$ occurs under Δ_C or vice versa.

Case $\Delta_C \subseteq \Delta_{\mathcal{A}}$ via Lemma C.5.
Case $\Delta_{\mathcal{A}} \subseteq \Delta_C$ via Lemma C.10.

$\square$

C.9 Garbage Collection

LEMMA C.58. $\mathcal{G}$ *is confluent.*

PROOF. We show that $\mathcal{G}$ satisfies the diamond property and hence, is confluent by Lemma B.10. Suppose that $e \rightarrow_{\mathcal{G}} e_1$ via the redux $\Delta_1 \rightarrow_{\mathcal{G}} \Delta_1'$ $e \rightarrow_{\mathcal{G}} e_2$ via the redux $\Delta_2 \rightarrow_{\mathcal{G}} \Delta_2'$. If Δ_1 and Δ_2 are disjoint, the terms e_1 and e_2 can be trivially joined in a single step. If $\Delta_1 \subseteq \Delta_2$ (or $\Delta_2 \subseteq \Delta_1$) then Lemma C.9 completes the proof. $\square$

LEMMA C.59. $\mathcal{G}$ *and* C *commute.*

PROOF. We show that $\mathcal{G}$ and C strongly commute. Let $\Delta_{\mathcal{G}} \rightarrow_{\mathcal{G}} \Delta_{\mathcal{G}}'$ and $\Delta_C \rightarrow_C \Delta_C'$ denote the reducts for $\mathcal{G}$ and $\mathcal{N}$ respectively. If the reducts are disjoint then terms can be trivially joined in a single step. Let us split cases on whether $\Delta_{\mathcal{G}}$ occurs under Δ_C or vice versa.

Case $\Delta_{\mathcal{G}} \subseteq \Delta_C$: via Lemma C.10.
Case $\Delta_C \subseteq \Delta_{\mathcal{G}}$: via Lemma C.9.

$\square$

C.10 Choice

LEMMA C.60. C *is confluent.*

PROOF. Lemma C.10 shows that C has the diamond property (as $C \subseteq \mathcal{R}$), and hence C is confluent via Lemma B.10. $\square$

D   SKEW CONFLUENCE

[This is a sketch of an incomplete proof of skew confluence.]

We now consider a version of $\mathcal{VC}$ that fully supports recursive substitution, and lifts the pesky no-recursion precondition on the confluence theorem. Rather than lifting the side condition $x \notin \mathrm{fvs}(v)$ on rule SUBST, which avoids using a recursive equation for substitution, we take another approach that we believe leads to a simpler proof: we introduce the familiar, conventional fixpoint operator $\mu x.\, hnf$, but allow it to be applied only to head values, not to general expressions. A new unification rule U-OCCURS-WRAP can turn a recursive equation into one that is not recursive (by packaging up the right-hand side within a fixpoint), after which rule SUBST may be applied. A corresponding new rule U-UNWRAP can expand a fixpoint by applying the conventional rewrite rule $\mu x.\, hnf \longrightarrow hnf\{\mu x.\, hnf / x\}$. While this rule allows infinite application, a sensible evaluation strategy would apply this rule "only when needed"—when it is on either side of an equation, or when it is in the function position of an application. If we regard any $\mathcal{VC}$ data structure as tree, the fixpoint construct in effect can label any subtree in such a way that any node beneath it can have a "back pointer" up to the labeled node by referring to the bound variable that serves as the label.

The overall plan is to adapt the proof strategy of Section 4.2 and Appendix C for skew confluence. To do this we need a new result: if two relations are skew confluent with respect to the same information content function and commute, then their union is also skew confluent. (In fact, it is not required that the two relations fully commute: a slightly weaker precondition suffices.) Using this result, our plan is to (i) define an appropriate information content function for $\mathcal{VC}$ expressions; (ii) prove that all the rewrite rules for $\mathcal{VC}$ are monotonic in this information content function; (iii) prove that the Unification rules (modified to permit recursive substitution) together with the Normalization rules are skew confluent; (iv) prove that this combined set of rules commutes in the necessary way with the rules for Application, Elimination, and Choice (which taken together are already known to be confluent); and (v) then apply our new result to show that the entire set of rewrite rules is skew confluent. At this time steps (iii) and (iv) are incomplete, so we emphasize that we do not yet have a complete proof of skew confluence for $\mathcal{VC}$.

D.1   Free Variables

We use the conventional notation $\mathrm{fvs}(e)$ to denote the set of variables that occur free in the expression e. Variables are bound by the constructs $\exists x.\, e$, $\lambda x.\, e$, and $\mu x.\, hnf$. Figure 15 contains a formal definition of this function for $\mathcal{VC}$.

We use the variation $\mathrm{fvsol}(e)$ to denote the set of variables that occur free in the expression e in at least one position that is not within the body of a lambda expression[15]. As an example, $\mathrm{fvsol}(\exists x.\, \langle x, f, g, \lambda y.\, \langle x, g, y, z \rangle \rangle) = \{f, g\}$ because:

- x is bound by $\exists$, so it is not free.
- f is free in a position not within the body of a lambda expression.
- g is free in at least one position not within the body of a lambda expression (it also happens to be free in a second position that is within the body of a lambda expression).
- y is bound by λ, so it is not free.
- z is free, but appears only in a position that is within the body of a lambda expression.

Figure 16 contains a formal definition of this function. Unlike $\mathrm{fvs}(e)$, $\mathrm{fvsol}(v)$ is only ever appplied to a value v.

---

[15]"fvsol$(\cdot)$" abbreviates "free variables outside lambda"

D.2 Substitution

We use the notation $e\{v/x\}$ to denote the expression that results from performing standard capture-avoiding substitution of the value v for every occurrence of the variable x within the expression e (it turns out that, for $\mathcal{VC}$, substitution of general expressions for variables is not required, only substitution of values for variables). Figure 17 contains a formal definition of how this notation applies to the $\mathcal{VC}$ grammar (compare [Barendregt 1984, §2.1.15]).

D.3 Modified grammar and rewrite rules

Let us modify the grammar for $\mathcal{VC}$ to have one additional kind of value, a *fixpoint* value $\mu x. \, hnf$:

$$\text{Values} \quad v ::= x \mid hnf \mid \mu x. \, hnf$$

and a modify the set of Unification rewrite rules $\mathcal{U}$ so that rule U-OCCURS

$$\text{U-OCCURS} \quad x = V[\,x\,] \longrightarrow \textbf{fail} \quad \text{if } V \neq \square$$

is replaced by the two rules[16][17]

$$\begin{array}{lll} \text{U-OCCURS-FAIL} & x = hnf; \, e \longrightarrow \textbf{fail} & \text{if } x \in \mathsf{fvsol}(hnf) \\ \text{U-OCCURS-WRAP} & x = hnf; \, e \longrightarrow x = \mu x. \, hnf; \, e & \text{if } x \in \mathsf{fvs}(hnf) \text{ and } x \notin \mathsf{fvsol}(hnf) \end{array}$$

Let us also add this rewrite rule:

$$\text{U-UNWRAP} \quad \mu x. \, hnf \longrightarrow hnf\{\mu x. \, hnf/x\}$$

As we will see, rule U-UNWRAP is confluent but not Noetherian.

The resulting grammar is not confluent; in particular, it suffers from the even-odd problem described in Section 4.1.2 [Ariola and Blom 2002, Example 4.1]. Therefore we will modify the proof of confluence for $\mathcal{U}$ to become a proof of *skew confluence* [Ariola and Blom 2002], and then go on to prove that $\mathcal{VC}$ itself, with this modification, is skew confluent.

D.4 Unwrapping of Fixpoints is Confluent but not Noetherian

LEMMA D.1. *The rule* U-UNWRAP *is confluent.*

PROOF. Suppose that $e \rightarrow_{\text{U-UNWRAP}} e_1$ and $e \rightarrow_{\text{U-UNWRAP}} e_2$ for distinct redexes within e.
If the redexes are disjoint, then Lemma C.4 applies.
Otherwise, without loss of generality assume the redex for $e \rightarrow_{\text{U-UNWRAP}} e_1$ contains the redex for $e \rightarrow_{\text{U-UNWRAP}} e_2$; let e must be of the form $C_1[\,\mu x. \, C_2[\,\mu y. \, hnf\,]\,]$ (α-conversion may be used to ensure that x and y are distinct variables).
Then $e_1 = C_1[\,C_2[\,\mu y. \, hnf\,]\{\mu x. \, C_2[\,\mu y. \, hnf\,]/x\}\,]$ and $e_2 = C_1[\,\mu x. \, C_2[\,hnf\{\mu y. \, hnf/y\}\,]\,]$.
From e_2 we can take just one more U-UNWRAP step, using the outermost redex $\mu x. \, C_2[\,cdots\,]$, obtaining $e' = C_1[\,(C_2[\,hnf\{\mu y. \, hnf/y\}\,])\{\mu x. \, C_2[\,hnf\{\mu y. \, hnf/y\}\,]/x\}\,]$.

[More to come.]

Thus we have $e_1 \twoheadrightarrow_{\text{U-UNWRAP}} e'$ and $e_2 \rightarrow_{\text{U-UNWRAP}} e'$, so U-UNWRAP is strongly confluent, and therefore by Lemma B.16 is confluent.   □

---

[16]These two rules allow equations to be recursive through lambda expressions and possibly also tuples, but not through tuples only; thus equations such as $f = \lambda y. \, \langle y, f \rangle$ and $x = \langle 1, \lambda y. \, \langle y, x \rangle \rangle$ can be processed by rule U-OCCURS-WRAP, but the equation $x = \langle 1, x \rangle$ can be processed only by rule U-OCCURS-FAIL. An alternate design using the single rule

$$\text{U-OCCURS-WRAP} \quad x = hnf \longrightarrow x = \mu x. \, hnf \quad \text{if } x \in \mathsf{fvs}(hnf)$$

plus rule U-UNWRAP could be used instead to support recursion through tuples only as well as through lambda expressions.
[17]U-OCCURS-FAIL is identical to U-OCCURS in its effect, but it is re-expressed using $\mathsf{fvsol}(\cdot)$, which we need anyway for U-OCCURS-WRAP. Now we can drop the context V, which was only used in U-OCCURS.

$$
\begin{aligned}
\text{fvs}(x) &= \{x\} \\
\text{fvs}(k) &= \{\,\} \\
\text{fvs}(op) &= \{\,\} \\
\text{fvs}(\langle v_1, \ldots, v_n \rangle) &= \text{fvs}(v_1) \cup \cdots \cup \text{fvs}(v_n) \\
\text{fvs}(\lambda x.\, e) &= \text{fvs}(e) \setminus \{x\} \\
\text{fvs}(\mu x.\, e) &= \text{fvs}(e) \setminus \{x\} \\
\text{fvs}(eq;\, e) &= \text{fvs}(eq) \cup \text{fvs}(e) \\
\text{fvs}(v = e) &= \text{fvs}(v) \cup \text{fvs}(e) \\
\text{fvs}(\exists x.\, e) &= \text{fvs}(e) \setminus \{x\} \\
\text{fvs}(\mathbf{fail}) &= \{\,\} \\
\text{fvs}(e_1 \mid e_2) &= \text{fvs}(e_1) \cup \text{fvs}(e_2) \\
\text{fvs}(v_1\ v_2) &= \text{fvs}(v_1) \cup \text{fvs}(v_2) \\
\text{fvs}(\mathbf{one}\{e\}) &= \text{fvs}(e) \\
\text{fvs}(\mathbf{all}\{e\}) &= \text{fvs}(e)
\end{aligned}
$$

Fig. 15. Definition of the free-variables function $\text{fvs}(e)$

$$
\begin{aligned}
\text{fvsol}(x) &= \{x\} \\
\text{fvsol}(k) &= \{\,\} \\
\text{fvsol}(op) &= \{\,\} \\
\text{fvsol}(\langle v_1, \ldots, v_n \rangle) &= \text{fvsol}(v_1) \cup \cdots \cup \text{fvsol}(v_n) \\
\text{fvsol}(\lambda x.\, e) &= \{\,\}
\end{aligned}
$$

Fig. 16. Definition of the free-variables-outside-lambdas function $\text{fvsol}(v)$

$$
\begin{aligned}
x\{v/x\} &\equiv v \\
y\{v/x\} &\equiv y && \text{if } y \neq x \\
k\{v/x\} &\equiv k \\
op\{v/x\} &\equiv op \\
\langle v_1, \ldots, v_n \rangle\{v/x\} &\equiv \langle v_1\{v/x\}, \ldots, v_n\{v/x\} \rangle \\
(\lambda y.\, e)\{v/x\} &\equiv \lambda y.\, e\{v/x\} && \text{if } y \notin \text{fvs}(x, v)\ [\text{use } \alpha] \\
(\mu y.\, v')\{v/x\} &\equiv \mu y.\, v'\{v/x\} && \text{if } y \notin \text{fvs}(x, v)\ [\text{use } \alpha] \\
(eq;\, e)\{v/x\} &\equiv eq\{v/x\};\, e\{v/x\} \\
(v' = e)\{v/x\} &\equiv v'\{v/x\} = e\{v/x\} \\
(\exists y.\, e)\{v/x\} &\equiv \exists y.\, e\{v/x\} && \text{if } y \notin \text{fvs}(x, v)\ [\text{use } \alpha] \\
\mathbf{fail}\{v/x\} &\equiv \mathbf{fail} \\
(e_1 \mid e_2)\{v/x\} &\equiv e_1\{v/x\} \mid e_2\{v/x\} \\
(v_1\ v_2)\{v/x\} &\equiv v_1\{v/x\}\ v_2\{v/x\} \\
(\mathbf{one}\{e\})\{v/x\} &\equiv \mathbf{one}\{e\{v/x\}\} \\
(\mathbf{all}\{e\})\{v/x\} &\equiv \mathbf{all}\{e\{v/x\}\}
\end{aligned}
$$

Fig. 17. Definition of the substitution notation $e\{v/x\}$

Information Content: $\mathcal{I}$	
INFO-FIX	$\mu x.\,v \longrightarrow \Omega$
INFO-SEQ	$eq;\,e \longrightarrow \Omega$
INFO-EXI	$\exists x.\,e \longrightarrow \Omega$
INFO-FAIL-L	$\mathbf{fail}\,\|\,e \longrightarrow \Omega$
INFO-FAIL-R	$e\,\|\,\mathbf{fail} \longrightarrow \Omega$
INFO-CHOICE-OMEGA	$\Omega\,\|\,e \longrightarrow \Omega$
INFO-CHOICE-ASSOC	$(e_1\,\|\,e_2)\,\|\,e_3 \longrightarrow \Omega$
INFO-APP-OMEGA	$\Omega\,v \longrightarrow \Omega$
INFO-APP-HNF	$hnf\,v \longrightarrow \Omega$
INFO-ONE	$\mathbf{one}\{e\} \longrightarrow \Omega$
INFO-ALL	$\mathbf{all}\{e\} \longrightarrow \Omega$

Fig. 18. Rewrite rules on $\mathcal{VC}_\Omega$ for defining the function $\omega_{\mathrm{VC}}(e)$

To see that U-UNWRAP is not Noetherian, observe that

$$\mu x.\,\langle 1, x\rangle \rightarrow_{\text{U-UNWRAP}} \langle 1, \mu x.\,\langle 1, x\rangle\rangle \rightarrow_{\text{U-UNWRAP}} \langle 1, \langle 1, \mu x.\,\langle 1, x\rangle\rangle\rangle \rightarrow_{\text{U-UNWRAP}} \cdots$$

is an unending sequence of reduction steps.

D.5 Information Content

We define a second grammar, for a second language $\mathcal{VC}_\Omega$, by adding one more value Ω, which indicates a lack of information as to just what will be computed. When Ω appears in a context where only a value is permitted, it indicates uncertainty as to what value will be provided; when Ω appears in a context where any expression permitted, it furthermore indicates uncertainty as to how many values will be provided (possibly none).

$$\text{Values} \quad v ::= x \mid hnf \mid \mu x.\,hnf \mid \Omega$$

For every term of $\mathcal{VC}$ there is a corresponding term of $\mathcal{VC}_\Omega$, identical in structure and appearance and containing no occurrence of Ω; we identify such terms and regard the set of terms of $\mathcal{VC}$ as simply a subset of the terms of $\mathcal{VC}_\Omega$.

The definition of substitution (Fig. 17) is extended in the expected trivial manner: $\Omega\{v/x\} \equiv \Omega$.

Definition D.2. (after [Ariola and Blom 2002, Definition 2.3]) Let T be a set of terms over a signature that includes the constant Ω. Define $\leqslant_\omega^T$ be the relation such that $\Omega \leqslant_\omega^T M$ for every term $M \in T$; then define $\leqslant_\omega$ to be the transitive, reflexive, and compatible closure of $\leqslant_\omega^T$.

Figure 18 shows a system $\mathcal{I}$ of rewrite rules on $\mathcal{VC}_\Omega$. These may be compared to similar rules for the $\lambda\circ$ calculus [Ariola and Blom 2002, Definition 5.20].

LEMMA D.3. *[Huet 1980, Lemma 3.1] The relation $\rightarrow_\mathcal{R}$ is locally confluent iff for every critical pair (e_1, e_2) of $\mathcal{R}$, e_1 and e_2 can be joined—that is, there exists e_3 such that $e_1 \rightarrow_\mathcal{R} e_3$ and $e_2 \rightarrow_\mathcal{R} e_3$.*

LEMMA D.4 ($\mathcal{I}$-CONFLUENT). *$\mathcal{I}$ is locally confluent and Noetherian; therefore $\mathcal{I}$ is confluent.*

PROOF. Consider all critical pairs of $\mathcal{I}$:
- Rules INFO-FAIL-L and INFO-FAIL-R produce the critical pair (Ω, Ω).
- Rules INFO-FAIL-L and INFO-CHOICE-OMEGA produce no critical pairs.
- Rules INFO-FAIL-L and INFO-CHOICE-ASSOC produce the critical pair $(\Omega\,\|\,e, \Omega)$.
- Rules INFO-FAIL-R and INFO-CHOICE-OMEGA produce the critical pair (Ω, Ω).

- Rules info-fail-r and info-choice-assoc produce the critical pairs (Ω, Ω) and $(\Omega \mathbin{\|} e, \Omega)$.
- Rules info-choice-omega and info-choice-assoc produce the critical pair $(\Omega \mathbin{\|} e, \Omega)$.
- No other pairs of rules produce any critical pairs.

The critical pair (Ω, Ω) can be trivially joined at Ω. The critical pair $(\Omega \mathbin{\|} e, \Omega)$ can be joined at Ω in one step by using rule info-choice-omega on $\Omega \mathbin{\|} e$.

All critical pairs can be joined; therefore by Lemma D.3 $\mathcal{I}$ is locally confluent.

Let the size of a term of $\mathcal{VC}_\Omega$ be the number of tokens it contains other than parentheses. Each of the rewrite rules in Fig. 18 strictly decreases the number of such tokens. Because any given term has a finite number n of such tokens, and the number of tokens cannot be less than zero, for every term every rewriting sequence from that term can have no more than n steps. So $\mathcal{I}$ is bounded and therefore Noetherian.

Because $\mathcal{I}$ is locally confluent and Noetherian, it is confluent by Newman's lemma. □

Because $\mathcal{I}$ is confluent and Noetherian, it follows that $\mathcal{I}$ defines unique normal forms for $\mathcal{VC}_\Omega$. Therefore we are justified in defining $\omega_{\mathrm{VC}}(e)$ to be the function that takes any term in $\mathcal{VC}_\Omega$ and returns the term that is its normal form under $\mathcal{I}$.

Definition D.5. The comparison $e \leqslant_{\omega_{\mathrm{VC}}} e'$ is defined to mean $\omega_{\mathrm{VC}}(e) \leqslant_\omega \omega_{\mathrm{VC}}(e')$.

[Need to prove that $\leqslant_{\omega_{\mathrm{VC}}}$ is monotonic with respect to $\leqslant_\omega$; this should be routine [Ariola and Blom 2002, Proposition 5.21].]

Lemma D.6 ($\mathcal{VC}$ monotonic). *Every rewrite rule in $\mathcal{VC}$ is monotonic with respect to $\leqslant_{\omega_{\mathrm{VC}}}$.*

Proof. For every rewrite rule in $\mathcal{A} \cup \mathcal{U} \cup \mathcal{N} \cup \mathcal{G} \cup \mathcal{C}$, the left-hand side is an expression that is mapped to Ω by the function ω_{VC}, and no matter what expression e is the result of applying ω_{VC} to the right-hand side, we have $\Omega \leqslant_\omega e$. □

D.6 Preliminaries about Skew Confluence

Why use skew confluence? Ordinary confluence is useful because if term e has an $\mathcal{R}$-normal form, then that normal form is *unique* if $\mathcal{R}$ is confluent. Ariola and Blom define a related notion, which we will refer to as $\mathcal{R}$-*skew-normal form*[18], and prove that under certain conditions, if term e has an $\mathcal{R}$-skew-normal form, then that normal form is unique if $\mathcal{R}$ is skew confluent.

A $\mathcal{R}$-skew-normal form is not a single term, but rather a possibly infinite set of *erased* terms. We summarize this idea, using our own terminology, as follows:

- Let an *erasure* of a term be a copy in which some number of subterms (possibly zero, and possibly the entire term) have been replaced with Ω, a special term that means "unknown" or "we don't know yet."
- We can compare erased terms with the partial order $\leqslant_\omega$, which is the transitive, reflexive, and compatible closure of the relation in which Ω is less than any other term. Observe that if e' is any erasure of e (including e itself) then $e' \leqslant_\omega e$.
- Define the $\mathcal{R}$-*information content* $\omega_\mathcal{R}(e)$ of a term e to be the unique erasure of e in which *every* redex has been replaced by Ω. Any non-Ω structure in the result is therefore "permanent": no further reductions under $\mathcal{R}$ can alter that structure.
- Define the *downward closure* $\Downarrow_{\leqslant_\omega} A$ of a set of terms A to be the set of all elements of A and all possible erasures of those elements, that is, $\Downarrow_{\leqslant_\omega} A = \{e' \mid e \in A, e' \leqslant_\omega\}$.
- Define the $\mathcal{R}$-*skew-normal form* of e to be the downward closure of the set of information contents of all possible $\mathcal{R}$-reducts of e, that is, $\Downarrow_{\leqslant_\omega} \{\omega(e') \mid e \twoheadrightarrow_\mathcal{R} e'\}$.

---

[18]Ariola and Blom call it the "infinite normal form"; this is a bit misleading because in fact not all such forms are infinite.

Taking the downward closure with respect to $\leqslant_\omega$ is crucial; without that step, it would not be possible to prove that skew-normal forms are unique for certain reduction relations.

Skew confluence is a natural extension of confluence, in this sense: if $\mathcal{R}$ is skew confluent, then an expression e has a unique $\mathcal{R}$-normal form if and only if its $\mathcal{R}$-skew-normal form is the (finite) set consisting of all possible erasures of that $\mathcal{R}$-normal form. (For example, $\langle 1, 2 \rangle$ is the unique normal form of $\exists x.\ x = 2;\ \langle 1, x \rangle$, and the $\mathcal{R}$-skew-normal form of that same term is $\{\langle 1, 2 \rangle, \langle \Omega, 2 \rangle, \langle 1, \Omega \rangle, \langle \Omega, \Omega \rangle, \Omega\}$.)

But working with possibly infinite sets directly is tricky. Fortunately, there is a simpler path, because Ariola and Blom prove an important theorem: Define the partial order $e \preceq_{\omega_\mathcal{R}} e'$ to mean $\omega_\mathcal{R}(e) \leqslant_\omega \omega_\mathcal{R}(e')$; then a reduction relation that is *monotonic* in $\preceq_{\omega_\mathcal{R}}$ ($e \rightarrow_\mathcal{R} e' \implies e \preceq_{\omega_\mathcal{R}} e'$) has unique skew-normal forms if and only if the reduction relation is skew confluent [Ariola and Blom 2002, Theorem 5.4]—and skew confluence is much easier to prove, using techniques that do not involve possibly infinite sets, but are fairly similar to proofs of ordinary confluence that use commutative diagrams and case analysis. They also prove that if a reduction relation is confluent and monotonic, then it is skew confluent [Ariola and Blom 2002, Corollary 5.5].

Definition D.7. Reduction relation $\mathcal{R}$ over the set of terms T is *skew confluent* using quasi order $\preceq$ if for all $a, b, c \in T$, if $a \twoheadrightarrow_\mathcal{R} b$ and $a \twoheadrightarrow_\mathcal{R} c$, there exists $d \in T$ such that $b \twoheadrightarrow_\mathcal{R} d$ and $c \preceq d$. As a diagram:

$$
\begin{array}{ccc}
a & \overset{\mathcal{R}}{\twoheadrightarrow} & b \\
{\scriptstyle\mathcal{R}}\downarrow & & \vdots{\scriptstyle\mathcal{R}} \\
c & \overset{\preceq}{\underset{\mathcal{R}}{\twoheadleftarrow\cdots\twoheadrightarrow}} & d
\end{array}
$$

[More to come.]

D.7 New Lemmas about Skew Confluence

LEMMA D.8. *If relation $\mathcal{R}$ is monotonic in some quasi order $\preceq$ and confluent, then it is skew confluent using $\preceq$.*

PROOF. By the definition of confluence,

$$
\begin{array}{ccc}
a & \overset{\mathcal{R}}{\twoheadrightarrow} & c \\
{\scriptstyle\mathcal{R}}\downarrow & & \vdots{\scriptstyle\mathcal{R}} \\
b & \underset{\mathcal{R}}{\cdots\twoheadrightarrow} & d
\end{array}
$$

Because $\mathcal{R}$ is monotonic, $\twoheadrightarrow_\mathcal{R} \subset \overset{\preceq}{\underset{\mathcal{R}}{\twoheadleftarrow\twoheadrightarrow}}$, therefore

$$
\begin{array}{ccc}
a & \overset{\mathcal{R}}{\twoheadrightarrow} & c \\
{\scriptstyle\mathcal{R}}\downarrow & & \vdots{\scriptstyle\mathcal{R}} \\
b & \overset{\preceq}{\underset{\mathcal{R}}{\twoheadleftarrow\cdots\twoheadrightarrow}} & d
\end{array}
$$

$\square$

Definition D.9. Let reduction relation $\mathcal{R}$ be monotonic in quasi order $\preceq$ and skew confluent using $\preceq$. Let reduction relation $\mathcal{R}^\leftarrow_\preceq$ be defined by a set of rewrite rules that are converses of those rewrite rules of $\mathcal{R}$ whose converses are used in the proof that $\mathcal{R}$ is skew confluent. Then we say that $\mathcal{R}$ is *skew confluent using $\preceq$ and $\mathcal{R}^\leftarrow_\preceq$*.

LEMMA D.10. *Let relation $\mathcal{R}$ be monotonic in quasi order $\preceq$ and skew confluent using $\preceq$ and $\mathcal{R}_{\preceq}^{\leftarrow}$; similarly let $\mathcal{S}$ be monotonic in the same quasi order $\preceq$ and skew confluent using $\preceq$ and $\mathcal{S}_{\preceq}^{\leftarrow}$. Suppose furthermore that $\mathcal{R}$ and $\mathcal{S}$ commute and that $\mathcal{R}$ and $\mathcal{S}_{\preceq}^{\leftarrow}$ commute. Then the following four diagrams hold:*

PROOF. (1) Because $\mathcal{R}$ is skew confluent, we have

Because $\twoheadleftarrow_{\mathcal{R}}^{\preceq}\twoheadrightarrow \subset \twoheadleftarrow_{\mathcal{R}\cup\mathcal{S}}^{\preceq}\twoheadrightarrow$, the first diagram follows.

(2) Because $\mathcal{R}$ and $\mathcal{S}$ commute, we have

Because $\twoheadrightarrow_{\mathcal{S}}\subset \twoheadleftarrow_{\mathcal{R}\cup\mathcal{S}}^{\preceq}\twoheadrightarrow$, the second diagram follows.

(3) The following diagram clearly holds if the sequence of reduction steps from b to d is the same as the sequence of reduction steps from b to a to d:

Because $\equiv\subset \twoheadleftarrow_{\mathcal{R}\cup\mathcal{S}}^{\preceq}\twoheadrightarrow$, the third diagram follows.

(4) Because $\mathcal{R}$ and $\mathcal{S}_{\preceq}^{\leftarrow}$ commute, we have

Because $\twoheadleftarrow_{\mathcal{S}}\subset \twoheadleftarrow_{\mathcal{R}\cup\mathcal{S}}^{\preceq}\twoheadrightarrow$, the fourth diagram follows.

[It may turn out to be impossible to prove for our specific application that $\mathcal{R}$ and $\mathcal{S}_{\preceq}^{\leftarrow}$ commute. In that case, it may be necessary to use a more complicated or more subtle precondition. The important thing is to prove the fourth diagram somehow.]

$\square$

LEMMA D.11. *Let relation $\mathcal{R}$ be monotonic in quasi order $\preceq$ and skew confluent using $\preceq$ and $\mathcal{R}_{\preceq}^{\leftarrow}$; similarly let $\mathcal{S}$ be monotonic in the same quasi order $\preceq$ and skew confluent using $\preceq$ and $\mathcal{S}_{\preceq}^{\leftarrow}$. Suppose furthermore that $\mathcal{R}$ and $\mathcal{S}$ commute and that $\mathcal{R}$ and $\mathcal{S}_{\preceq}^{\leftarrow}$ commute. Then*

$$a \xleftrightarrow[\mathcal{R} \cup \mathcal{S}]{\leq} b$$
$$\mathcal{R} \downarrow \qquad \qquad \downarrow \mathcal{R}$$
$$c \xleftrightarrow[\mathcal{R} \cup \mathcal{S}]{\leq} d$$

PROOF. By induction on the size of the top edge of the diagram. At each step one of the four diagrams from Lemma D.10 will be used.

[More to come.]

□

LEMMA D.12. *Let relation $\mathcal{R}$ be monotonic in quasi order $\leq$ and skew confluent using $\leq$ and $\mathcal{R}_{\leq}^{\leftarrow}$; similarly let $\mathcal{S}$ be monotonic in the same quasi order $\leq$ and skew confluent using $\leq$ and $\mathcal{S}_{\leq}^{\leftarrow}$. Then*

$$a \xleftrightarrow[\mathcal{R} \cup \mathcal{S}]{\leq} b$$
$$\mathcal{R} \cup \mathcal{S} \downarrow \qquad \qquad \downarrow \mathcal{R} \cup \mathcal{S}$$
$$c \xleftrightarrow[\mathcal{R} \cup \mathcal{S}]{\leq} d$$

PROOF. By case analysis on whether left edge uses $\mathcal{R}$ or $\mathcal{S}$; then project that left edge into $\mathcal{R}^*$ or $\mathcal{S}^*$ respectively and apply Lemma D.11.

[More to come.]

□

LEMMA D.13. *Let relation $\mathcal{R}$ be monotonic in quasi order $\leq$ and skew confluent using $\leq$ and $\mathcal{R}_{\leq}^{\leftarrow}$; similarly let $\mathcal{S}$ be monotonic in the same quasi order $\leq$ and skew confluent using $\leq$ and $\mathcal{S}_{\leq}^{\leftarrow}$. Then*

$$a \xleftrightarrow[\mathcal{R} \cup \mathcal{S}]{\leq} b$$
$$\mathcal{R} \cup \mathcal{S} \downarrow \qquad \qquad \downarrow \mathcal{R} \cup \mathcal{S}$$
$$c \xleftrightarrow[\mathcal{R} \cup \mathcal{S}]{\leq} d$$

PROOF. By induction on the size of the left edge of the diagram.

Base case This diagram clearly holds by letting the bottom edge be the same as the top edge:

$$a \xleftrightarrow[\mathcal{R} \cup \mathcal{S}]{\leq} b$$
$$\equiv \qquad \qquad \equiv$$
$$c \xleftrightarrow[\mathcal{R} \cup \mathcal{S}]{\leq} d$$

and it implies this diagram:

$$a \xleftrightarrow[\mathcal{R} \cup \mathcal{S}]{\leq} b$$
$$(\mathcal{R} \cup \mathcal{S})^0 \downarrow \qquad \qquad \downarrow \mathcal{R} \cup \mathcal{S}$$
$$c \xleftrightarrow[\mathcal{R} \cup \mathcal{S}]{\leq} d$$

Inductive case Assume the diagram holds for left edges of all sizes up to n. Then this diagram:

$$
\begin{array}{ccc}
a & \overset{\preceq}{\underset{\mathcal{R}\cup\mathcal{S}}{\Longleftarrow\!\Longrightarrow}} & b \\
{\scriptstyle(\mathcal{R}\cup\mathcal{S})^{n+1}}\Big\downarrow & & \Big\downarrow{\scriptstyle\mathcal{R}\cup\mathcal{S}} \\
c & \overset{\preceq}{\underset{\mathcal{R}\cup\mathcal{S}}{\Longleftarrow\!\Longrightarrow}} & d
\end{array}
$$

follows from this diagram:

$$
\begin{array}{ccc}
a & \overset{\preceq}{\underset{\mathcal{R}\cup\mathcal{S}}{\Longleftarrow\!\Longrightarrow}} & b \\
{\scriptstyle(\mathcal{R}\cup\mathcal{S})^{n}}\Big\downarrow & & \Big\downarrow{\scriptstyle\mathcal{R}\cup\mathcal{S}} \\
c & \overset{\preceq}{\underset{\mathcal{R}\cup\mathcal{S}}{\Longleftarrow\!\Longrightarrow}} & d \\
{\scriptstyle\mathcal{R}\cup\mathcal{S}}\Big\downarrow & {\scriptstyle D.12} & \Big\downarrow{\scriptstyle\mathcal{R}\cup\mathcal{S}} \\
e & \overset{\preceq}{\underset{\mathcal{R}\cup\mathcal{S}}{\Longleftarrow\!\Longrightarrow}} & f
\end{array}
$$

where the top half is the inductive hypothesis and the bottom half follows from Lemma D.12.

□

LEMMA D.14. *Let relation $\mathcal{R}$ be monotonic in quasi order $\preceq$ and skew confluent using $\preceq$ and $\mathcal{R}_{\preceq}^{\leftarrow}$; similarly let $\mathcal{S}$ be monotonic in the same quasi order $\preceq$ and skew confluent using $\preceq$ and $\mathcal{S}_{\preceq}^{\leftarrow}$. If $\mathcal{R}$ commutes with $\mathcal{S}$, then $\mathcal{T} = \mathcal{R} \cup \mathcal{S}$ is monotonic in $\preceq$ and skew confluent using $\preceq$ and $\mathcal{T}_{\preceq}^{\leftarrow}$.*

PROOF.   □

D.8 Proof that $\mathcal{VC}$ Is Skew Confluent

[This is just a brief proof sketch.]

First prove that the modified $\mathcal{U}$ is skew confluent. (In doing so we will define $\mathcal{U}_{\lessgtr_{\omega VC}}^{\leftarrow}$.)

Then use existing proofs to demonstrate that $\mathcal{A} \cup \mathcal{N} \cup \mathcal{G} \cup \mathcal{C}$ is confluent. Because they are also monotonic, they are therefore skew confluent, and $(\mathcal{A} \cup \mathcal{N} \cup \mathcal{G} \cup \mathcal{C})_{\lessgtr_{\omega VC}}^{\leftarrow}$ is trivial.

Prove that $\mathcal{A} \cup \mathcal{N} \cup \mathcal{G} \cup \mathcal{C}$ commutes with $\mathcal{U}_{\lessgtr_{\omega VC}}^{\leftarrow}$.

Then apply Lemma D.14 to show that $\mathcal{U} \cup (\mathcal{A} \cup \mathcal{N} \cup \mathcal{G} \cup \mathcal{C})$ is skew confluent.

Domains

$$W \quad = \quad \mathbb{Z} + \langle W \rangle + (W \to W^*)$$
$$\langle W \rangle \quad = \quad \text{a finite tuple of values } W$$
$$Env \quad = \quad Ident \to W$$

Semantics of expressions and values

$$\mathcal{E}[\![e]\!] \quad : \quad Env \to W^*$$
$$\mathcal{E}[\![v]\!]\, \rho \quad = \quad unit(\mathcal{V}[\![v]\!]\, \rho)$$
$$\mathcal{E}[\![\mathbf{fail}]\!]\, \rho \quad = \quad empty$$
$$\mathcal{E}[\![e_1 \mathbin{|} e_2]\!]\, \rho \quad = \quad \mathcal{E}[\![e_1]\!]\, \rho \;\cupplus\; \mathcal{E}[\![e_2]\!]\, \rho$$
$$\mathcal{E}[\![e_1 = e_2]\!]\, \rho \quad = \quad \mathcal{E}[\![e_1]\!]\, \rho \;\cap\!\!\!\cap\; \mathcal{E}[\![e_2]\!]\, \rho$$
$$\mathcal{E}[\![e_1; e_2]\!]\, \rho \quad = \quad \mathcal{E}[\![e_1]\!]\, \rho \;\fatsemi\; \mathcal{E}[\![e_2]\!]\, \rho$$
$$\mathcal{E}[\![v_1\, v_2]\!]\, \rho \quad = \quad apply(\mathcal{V}[\![v_1]\!]\, \rho,\, \mathcal{V}[\![v_2]\!]\, \rho)$$
$$\mathcal{E}[\![\exists x.\, e]\!]\, \rho \quad = \quad \bigcup_{w \in W} \mathcal{E}[\![e]\!]\, (\rho[x \mapsto w])$$
$$\mathcal{E}[\![\mathbf{one}\{e\}]\!]\, \rho \quad = \quad one(\mathcal{E}[\![e]\!]\, \rho)$$
$$\mathcal{E}[\![\mathbf{all}\{e\}]\!]\, \rho \quad = \quad unit(all(\mathcal{E}[\![e]\!]\, \rho))$$

$$\mathcal{V}[\![v]\!] \quad : \quad Env \to W$$
$$\mathcal{V}[\![x]\!]\, \rho \quad = \quad \rho(x)$$
$$\mathcal{V}[\![k]\!]\, \rho \quad = \quad k$$
$$\mathcal{V}[\![op]\!]\, \rho \quad = \quad O[\![op]\!]$$
$$\mathcal{V}[\![\lambda x.\, e]\!]\, \rho \quad = \quad \lambda w.\mathcal{E}[\![e]\!]\, (\rho[x \mapsto w])$$
$$\mathcal{V}[\![\langle v_1, \cdots, v_n \rangle]\!]\, \rho \quad = \quad \langle \mathcal{V}[\![v_1]\!]\, \rho, \cdots, \mathcal{V}[\![v_n]\!]\, \rho \rangle$$

$$O[\![op]\!] \quad : \quad W$$
$$O[\![\mathbf{add}]\!] \quad = \quad \lambda w.\, \mathbf{if}\ (w = \langle k_1, k_2 \rangle)\ \mathbf{then}\ unit(k_1 + k_2)\ \mathbf{else}\ \text{WRONG}$$
$$O[\![\mathbf{gt}]\!] \quad = \quad \lambda w.\, \mathbf{if}\ (w = \langle k_1, k_2 \rangle \wedge k_1 > k_2)\ \mathbf{then}\ unit(k_1)\ \mathbf{else}\ empty$$
$$O[\![\mathbf{int}]\!] \quad = \quad \lambda w.\, \mathbf{if}\ (w = k)\ \mathbf{then}\ unit(k)\ \mathbf{else}\ empty$$

$$apply \quad : \quad (W \times W) \to W^*$$
$$apply(k, w) \quad = \quad \text{WRONG} \qquad\qquad k \in \mathbb{Z}$$
$$apply(\langle v_0, \cdots, v_n \rangle, k) \quad = \quad unit(v_k) \qquad\qquad 0 \leqslant k \leqslant n$$
$$= \quad empty \qquad\qquad \text{otherwise}$$
$$apply(f, w) \quad = \quad f(w) \qquad\qquad f \in W \to W^*$$

Fig. 19. Expression semantics

E   A DENOTATIONAL SEMANTICS FOR $\mathcal{VC}$

It is highly desirable to have a denotational semantics for $\mathcal{VC}$. A denotational semantics says directly what an expression *means* rather than how it *behaves*, and that meaning can be very perspicuous. Equipped with a denotational semantics we can, for example, prove that the left hand side and right hand side of each rewrite rule have the same denotation; that is, the rewrites are meaning-preserving.

Domains

$W^* = (\text{WRONG} + \mathcal{P}(W))_\bot$

Operations

Empty	$empty$	:	W^*	
	$empty$	=	$\{\,\}$	
Unit	$unit$	:	$W \to W^*$	
	$unit(w)$	=	$\{w\}$	
Union	$\uplus$	:	$W^* \to W^* \to W^*$	
	$s_1 \uplus s_2$	=	$s_1 \cup s_2$	
Intersection	$\sqcap$	:	$W^* \to W^* \to W^*$	
	$s_1 \sqcap s_2$	=	$s_1 \cap s_2$	
Sequencing	$\,\fatsemi\,$	:	$W^* \to W^* \to W^*$	
	$s_1 \fatsemi s_2$	=	s_2	if s_1 is non-empty
		=	$\{\,\}$	otherwise
One	one	:	$W^* \to W^*$	The result is either empty or a singleton
	$one(s)$	=	???	
All	all	:	$W^* \to \langle W \rangle$	
	$all(s)$	=	???	

All operations over W^* implicitly propagate $\bot$ and WRONG. E.g.

$s_1 \uplus s_2$	=	$\bot$	if $s_1 = \bot$ or $s_2 = \bot$
	=	WRONG	if $(s_1 = \text{WRONG and } s_2 \neq \bot)$ or $(s_2 = \text{WRONG and } s_1 \neq \bot)$
	=	$s_1 \cup s_2$	otherwise

Fig. 20. Set semantics for W^*

But a denotational semantics for a functional logic language is tricky. Typically one writes a denotation function something like

$$\mathcal{E}[\![e]\!] : Env \to W$$

where $Env = Ident \to W$. So $\mathcal{E}$ takes an expression e and an environment $\rho : Env$ and returns the value, or denotation, of the expression. The environment binds each free variable of e to its value. But what is the semantics of $\exists x.\, e$? We need to extend ρ with a binding for x, but what is x bound to? In a functional logic language x is given its value by various equalities scattered throughout e.

This section sketches our approach to this challenge. It is not finished work, and does not count as a contribution of our paper. We offer it because we have found it an illuminating alternative way to understand $\mathcal{VC}$, one that complements the rewrite rules that are the substance of the paper.

E.1 A first attempt at a denotational semantics

Our denotational semantics for $\mathcal{VC}$ is given in Fig. 19.

- We have one semantic function (here $\mathcal{E}$ and $\mathcal{V}$) for each syntactic non terminal (here e and v respectively.)
- Each function has one equation for each form of the construct.

- Both functions take an environment ρ that maps in-scope identifiers to a *single* value; see the definition $Env = Ident \rightarrow W$.
- The value function $\mathcal{V}$ returns a *single value W*, while the expression function $\mathcal{E}$ returns a *collection of values W^** (Appendix E.1).

The semantics is parameterized over the meaning of a "collection of values W^*". To a first approximation, think of W^* a (possibly infinite) set of values W, with union, intersection etc having their ordinary meaning.

Our first interpretation, given in Figure 20, is a little more refined: W^* includes $\perp$ and WRONG as well as a set of values. Our second interpretation is given in Figure 21, and discussed in Appendix E.4.

The equations themselves, in Fig. 19 are beautifully simple and compositional, as a denotational semantics should be.

The equations for $\mathcal{V}$ are mostly self-explanatory, but an equation like $\mathcal{V}[\![k]\!] \, \rho = k$ needs some explanation: the k on the left hand side (e.g. "3") is a piece of *syntax*, but the k on the right is the corresponding element of the *semantic world of values W* (e.g. 3). As is conventional, albeit a bit confusing, we use the same k for both. Same for op, where the semantic equivalent is the corresponding mathematical function.

The equations for $\mathcal{E}$ are more interesting.

- Values $\mathcal{E}[\![v]\!] \, \rho$: compute the single value for v, and return a singleton sequence of results. The auxiliary function *unit* is defined at the bottom of Fig. 19.
- In particular, values include lambdas. The semantics says that a lambda evaluates to a *singleton* collection, whose only element is a function value. But that function value has type $W \rightarrow W^*$; that is, it is a function that takes a single value and returns a *collection* of values.
- Function application $\mathcal{E}[\![v_1 \, v_2]\!] \, \rho$ is easy, because $\mathcal{V}$ returns a single value: just apply the meaning of the function to the meaning of the argument. The *apply* function is defined in Figure 19.
- Choice $\mathcal{E}[\![e_1 \mid e_2]\!] \, \rho$: take the union (written $\uplus$) of the values returned by e_1 and e_2 respectively. For bags this union operator is just bag union (Figure 20).
- Unification $\mathcal{E}[\![e_1 \mid e_2]\!] \, \rho$: take the *intersection* of the values returned by e_1 and e_2 respectively. For bags, this "intersection" operator $\Cap$ is defined in Fig. 20. In this definition, the equality is mathematical equality of functions; which we can't implement for functions; see Appendix E.1.
- Sequencing $\mathcal{E}[\![e_1; e_2]\!] \, \rho$. Again we use an auxiliary function $\fatsemi$ to combine the meanings of e_1 and e_2. For bags, the function $\fatsemi$ (Fig. 20 again) uses a bag comprehension. Again it does a cartesian product, but without the equality constraint of $\Cap$.
- The semantics of $(\mathbf{one}\{e\})$ simply applies the semantic function $one : W^* \rightarrow W^*$ to the collection of values returned by e. If e returns no values, so does $(\mathbf{one}\{e\})$; but if e returns one or more values, $(\mathbf{one}\{e\})$ returns the first. Of course that begs the question of what "the first" means – for bags it would be non-deterministic. We will fix this problem in Appendix E.4, but for now we simply ignore it.
- The semantics of $(\mathbf{all}\{e\})$ is similar, but it always returns a singleton collection (hence the *unit* in the semantics of $\mathbf{all}$) whose element is a (possibly-empty) tuple that contains all the values in the collection returned by e.

The fact that unification "=" maps onto intersection, and choice "$\mid$" onto union, is very satisfying.

The big excitement is the treatment of $\exists$. We must extend ρ, but what should we bind x to? (Compare the equation for $\mathcal{V}[\![\lambda x. e]\!]$, where we have a value w to hand.) Our answer is simple: *try all possible values, and union the results*:

$$\mathcal{E}[\![\exists x. e]\!] \, \rho \;=\; \bigcup_{w \in W} \mathcal{E}[\![e]\!] \, (\rho[x \mapsto w])$$

That $\bigcup_{w \in W}$ means: enumerate all values in $w \in W$, in some arbitrary order, and for each: bind x to w, find the semantics of e for that value of x, namely $\mathcal{E}[\![e]\!] \, (\rho[x \mapsto w])$, and take the union (in the sense of $\Cup$) of the results.

Of course we can't possibly implement it like this, but it makes a great specification. For example $\exists x. \, x = 3$ tries all possible values for x, but only one of them succeeds, namely 3, so the semantics is a singleton sequence [3].

E.2 The denotational semantics is un-implementable

This semantics is nice and simple, but we definitely can't implement it! Consider

$$\exists x. \, (x^2 - x - 6) = 0; \; x$$

The semantics will iterate over all possible values for x, returning all those that satisfy the equality; including 3, for example. But unless our implementation can guarantee to solve quadratic equations, we can't expect it to return 3. Instead it'll get stuck.

Another way in which the implementation might get stuck is through unifying functions:

$$(\lambda x. \, x + x) = (\lambda y. \, y * 2) \quad \text{or even} \quad (\lambda x. \, x + 1) = (\lambda y. \, y + 1)$$

But not all unification-over-functions is ruled out. We do expect the implementation to succeed with

$$\exists f. \, ((\lambda x. \, x + 1) = f); \; f \, 3$$

Here the $\exists$ will "iterate" over all values of f, and the equality will pick out the (unique) iteration in which f is bound to the incrementing function.

So our touchstone must be:

- If the implementation returns a value at all, it must be the value given by the semantics.
- Ideally, the verifier will guarantee that the implementation does not get stuck, or go WRONG.

E.3 Getting WRONG right

Getting WRONG right is a bit tricky.

- What is the value of $(3 = \langle\rangle)$? The intersection semantics would say *empty*, the empty collection of results, but we might want to say WRONG.
- Should WRONG be an element of W or of W^*? We probably want (**one**$\{3 \mid \textbf{wrong}\}$) to return a *unit*(3) rather then WRONG?
- What about $fst(\langle 3, \textbf{wrong}\rangle)$? Is that wrong or 3?

There is probably more than one possible choice here.

E.4 An order-sensitive denotational semantics

There is a Big Problem with this approach. Consider $\exists x. \, x = (4 \mid 3)$. The existential enumerates all possible values of x *in some arbitrary order*, and takes the union (*i.e.*, "concatentation") of the results from each of these bindings. Suppose that $\exists$ enumerates 3 before 4; then the semantics of this expression is the sequence $[3, 4]$, and not $[4, 3]$ as it should be. And yet returning a sequence (not a set nor a bag) is a key design choice in Verse. What can we do?

Figure 21 give a new denotational semantics that *does* account for order. The key idea (due to Joachim Breitner) is this: return a sequence of *labelled* values; and then sort that sequence (in *one* and *all*) into canonical order before exposing it to the programmer.

We do not change the equations for $\mathcal{E}$, $\mathcal{V}$, and O at all; they remain precisely as they are in Figure 19. However the semantics of a collection of values, W^*, does change, and is given in Figure 21:

Domains

$$
\begin{aligned}
W^* &= (\text{WRONG} + \mathcal{P}(LW))_\bot \\
W^? &= \{W\} \qquad\qquad \text{Set with 0 or 1 elements} \\
LW &= [L] \times W \qquad\quad \text{Sequence of } L \text{ and a value} \\
L &= \mathbf{L} + \mathbf{R}
\end{aligned}
$$

Operations

Empty	$empty$	:	W^*
	$empty$	=	$\varnothing$
Singleton	$unit$	:	$W \to W^*$
	$unit(w)$	=	$\{([], w)\}$
Union	$\uplus$	:	$W^* \to W^* \to W^*$
	$s_1 \uplus s_2$	=	$\{(\mathbf{L}:l, w) \mid (l, w) \in s_1\} \cup \{(\mathbf{R}:l, w) \mid (l, w) \in s_2\}$
Intersection	$\sqcap$	:	$W^* \to W^* \to W^*$
	$s_1 \sqcap s_2$	=	$\{(l_1 \bowtie l_2, w_1) \mid (l_1, w_1) \in s_1, (l_2, w_2) \in s_2, w_1 = w_2\}$
Sequencing	$\mathbin{\substack{\circ\\\circ}}$	:	$W^* \to W^* \to W^*$
	$s_1 \mathbin{\substack{\circ\\\circ}} s_2$	=	$\{(l_1 \bowtie l_2, w_2) \mid (l_1, w_1) \in s_1, (l_2, w_2) \in s_2\}$
One	one	:	$W^* \to W^*$
	$one(s)$	=	$head(sort(s))$
All	all	:	$W^* \to W^*$
	$all(s)$	=	$tuple(sort(s))$
Head	$head$	:	$([W] + \text{WRONG}) \to W^*$
	$head(\text{WRONG})$	=	WRONG
	$head[]$	=	$empty$
	$head(w : s)$	=	$unit(w)$
To tuple	$tuple$	:	$([W] + \text{WRONG}) \to \langle W \rangle$
	$tuple(\text{WRONG})$	=	WRONG
	$tuple[w_1, \cdots, w_n]$	=	$\langle w_1, \cdots, w_n \rangle$

$$
\begin{aligned}
\text{Sort} \quad sort &: LW^* \to ([W] + \text{WRONG})_\bot \\
sort(s) &= [] && \text{if } s \text{ is empty} \\
&= \text{WRONG} && \text{if } ws \text{ has more than one element} \\
&= ws && \text{otherwise} \\
&\quad \bowtie sort\{(l, w) \mid (\mathbf{L}:l, w) \in s\} \\
&\quad \bowtie sort\{(l, w) \mid (\mathbf{R}:l, w) \in s\} \\
&\text{where } ws = [w \mid ([], w) \in s]
\end{aligned}
$$

Fig. 21. Labelled set semantics for W^*

- A collection of values W^* is now $\perp$ or WRONG (as before), or a *set of labelled values*, each of type LW.
- A labelled value (of type LW) is just a pair $([L] \times W)$ of a *label* and a value.
- A label is a sequence of tags L, where a tag is just **L** or **R**, similar to Section 5.2.
- The union (or concatenation) operation $\uplus$, defined in Fig. 21, adds a **L** tag to the labels of the values in the left branch of the choice, and a **R** tag to those coming from the right. So the labels specify where in the tree the value comes from.
- Sequencing $\mathbin{\S}$ and $\mathbin{\owedge}$ both concatenate the labels from the values they combine.
- Finally *sort* puts everything in the "right" order: first the values with an empty label, then the values whose label starts with **L** (notice the recursive sort of the trimmed-down sequence), and then those that start with **R**. Notice that *sort* removes all the labels, leaving just a bare sequence of values W^*.
- Note that if *sort* encounters a set with more than one unlabelled element then this considered WRONG. This makes ambiguous expressions, like **one**$\{\exists x.\, x\}$, WRONG.

Let us look at our troublesome example $\exists x.\, x = (4 \mid 3)$, and assume that $\exists$ binds x to 3 and then 4. The meaning of this expression will be

$$\mathcal{E}[\![\exists x.\, x = (4 \mid 3)]\!]\,\epsilon \quad = \quad [(\mathbf{R}, 3), (\mathbf{L}, 4)]$$

Now if we take **all** of that expression we will get a singleton sequence containing $\langle 4, 3 \rangle$, because **all** does a sort, stripping off all the tags.

$$\mathcal{E}[\![\mathbf{all}\{\exists x.\, x = (4 \mid 3)\}]\!]\,\epsilon \quad = \quad [([], \langle 4, 3 \rangle)]$$

E.5 Related work

[Christiansen et al. 2011] gives another approach to a denotational semantics for a functional logic language. We are keen to learn of others.

F UPDATEABLE REFERENCES

The full Verse language has updatable references (à la ML). There are three new primitive operations, **alloc**, **read**, and **write**. The **alloc** creates a new reference with an initial value, **read** extracts the value from a reference, and **write** sets the value of a referene.

Modifying these references is transactional in the sense that if a computation fails, then any updates will not be visible outside the construct that handles the failures. *E.g.*,
$r := \textbf{alloc}(0);\ (\textbf{if}\ (\textbf{write}\langle r, 1\rangle;\ \textbf{fail})\ \textbf{then}\ 1\ \textbf{else}\ 2);\ \textbf{read}(r)$
will have the value 0, because the **write** is part of an expression that fails, and so its effect is not visible.

To add updateable references we extend the system with syntax and rules from figure 22. The **store** h **in** $\{e\}$ indicates that e should be reduced using the heap h. A heap is simply a mapping from references to values (one mapping being $r \mapsto v$). A reference is some opaque type that supports equality (unification) and creation of a new reference.

The interaction of the new primitives with the store can be seen from the axioms. The **alloc**(v) operation creates a new reference and adds a binding with v to the store. The **read**(r) operation retrieves the value for reference r from the store, and **write**$\langle r, v\rangle$ updates the reference r with v in the store. All of these operations use the context S which ensures that there are no store operations to the left of the hole, *i.e.*, a store operation in the hole is the next one that should execute.

The interesting rules involve choice and **split** because store operations are transactional in the sense that when an expressions fails, none of its store operations will happen.

When reducing **split**$\{e\}\langle f, g\rangle$ in an S hole, rule st-split-dup, the store is duplicated. Any store operations inside the **split** will happen in this local copy of the store. Note the two occurrences of h in the right hand side of st-split-dup. If the reduction of e results in **fail** then rule fail-elim is used, and the store from the failing computation is simply thrown away. If the reduction of e results in a value (with or without more alternatives) then rule st-split is used. This rule replaces the outer store with the inner store, since we know the inner computation has succeeded.

Similarly, the reduction of $e_1 \mathbin{\|} e_2$ will duplicate the store into the first branch, st-choice-dup. Here e_1 must not contain any store operation nor be a value. And again, similarly, st-choice commits the new store and throws away the old.

The use of oe in the rules is to ensure that the rules cannot get stuck in a loop. Using e instead of oe would mean that failing or committing would make the expression match the duplication rule again. It also prevents the duplication rule from repeatedly duplicating the **store**.

Note that **store** is part of the X context, which means that the **store** can float inside existentials. This is necessary for the store rules to fire since the S context does not allow going under existentials.

The semantics of **for**(d) **do** e with respect to store effects is somewhat intricate. The expression d is possibly multi-valued; any effects that happens when computing the first value of d will be visible the first time e is computed. Both these effects are then visible when computing the second value of d, and so on. If any iteration of d fails, then the effects of that computation are not visible outside d. This means that the desugaring of **for** into **split** needs to be more elaborate.
$$\textbf{for}(\exists x_1 \cdots x_n.\ d)\ \textbf{do}\ e$$
means
$$f\,\langle\rangle := \langle\rangle;$$
$$g(v)(r) := (v = \langle x_1, \cdots, x_n\rangle;\ cons\langle e, \textbf{split}\{r\,\langle\rangle\}\langle f, g\rangle\rangle);$$
$$\textbf{split}\{\exists x_1 \cdots x_n.\ d;\ \langle x_1, \cdots, x_n\rangle\}\langle f, g\rangle$$
To support limited store operations (*e.g.*, **read**, but not **write**) we can equip the store with a set of currently allowed operations. We also need some extra primitives that modify this set.

Syntax extension

References	r		
Expressions	e	$::=$	$\cdots \mid$ **store** h **in** $\{e\}$
Primops	op	$::=$	$\cdots \mid$ **alloc** $\mid$ **read** $\mid$ **write**
Head values	hnf	$::=$	$\cdots \mid r$
Execution contexts	X	$::=$	$\cdots \mid$ **store** h **in** $\{X\}$
Scope contexts	SC	$::=$	$\cdots \mid$ **store** h **in** $\{SC\}$
Heap	h	$::=$	$\epsilon \mid r \mapsto v, h$
Heap context	H	$::=$	$\square, h \mid r \mapsto v, H$
Store contexts	S	$::=$	$\square \mid v = S \mid S; e \mid se; S \mid \exists x. S$
Store-op free exprs	se	$::=$	$v \mid se_1 = se_2 \mid se_1; se_2 \mid \exists x. se \mid sp(v)$
Results	w	$::=$	$v \mid v \,\|\, e$
Non-store primops	sp	$::=$	any, except **alloc**, **read**, **write**
Non-store expression	oe	$::=$	like e, but not w, **store**, or **fail**

Axiom extensions

Normalization change

EXI-FLOAT $X[\exists x. e] \longrightarrow \exists x. X[e]$ if $X \neq \square$, $x \notin \text{fvs}(X)$, use α
                                                                      if there is **store** in X then $e \in ce$

Reference ops

REF-ALLOC   **store** h **in** $\{S[\textbf{alloc}(v)]\} \longrightarrow$ **store** $r \mapsto v, h$ **in** $\{S[r]\}$
                                                                                           $\text{fvs}(v)\#\text{bvs}(S), r$ fresh

REF-READ   **store** $H[r \mapsto v]$ **in** $\{S[\textbf{read}(r)]\} \longrightarrow$ **store** $H[r \mapsto v]$ **in** $\{S[v]\}$
                                                                                           $\text{fvs}(v)\#\text{bvs}(S)$, use α

REF-WRITE   **store** $H[r \mapsto v_1]$ **in** $\{S[\textbf{write}\langle r, v_2 \rangle]\} \longrightarrow$ **store** $H[r \mapsto v_2]$ **in** $\{S[\langle\rangle]\}$
                                                                                           $\text{fvs}(v_2)\#\text{bvs}(S)$

Store duplication

ST-SPLIT-DUP   **store** h **in** $\{S[\textbf{split}\{oe\}\langle f, g\rangle]\} \longrightarrow$ **store** h **in** $\{S[\textbf{split}\{\textbf{store}\ h\ \textbf{in}\ \{oe\}\}\langle f, g\rangle]\}$
                                                                                           $\text{fvs}(h)\#\text{bvs}(S)$, use α

ST-CHOICE-DUP   **store** h **in** $\{oe \,\|\, e\} \longrightarrow$ **store** h **in** $\{\textbf{store}\ h\ \textbf{in}\ \{oe\} \,\|\, e\}$

Store commit

ST-SPLIT   **store** h_1 **in** $\{S[\textbf{split}\{\textbf{store}\ h_2\ \textbf{in}\ \{w\}\}\langle f, g\rangle]\} \longrightarrow$ **store** h_2 **in** $\{S[\textbf{split}\{w\}\langle f, g\rangle]\}$
                                                                                           $\text{fvs}(h_2)\#\text{bvs}(S)$

ST-CHOICE   **store** h_1 **in** $\{S[(\textbf{store}\ h_2\ \textbf{in}\ \{w\}) \,\|\, e]\} \longrightarrow$ **store** h_2 **in** $\{S[w \,\|\, e]\}$
                                                                                           $\text{fvs}(h_2)\#\text{bvs}(S)$

Unification

Extension with the obvious axioms making equal references unify, and anything else fail.

Top level

Start top level reduction of e with **store** ϵ **in** $\{e\}$.

Fig. 22. The Verse calculus: store axioms