



SPARCL: A Language for Partially-Invertible Computation

KAZUTAKA MATSUDA, Tohoku University, Japan

MENG WANG, University of Bristol, United Kingdom

Invertibility is a fundamental concept in computer science, with various manifestations in software development (serializer/deserializer, parser/printer, redo/undo, compressor/decompressor, and so on). Full invertibility necessarily requires bijectivity, but the direct approach of composing bijective functions to develop invertible programs is too restrictive to be useful. In this paper, we take a different approach by focusing on *partially-invertible* functions—functions that become invertible if some of their arguments are fixed. The simplest example of such is addition, which becomes invertible when fixing one of the operands. More involved examples include entropy-based compression methods (e.g., Huffman coding), which carry the occurrence frequency of input symbols (in certain formats such as Huffman tree), and fixing this frequency information makes the compression methods invertible.

We develop a language SPARCL for programming such functions in a natural way, where partial-invertibility is the norm and bijectivity is a special case, hence gaining significant expressiveness without compromising correctness. The challenge in designing such a language is to allow ordinary programming (the “partially” part) to interact with the invertible part freely, and yet guarantee invertibility by construction. The language SPARCL is linear-typed, and has a type constructor to distinguish data that are subject to invertible computation and those that are not. We present the syntax, type system, and semantics of the language, and prove that SPARCL correctly guarantees invertibility for its programs. We demonstrate the expressiveness of SPARCL with examples including tree rebuilding from preorder and inorder traversals and Huffman coding.

CCS Concepts: • **Software and its engineering** → **Functional languages; Domain specific languages.**

Additional Key Words and Phrases: reversible computation, linear types

ACM Reference Format:

Kazutaka Matsuda and Meng Wang. 2020. SPARCL: A Language for Partially-Invertible Computation. *Proc. ACM Program. Lang.* 4, ICFP, Article 118 (August 2020), 31 pages. <https://doi.org/10.1145/3409000>

1 INTRODUCTION

Invertible computation, also known as reversible computation in physics and more hardware-oriented contexts, is a fundamental concept in computing. It involves computations that run both forwards and backwards so that the forward/backward semantics form a bijection. (In this paper, we *do not* concern ourselves with the totality of functions. We call a function a bijection if it is bijective on its actual domain and range, instead of its formal domain and codomain.) Early studies of invertible computation arise from the effort to reduce heat dissipation caused by information-loss in the traditional (unidirectional) computation model [Landauer 1961]. More modern interpretations of the problem include software concerns that are not necessarily connected to the physical realization. Examples of such include developing pairs of programs that are each other’s inverses: serializer and deserializer [Kennedy and Vytiniotis 2012], parser and printer [Matsuda and Wang 2013, 2018b;

118

Authors’ addresses: Kazutaka Matsuda, Tohoku University, 6-3-09 Aramaki, Aza-Aoba, Aoba-ku, Sendai, Japan, kztok@ecei.tohoku.ac.jp; Meng Wang, University of Bristol, BS8 1TH, Bristol, United Kingdom, meng.wang@bristol.ac.uk.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2020 Copyright held by the owner/author(s).

2475-1421/2020/8-ART118

<https://doi.org/10.1145/3409000>

Rendel and Ostermann 2010], compressor and decompressor [Srivastava et al. 2011], and also auxiliary processes in other program transformations such as bidirectionalization [Matsuda et al. 2007].

Invertible (reversible) programming languages are languages that offer primitive support to invertible computations. Examples include Janus [Lutz 1986; Yokoyama et al. 2008], Frank’s R [Frank 1997], Ψ -Lisp [Baker 1992], RFun [Yokoyama et al. 2011], Π/Π° [James and Sabry 2012] and Inv [Mu et al. 2004b]. The basic idea of these programming languages is to support deterministic forward and backward computation by local inversion: if a forward execution issues (invertible) commands c_1 , c_2 , and c_3 in this order, a backward execution issues corresponding inverse commands in the reverse order, as c_3^{-1} , c_2^{-1} , and c_1^{-1} . This design has a strong connection to the underlying physical reversibility, and is known to be able to achieve reversible Turing completeness [Bennett 1973]; i.e., all computable bijections can be defined.

However, this requirement of local invertibility does not always match how high-level programs are naturally expressed. As a concrete example, let us see the following toy program that computes the difference of two adjacent elements in a list, where the first element in the input list is kept in the output. For example, we have $\text{subs } [1, 2, 5, 2, 3] = [1, 1, 3, -3, 1]$.

```

subs :: [Int] → [Int]
subs xs = goSubs 0 xs
goSubs :: Int → [Int] → [Int]
goSubs _ [] = []
goSubs n (x : xs) = (x - n) : goSubs x xs

```

Despite being simple, these kind of transformations are nevertheless useful. For example, a function similar to *subs* can be found in the pre-processing step of image compression algorithms such as those used for PNG.¹ Another example is the encoding of bags (multisets) of integers, where *subs* can be used to convert sorted lists to lists of integers without any constraints [Kennedy and Vytiniotis 2012].

The function *subs* is invertible. We can define its inverse as below.

```

subs-1 :: [Int] → [Int]
subs-1 ys = goSubs' 0 ys
goSubs' :: Int → [Int] → [Int]
goSubs' _ [] = []
goSubs' n (y : ys) = let x = y + n in x : goSubs' x ys

```

However, *subs* cannot be expressed directly in existing reversible programming languages. The problem is that, though *subs* is perfectly invertible, its sub-component *goSubs* is not (its first argument is discarded in the empty-list case, and thus the function is not injective). Similar problems are also common in adaptive compression algorithms, where the model (such as a Huffman tree or a dictionary) grows in the same way in both compression and decompression, and the encoding process itself is only invertible after fixing the model at the point.

In the neighboring research area of program inversion, which studies program transformation techniques that derive f^{-1} from f ’s definition, functions like *goSubs* are identified as *partially* invertible. Note that this notion of partiality is inspired by partial evaluation, and *partial* inversion [Almendros-Jiménez and Vidal 2006; Nishida et al. 2005] allows static (or fixed) parameters whose values are known in inversion and therefore not required to be invertible (for example the first argument of *goSubs*). (To avoid potential confusion, in this paper, we avoid the use of “partial”

¹<https://www.w3.org/TR/2003/REC-PNG-20031110/>

when referring to totality, and use the phrase “not-necessarily-total” instead.) However, program inversion by itself does not readily give rise to a design of invertible programming language. Like most program transformations, program inversion may fail, and often for reasons that are not obvious to users. Indeed, the method by Nishida et al. [2005] fails for *subs*, and the outcome of Almendros-Jiménez and Vidal [2006]’s method depends on the processing order of the expressions.

In this paper, we propose a novel programming language SPARCL² that for the first time addresses the practical needs of partially invertible programs. The core idea of our proposal is based on a language design that allows invertible and conventional unidirectional computations, which are distinguished by types, to coexist and interact in a single definition. Specifically, inspired by Matsuda and Wang [2018c] our type system contains a special type constructor $(-)^{\bullet}$ (pronounced as “invertible”), where A^{\bullet} represents A -typed values that are subject to invertible computation. However, having invertible types like A^{\bullet} only solves half of the problem. For the applications we consider, exemplified by *subs*, the unidirectional parts (the first argument of *goSubs*) may depend on the invertible part (the second argument of *goSubs*), which complicates the design. (This is the very reason why Nishida et al. [2005]’s partial inversion fails for *subs*. In other words, a binding-time analysis [Gomard and Jones 1991] is not enough [Almendros-Jiménez and Vidal 2006].) This interaction demands conversion from invertible values of type A^{\bullet} to ordinary ones of type A , which only becomes feasible when we leverage the fact that such values can be seen as static (in the sense of partial inversion [Almendros-Jiménez and Vidal 2006]) if the values are already known in both forward and backward directions. The nature of reversibility suggests linearity or relevance [Walker 2004], as discarding of inputs is intrinsically irreversible. In fact, reversible functional programming languages [Baker 1992; James and Sabry 2012; Matsuda and Wang 2013; Mu et al. 2004b; Yokoyama et al. 2011] commonly assume a form of linearity or relevance, and in SPARCL this assumption is made explicit by a linear type system based on λ^q , (the core system of Linear Haskell [Bernardy et al. 2018]). As a teaser, an invertible version of *subs* in SPARCL is as below.³

```

subs : (List Int)•  $\multimap$  (List Int)•
subs xs = goSubs 0 xs

goSubs : Int  $\rightarrow$  (List Int)•  $\multimap$  (List Int)•
goSubs _ Nil• = Nil• with null
goSubs n (Cons x xs)• =
  let (x, r)• = pin x ( $\lambda x'.goSubs x' xs$ ) in    -- x' : Int is a static version of x : Int•.
  Cons• (sub n x) r with not  $\circ$  null
sub : Int  $\rightarrow$  Int•  $\multimap$  Int•
sub n = lift ( $\lambda x.x - n$ ) ( $\lambda x.x + n$ )

```

In SPARCL, invertible functions from A to B are represented as functions of type $A^{\bullet} \multimap B^{\bullet}$, where \multimap is the constructor for linear functions. Partial invertibility is conveniently expressed by taking additional parameters as in $\text{Int} \rightarrow \text{Int}^{\bullet} \multimap \text{Int}^{\bullet}$ and $\text{Int} \rightarrow (\text{List Int})^{\bullet} \multimap (\text{List Int})^{\bullet}$. The `pin` : $A^{\bullet} \multimap (A \rightarrow B^{\bullet}) \multimap (A \otimes B)^{\bullet}$ operator converts invertible objects into unidirectional ones. It captures a snapshot of its invertible argument and uses the snapshot as a static value in the body to create a safe local scope for the recursive call. Both the invertible argument and evaluation result of the body are returned as the output to preserve invertibility. The `with` conditions associated with the branches can be seen as postconditions which will be used for invertible case-branching. We leave the detailed discussion of the language constructs to later sections, but would like to highlight the

²The name stands for “a system for partially-reversible computation with linear types”.

³We use a Haskell-like syntax in this paper for readability, although our prototype implementation (<https://github.com/kztkm/sparcl>) uses simple non-indentation-sensitive syntax that requires more keywords for parsing.

fact that looking beyond the surface syntax, the definition is identical in structure to how *subs* is defined in a conventional language: *goSubs* has the same recursive pattern with two cases for empty and non-empty lists. This close resemblance to the conventional programming style is what we strive for in the design of SPARCL.

What SPARCL brings to the table is bijectivity guaranteed by construction (potentially with partially-invertible functions as auxiliary functions). We can run SPARCL programs in both directions, for example as below, and it is guaranteed that `fwd e v` results in v' if and only if `bwd e v'` results in v (`fwd` and `bwd` are primitives for forward and backward executions).

```
> fwd subs [1, 2, 5, 2, 3]
[1, 1, 3, -3, 1]
> bwd subs [1, 1, 3, -3, 1]
[1, 2, 5, 2, 3]
```

This guarantee of bijectivity is clearly different from the case of (functional) logic programming languages such as Prolog and Curry. Those languages relies on (lazy) generate-and-test [Antoy et al. 2000] to find inputs corresponding to a given output, a technique that may be adopted in the context of inverse computation [Abramov et al. 2006]. However, the generate-and-test strategy has the undesirable consequence of making reversible programming less apparent: it is unclear to programmers whether they are writing bijective programs that may be executed deterministically. Moreover, lazy generation of inputs may cause unpredictable overhead, whereas in reversible languages [Baker 1992; Frank 1997; James and Sabry 2012; Lutz 1986; Mu et al. 2004b; Yokoyama et al. 2008, 2011] including SPARCL, the forward and backward executions of a program take the same steps.

In summary, our main contributions are:

- We design SPARCL, a novel invertible programming language that captures the notion of partial invertibility. It is the first language that handles both clear separation and close integration of unidirectional and invertible computations, enabling new ways of structuring invertible programs. We formally specify the syntax, type system and semantics of its core system named $\lambda_{\rightarrow}^{\text{PI}}$ (Section 3).
- We state and prove several properties about $\lambda_{\rightarrow}^{\text{PI}}$ (Section 3.6), including subject reduction, bijectivity, and reversible Turing completeness [Bennett 1973]. We do not state the progress property directly, which is implied by our definitional [Reynolds 1998] interpreter written in Agda (available from <https://github.com/kztk-m/sparcl-agda>).
- We demonstrate the utility of SPARCL with nontrivial examples: tree rebuilding from inorder and preorder traversals [Mu and Bird 2003], and a simplified version of Huffman coding (Section 4).

In addition, a prototype implementation of SPARCL is available from <https://github.com/kztk-m/sparcl>, which also contains more examples. All the artifacts are linked from the SPARCL web page (<https://bx-lang.github.io/EXHIBIT/sparcl.html>).

2 OVERVIEW

In this section, we informally introduce the essential constructs of SPARCL and demonstrate their use with small examples.

2.1 Linear-Typed Programming

Linearity (or weaker relevance) is commonly adopted in reversible functional languages [Baker 1992; James and Sabry 2012; Matsuda and Wang 2013; Mu et al. 2004b; Yokoyama et al. 2011] to

exclude non injective functions such as constant functions. SPARCL is no exception (we will revisit its importance in Section 2.3) and adopts a linear type system based on λ_{\rightarrow}^q (the core system of Linear Haskell [Bernardy et al. 2018]). A feature of the type system is its function type $A \rightarrow_p B$, where the arrow is annotated by the argument's multiplicity (1 or ω). Here, $A \rightarrow_1 B$ denotes *linear functions* that use the input *exactly once*, while $A \rightarrow_{\omega} B$ denotes *unrestricted functions* that have no restriction on the use of its input. The following are some examples of linear and unrestricted functions.

$$\begin{array}{lll} id : a \rightarrow_1 a & double : \text{Int} \rightarrow_{\omega} \text{Int} & const : a \rightarrow_1 b \rightarrow_{\omega} a \\ id \ x = x & double \ x = x + x & const \ x \ y = x \end{array}$$

The purpose of the type system is to ensure bijectivity. But having linearity alone is not sufficient. We will come back to this point after showing invertible programming in SPARCL. We write \rightarrow for \rightarrow_1 and simply \rightarrow for \rightarrow_{ω} . Readers who are familiar with linear-type systems that have the exponential operator ! [Wadler 1993] may view $A \rightarrow_{\omega} B$ as $!A \rightarrow B$.

A small deviation from the (simply-typed fragment of) λ_{\rightarrow}^q , is that SPARCL is equipped with rank-1 polymorphism with qualified typing [Jones 1995] and type inference [Matsuda 2020]. For example, the system infers the following types for the following functions.

$$\begin{array}{lll} id : a \rightarrow_p a & const : a \rightarrow_p b \rightarrow a & app : (p \leq q) \Rightarrow (a \rightarrow_p b) \rightarrow_r a \rightarrow_q b \\ id \ x = x & const \ x \ y = x & app \ f \ x = f \ x \end{array}$$

In first two examples, p is arbitrary (i.e., 1 or ω); in the last example, the predicate $p \leq q$ states an ordering of multiplicity, where $1 \leq \omega$.⁴ This predicate states that if an argument is linear then it cannot be passed to an unrestricted function, as an unrestricted function may use its argument arbitrary many times. A more in-depth discussion of the surface type system is beyond the scope of this paper, but note that unlike the implementation of Linear Haskell,⁵ there is no defaulting nor compromise on unification of multiplicities thanks to the use of qualified typing.

2.2 Multiplication

One of the simplest examples of partially-invertible programs is multiplication [Nishida et al. 2005]. Suppose that we have a datatype of natural numbers defined as below.

```
data Nat = Z | S Nat
```

In SPARCL, constructors have linear types: $Z : \text{nat}$ and $S : \text{Nat} \rightarrow \text{Nat}$.

We define multiplication in term of addition, which is also partially-invertible.⁶

$$\begin{array}{l} add : \text{Nat} \rightarrow \text{Nat}^{\bullet} \rightarrow \text{Nat}^{\bullet} \\ add \ Z \quad y = y \\ add \ (S \ x) \ y = S^{\bullet} \ (add \ x \ y) \end{array}$$

As mentioned in the introduction, we use the type constructor $(-)^{\bullet}$ to distinguish data that are subject to invertible computation (such as Nat^{\bullet}) and those that are not (such as Nat): when the latter is fixed, a partially invertible function is turned into a (not-necessarily-total) bijection, for example, $add \ n : \text{Nat}^{\bullet} \rightarrow \text{Nat}^{\bullet}$. (For those who read the paper with colors, the arguments of $(-)^{\bullet}$ are highlighted in dark red.) Values of $(-)^{\bullet}$ -types are constructed by *lifted constructors* such as $S^{\bullet} : \text{Nat}^{\bullet} \rightarrow \text{Nat}^{\bullet}$. In the forward direction, S^{\bullet} applies S to the input, and in the backward

⁴For curious readers, we note that the inequality predicate is sufficient for typing our core system (Section 3) where constructors have linear types [Matsuda 2020].

⁵Confirmed for commit 1c80dcb424e1401f32bf7436290dd698c739d906 at May 14, 2019.

⁶This type is an instance of the most general type $\text{Nat} \rightarrow_p \text{Nat}^{\bullet} \rightarrow_q \text{Nat}^{\bullet}$ of add ; recall that there is no problem in using unrestricted inputs only once. We want to avoid overly polymorphic functions for simplicity of presentation.

direction, it strips one S (and throws an error if Z is found). In general, since constructors by nature are always bijective (though not necessarily total in the backward direction), every constructor $C : \sigma_1 \multimap \dots \multimap \sigma_n \multimap \tau$ automatically give rise to a corresponding lifted version $C^\bullet : \sigma_1^\bullet \multimap \dots \multimap \sigma_n^\bullet \multimap \tau^\bullet$.

A partially invertible multiplication function can be defined by using *add* as below.

$$\begin{aligned} \text{mul} &: \text{Nat} \rightarrow \text{Nat}^\bullet \multimap \text{Nat}^\bullet \\ \text{mul } z \text{ Z}^\bullet &= \text{Z}^\bullet \quad \text{with } \text{isZ} \\ \text{mul } z (\text{S } x)^\bullet &= \text{add } z (\text{mul } z x) \text{ with } \text{not} \circ \text{isZ} \end{aligned}$$

An interesting feature in the *mul* program is the *invertible pattern matching* [Yokoyama et al. 2008] indicated by patterns Z^\bullet and $(S x)^\bullet$ (here, we annotate patterns instead of constructors). Invertible pattern matching is a branching mechanism that executes bidirectionally: the forward direction basically performs the standard pattern matching, the backward direction employs an additional **with** clause to determine the branch to be taken. For example, $\text{mul } n : \text{Nat}^\bullet \multimap \text{Nat}^\bullet$, in the forward direction values are matched against the forms Z and $S x$; in the backward direction, the **with** conditions are checked: if $\text{isZ} : \text{Nat} \rightarrow \text{Bool}$ returns True, the first branch is chosen, otherwise the second branch is chosen. When the second branch is taken, the backward computation of *add* n is performed, which essentially subtracts n , followed by recursively applying the backward computation of *mul* n to the result. As the last step, the final result is enclosed with S and returned. In other words, the backward behavior of *mul* n recursively tries to subtract n , and returns the count of successful subtractions.

In SPARCL, **with** conditions are exclusive, which is enforced at run-time by using **with** conditions as assertions. Specifically, the branch's **with** condition is a positive assertion while all the other branches' ones are negative assertions. Thus, the forward computation fails when the branch's **with** condition is not satisfied, or any of the other **with** conditions is also satisfied. This exclusiveness enables the backward computation to uniquely identify the branch [Lutz 1986; Yokoyama et al. 2008]. Sometimes we may omit the **with** condition of the last branch, as it can be inferred as the negation of the conjunction of all the others. For example, in the definition of *goSubs* the second branch's **with** condition is *not* \circ *null*. One could use sophisticated types such as refinement types to statically enforce exclusiveness instead of assertion checking. However, we stick to simple types in this paper as our primal goal is to establish the basic design of SPARCL.

An astute reader may wonder what bijection *mul* Z defines, as zero times n is zero for any n . In fact, it defines a non-total bijection that in the forward direction the domain of the function contains only Z , i.e., the trivial bijection between $\{Z\}$ and $\{Z\}$.

2.3 Why Linearity Itself is Insufficient but Still Matters

The primal role of linearity is to prohibit values from being discarded or copied, and SPARCL is no exception. However, linearity itself is insufficient for partially-invertible programming. It is true that a linear calculus concerning tensor products (\otimes) and linear functions (\multimap) (even with exponentials (!)) can be modelled in the Int-construction [Joyal et al. 1996] of the category of not-necessarily-total bijections [Abramsky 2005; Abramsky et al. 2002]. However, it is also known that such a system cannot be easily extended to include sum-types nor invertible pattern matching [Abramsky 2005, Section 7]. Moreover, linearity does not express partiality as in partially-invertible computations. These are the reasons why we separate the invertible world and the unidirectional world by using $(-)^*$, inspired by staged languages [Davies and Pfenning 2001; Moggi 1998; Nielson and Nielson 1992]. Readers familiar with staged languages may see A^* as invertible *code* of type A , which will be executed forwards or backwards at the second stage to compute or uncompute A -typed values.

On the other hand, $(-)^{\bullet}$ does not replace the need for linearity either. Without linearity, $(-)^{\bullet}$ -typed values may be discarded or duplicated, which may lead to non-bijection. Unlike discarding, the exclusion of duplication is debatable as the inverse of duplication can be given as equality check [Glück and Kawabe 2003]. So it is our design choice to exclude duplication (contraction) in addition to discarding (weakening) to avoid unpredictable failures that may be caused by the equality checks. Without contraction, users are still able to implement duplication for datatypes with decidable equality (see Section 4.1.3). However, this design requires duplication (and the potential of failing) to be made explicit, which improves the predictability of the system. Having explicit duplication is not uncommon in this context [Mu et al. 2004b; Yokoyama et al. 2011].

Another design choice we made is to admit types like $(A \multimap B)^{\bullet}$ and $(A^{\bullet})^{\bullet}$ to allow a neat formulation (avoiding the use of kinds and subkinding to distinguish types that can be used in $(-)^{\bullet}$ from general ones). Such types are not very useful though, as function- or invertible-typed values cannot be inspected during invertible computations.

2.4 Running Reversible Computation

SPARCL provides primitive functions to execute invertible functions in either directions: **fwd** : $(A^{\bullet} \multimap B^{\bullet}) \rightarrow A \rightarrow B$ and **bwd** : $(A^{\bullet} \multimap B^{\bullet}) \rightarrow B \rightarrow A$. For example, we have:

```
> fwd (add (S Z)) (S Z)           -- (1 +) 1
S (S Z)                          -- = 2
> bwd (add (S Z)) (S (S Z))      -- (1 +)-1 2
S Z                               -- = 1
> fwd (mul (S (S Z))) (S (S (S Z))) -- (2 ×) 3
S (S (S (S (S (S Z))))))         -- = 6
> bwd (mul (S (S Z))) (S (S (S (S (S (S Z)))))) -- (2 ×)-1 6
S (S (S Z))                      -- = 3
```

Of course, the forward and backward computations may not be total. For example, the following expression (legitimately) fails.

```
> bwd (mul (S (S Z))) (S (S (S Z))) -- (2 ×)-1 3
Runtime Error:...
```

The guarantee SPARCL offers is that derived bijections are total with respect to the functions' actual domains and ranges; i.e., **fwd** $e v$ results in u , then **bwd** $e u$ results in v , and vice versa (Section 3.6.2).

Linearity plays a role here. Linear calculi are considered *resource*-aware in the sense that linear variables will be lost once used. In our case, resources are A^{\bullet} -typed values, as A^{\bullet} represents (a part of) an input or (a part of) an output of a bijection being constructed, which must be retained throughout the computation. This is why the first argument of **fwd/bwd** is unrestricted rather than linear. Very roughly speaking, an expression that can be passed to an unrestricted function cannot contain linear variables, or “resources”. Thus, a function of type $A^{\bullet} \multimap B^{\bullet}$ passed to **fwd/bwd** cannot use any resources other than *one* value of type A^{\bullet} to produce *one* value of type B^{\bullet} . In other words, all and only information from A^{\bullet} is retained in B^{\bullet} , guaranteeing bijection. As a result, SPARCL's type system effectively rejects code like **bwd** $(\lambda x. Z^{\bullet})$ and **bwd** $(\lambda x. \text{if fwd } (\lambda ()^{\bullet}. x) () \text{ then } Z^{\bullet} \text{ else } Z^{\bullet})$ as x 's multiplicity is ω in both cases. In the former case, x is discarded and multiplicity in our system is either 1 or ω . In the latter case, x appears in the first argument of **fwd**, which is unrestricted.

2.5 Importing Existing Invertible Functions

Bijectivity is not uncommon in computer science or mathematics, and there already exist many established algorithms that are bijective. Examples include nontrivial results in number theory or category theory, and manipulation of primitive or sophisticated data structures such as Burrows-Wheeler transformations on suffix arrays.

Instead of (re)writing them in SPARCL, the language provides a mechanism to directly import existing bijections (as a pair of functions) to construct valid SPARCL programs: $\mathbf{lift} : (A \rightarrow B) \rightarrow (B \rightarrow A) \rightarrow A^\bullet \multimap B^\bullet$ converts a pair of functions into a function on $(-)^{\bullet}$ -typed values, expecting that the pair of functions form mutual inverses. For example, by \mathbf{lift} , we can define *addInt* as below

$$\begin{aligned} \mathit{addInt} &: \mathit{Int} \rightarrow \mathit{Int}^\bullet \multimap \mathit{Int}^\bullet \\ \mathit{addInt} \ n &= \mathbf{lift} \ (\lambda x. x + n) \ (\lambda x. x - n) \end{aligned}$$

The use of \mathbf{lift} allows one to create primitive bijections to be composed by the various constructs in SPARCL. Another interesting use of \mathbf{lift} is to implement in-language inversion.

$$\begin{aligned} \mathit{invert} &: (A^\bullet \multimap B^\bullet) \rightarrow (B^\bullet \multimap A^\bullet) \\ \mathit{invert} \ h &= \mathbf{lift} \ (\mathbf{bwd} \ h) \ (\mathbf{fwd} \ h) \end{aligned}$$

2.6 Composing Partially-Invertible Functions

Partially-invertible functions in SPARCL expect arguments of both $(-)^{\bullet}$ and non- $(-)^{\bullet}$ types, which sometimes makes the calling of such functions interesting. This phenomenon is particularly noticeable in recursive calls where values of type A^\bullet may need to be fed into function calls expecting values of type A . In this case, it becomes necessary to convert A^\bullet -typed values to A -typed one. To avoid the risk of violating invertibility, such conversions are carefully managed in SPARCL through a special function $\mathbf{pin} : A^\bullet \multimap (A \rightarrow B^\bullet) \multimap (A \otimes B)^\bullet$, inspired by the *depGame* function in Kennedy and Vytiniotis [2012] and reversible updates [Axelsen et al. 2007] in reversible imperative languages [Frank 1997; Glück and Yokoyama 2016; Lutz 1986; Yokoyama et al. 2008]. The function \mathbf{pin} creates a static snapshot of its first argument (A^\bullet) and uses the snapshot (A) in its second argument. Bijectivity is guaranteed as the original A^\bullet value is retained in the output $(A \otimes B)^\bullet$ together with the evaluation result of the second argument (B^\bullet). We will define the function \mathbf{pin} and formally state the correctness property in Section 3.

Let us revisit the example in Section 1. The partially-invertible version of *goSubs* can be implemented via \mathbf{pin} as below.

$$\begin{aligned} \mathit{goSubs} &: \mathit{Int} \rightarrow (\mathit{List} \ \mathit{Int})^\bullet \multimap (\mathit{List} \ \mathit{Int})^\bullet \\ \mathit{goSubs} \ _ \ \mathit{Nil}^\bullet &= \mathit{Nil}^\bullet \ \mathbf{with} \ \mathit{null} \\ \mathit{goSubs} \ n \ (\mathit{Cons} \ x \ xs)^\bullet &= (\mathbf{case} \ \mathbf{pin} \ x \ (\lambda x'. \mathit{goSubs} \ x' \ xs) \ \mathbf{of} \\ &\quad (x, r)^\bullet \rightarrow \mathit{Cons}^\bullet \ (\mathit{sub} \ n \ x) \ r \ \mathbf{with} \ \lambda _ . \mathit{True}) \ \mathbf{with} \ \mathit{not} \ \circ \ \mathit{null} \end{aligned}$$

Here, we used \mathbf{pin} to convert $x : \mathit{Int}^\bullet$ to $x' : \mathit{Int}$ in order to pass it to the recursive call of *goSubs*. In the backward direction, *goSubs* n executes as follows⁷.

$$\begin{aligned} &\mathbf{bwd} \ (\mathit{goSubs} \ 0) \ [1, 1, 3, -3, 1] \\ &= \{ \mathit{Cons} \ \mathit{branch} \ \mathit{is} \ \mathit{taken}; \ \mathit{Cons} \ (\mathit{sub} \ 0 \ x) \ r = [1, 1, 3, -3, 1] \implies x = 1, r = [1, 3, -3, 1]. \} \\ &\quad \mathit{Cons} \ 1 \ (\mathbf{bwd} \ (\mathit{goSubs} \ 1) \ [1, 3, -3, 1]) \\ &= \{ \mathit{Cons} \ \mathit{branch} \ \mathit{is} \ \mathit{taken}; \ \mathit{Cons} \ (\mathit{sub} \ 1 \ x) \ r = [1, 3, -3, 1] \implies x = 2, r = [3, -3, 1]. \} \\ &\quad \mathit{Cons} \ 1 \ (\mathit{Cons} \ 2 \ (\mathbf{bwd} \ (\mathit{goSubs} \ 2) \ [3, -3, 1])) \\ &= \{ \mathit{Cons} \ \mathit{branch} \ \mathit{is} \ \mathit{taken}; \ \mathit{Cons} \ (\mathit{sub} \ 2 \ x) \ r = [3, -3, 1] \implies x = 5, r = [-3, 1]. \} \end{aligned}$$

⁷This execution trace is (overly) simplified for illustration purpose. See Section 3.5 for the actual operational semantics.

<pre> subs : (List Int)[•] → (List Int)[•] subs xs = goSubs 0 xs goSubs : Int → (List Int)[•] → (List Int)[•] goSubs _ Nil[•] = Nil[•] with null goSubs n (Cons x xs)[•] = let (x, r)[•] = pin x (λx'.goSub x' xs) in Cons[•] (sub n x) r with not ∘ null sub : Int → Int[•] → Int[•] sub n = lift (λx.x - n) (λx.x + n) </pre> <p style="text-align: center;">(a) partially-invertible version</p>	<pre> subsF : (List Int)[•] → (List Int)[•] subsF xs = let (0, r)[•] = goSubsF 0[•] xs in r goSubsF : Int[•] → (List Int)[•] → (Int ⊗ List Int)[•] goSubsF n Nil[•] = (n, Nil[•])[•] with null ∘ snd goSubsF n (Cons x xs)[•] = let (x, r)[•] = goSubsF x xs in let (n, x')[•] = subF (n, x)[•] in (n, Cons[•] x' r)[•] with not ∘ null ∘ snd subF : (Int ⊗ Int)[•] → Int ⊗ Int[•] subF = lift (λ(n, x).(n, x - n)) (λ(n, x).(n, x + n)) </pre> <p style="text-align: center;">(b) fully-invertible version</p>
--	---

Fig. 1. Side-by-side comparison of partially-invertible (a) and fully-invertible (b) versions of *subs*

```

Cons 1 (Cons 2 (Cons 5 (bwd (goSubs 5) [-3, 1])))
= ...
= Cons 1 (Cons 2 (Cons 5 (Cons 2 (Cons 3 (bwd (goSubs 3) []))))))
= { Nil branch is taken. }
Cons 1 (Cons 2 (Cons 5 (Cons 2 (Cons 3 Nil)))) = [1, 2, 5, 2, 3]

```

Note that the first arguments of (recursive) calls of *goSubs* (which are static) have the same values (1, 2, 5, 2, and 3) in both forward/backward executions, distinguishing their uses from those of the invertible arguments. As one can see, *goSubs n* behaves exactly like the hand-written *goSubs'* in *subs*⁻¹ which is reproduced below.

```

goSubs' _ [] = []
goSubs' n (y : ys) = let x = y + n in x : goSubs' x ys

```

The use of **pin** commonly results in an invertible **case** with a single branch, as we see in *goSubs* above. We capture this pattern with an invertible **let** as a shorthand notation, which enables us to write **let** $p^{\bullet} = e_1$ in e_2 for **case** e_1 of $\{p^{\bullet} \rightarrow e_2$ with $\lambda_.$ True $\}$. The definition of *goSubs* shown in Section 1 uses this shorthand notation, which is reproduced in Fig. 1a.

We would like to emphasise that partial invertibility, as supported in SPARCL, is key to concise function definitions. In Fig. 1, we show side-by-side two versions of the same program written in the same language: the one on the left allows partial invertibility whereas the one on the right requires all functions (include the intermediate ones) to be fully-invertible (note the different types in the two versions of *goSubs* and *sub*). As a result, *goSubsF* is much harder to define and the code becomes fragile and error-prone. This advantage of SPARCL, which is already evident in this small example, becomes decisive when dealing with larger programs, especially those requiring complex manipulation of static values (for example, Huffman Coding in Section 4.2).

We end this section with a theoretical remark. One might wonder why $(-)^{\bullet}$ is not a monad. It is true that $(-)^{\bullet}$ forms a functor, but the functor is not endo. Recall that A^{\bullet} represents bijective code of type A ; that is, A^{\bullet} and its component A belong to different categories (though we have not formally described them).⁸ One then might wonder whether $(-)^{\bullet}$ is a relative monad [Altenkirch et al. 2010]. To form a relative monad, one needs to find a functor that has the same domain and codomain as (the functor corresponding to) $(-)^{\bullet}$. It is unclear whether there exists such a functor

⁸For curious readers, we note our conjecture that $(-)^{\bullet}$ corresponds to the Yoneda embedding for the CPO-enriched category of (strict) bijections, analogous to Moggi [1998], although denotational semantics is outside the scope of this paper.

other than $(-)^{\bullet}$ itself; in this case, the relative monad operations do not provide any additional expressive power.

2.7 Implementations

We have implemented a proof-of-concept interpreter for SPARCL including the linear type system, which is available from <https://github.com/kztk-m/sparcl>. The implementation adds two small but useful extensions to what is presented in this paper. First, the implementation allows non-linear constructors, such as $\text{MkUn} : a \rightarrow \text{Un } a$ which serves as $!$ and helps us to write a function that returns both linear and unrestricted results. Misusing such constructors in invertible pattern matching is guarded against by the type system (otherwise it may lead to discarding or copying of invertible values). Second, the implementation uses the first-match principle for both forward and backward computations. That is, both patterns and **with** conditions are examined from top to bottom. Recall also that the implementation uses a non-indentation-sensitive syntax for simplicity as mentioned in Section 1.

It is worth noting that the implementation uses Matsuda [2020]’s type inference to infer linear types effectively without requiring any annotations. Hence, the type annotations in this paper are more for documentation purposes.

As part of our effort to prove type safety (subject reduction and progress), we also produced a parallel implementation in Agda to serve as proofs (Section 3.6), available from <https://github.com/kztk-m/sparcl-agda>.

3 CORE SYSTEM: $\lambda_{\rightarrow}^{\text{PI}}$

This section introduces $\lambda_{\rightarrow}^{\text{PI}}$, the core system that SPARCL is built on. Our design mixes ideas of linear-typed programming and meta-programming. As mentioned in Section 2.1, the language is based on (the simple multiplicity fragment of) $\lambda_{\rightarrow}^{\text{q}}$ [Bernardy et al. 2018], and, as mentioned in Section 2.3, it is also two-staged [Moggi 1998; Nielson and Nielson 1992] with different meta and object languages. Specifically, the meta stage is a usual call-by-value language (i.e., *unidirectional*) and the object stage is an *invertible* language. By having the two stages, partial invertibility is made explicit in this formalization.

In what follows, we use a number of notational conventions. A vector notation \bar{t} denotes a sequence such as t_1, \dots, t_n or $t_1; \dots; t_n$, where each t_i can be of any syntactic category and the delimiter (such as “,” and “;”) can differ depending on the context; we also refer to the length of the sequence by $|\bar{t}|$. In addition, we may refer to an element in the sequence \bar{t} as t_i . A simultaneous substitution of x_1, \dots, x_n in t with s_1, \dots, s_n is denoted as $t[s_1/x_1, \dots, s_n/x_n]$, which may also be written as $t[\overline{s/x}]$.

3.1 Central Concept: Bijections at the Heart

The surface language of SPARCL is designed for programming partially invertible functions, which are turned into bijections (by fixing the static arguments) for execution. This fact is highlighted in the core system $\lambda_{\rightarrow}^{\text{PI}}$, where we have a primitive *bijection type* $A \rightleftharpoons B$, which is inhabited by bijections constructed from functions of type $A^{\bullet} \multimap B^{\bullet}$. Technically, having a dedicated bijection type facilitates reasoning. For example, we may now straightforwardly state that “values of a bijection type $A \rightleftharpoons B$ are bijections between A and B ” (Corollary 3.4).

Accordingly, the **fwd** and **bwd** functions for execution in SPARCL are divided into application operators \triangleright and \triangleleft that apply bijection-typed values and an **unlift** operator for constructing bijections from functions of type $A^{\bullet} \multimap B^{\bullet}$. For example, we have **unlift** ($\text{add } (S \ Z)$) : $\text{Nat} \rightleftharpoons \text{Nat}$ (where $\text{add} : \text{Nat} \rightarrow \text{Nat}^{\bullet} \multimap \text{Nat}^{\bullet}$ is defined in Section 2), and the bijection can be executed as **unlift** ($\text{add } (S \ Z)$) \triangleright

S Z resulting in S (S Z) and **unlift** (add (S Z)) \triangleleft S (S Z) resulting in S Z. In fact, the operators **fwd** and **bwd** are now derived in $\lambda_{\rightarrow}^{\text{PI}}$, as **fwd** = $\lambda_{\omega} h. \lambda_{\omega} x. \mathbf{unlift} \ h \triangleright x$ and **bwd** = $\lambda_{\omega} h. \lambda_{\omega} x. \mathbf{unlift} \ h \triangleleft x$.

3.2 Syntax

The syntax of $\lambda_{\rightarrow}^{\text{PI}}$ is given as below.

Expressions: $e ::= x \mid \lambda_{\pi} x. e \mid e_1 e_2 \mid C \bar{e} \mid \mathbf{case}_{\pi} e_0 \mathbf{of} \{ \overline{p \rightarrow e} \}$
 $\mid C^{\bullet} \bar{e} \mid \mathbf{case} e_0 \mathbf{of} \{ p^{\bullet} \rightarrow e \mathbf{with} e' \} \mid \mathbf{pin} e_1 e_2 \mid \mathbf{unlift} e \mid e_1 \triangleright e_2 \mid e_1 \triangleleft e_2$

Patterns: $p ::= C \bar{x}$

Multiplicities: $\pi ::= 1 \mid \omega$

There are two lines for the various constructs of expressions. The ones in the first line are standard except the annotations in λ and **case** that determine the multiplicity of the variables introduced by the binders: $\pi = 1$ means that the bound variable is linear, and $\pi = \omega$ means there is no restriction. These annotations are omitted in the surface language as they are inferred. The second line consists of constructs that deal with invertibility. As mentioned above, **unlift** e , $e_1 \triangleright e_2$, and $e_1 \triangleleft e_2$ handles bijections which can be used to encode **fwd** and **bwd** in SPARCL. We have already seen lifted constructors, invertible **case**, and **pin** in Section 2. For simplicity, we assume that **pin**, C and C^{\bullet} are fully-applied. Lifted constructor expressions $C^{\bullet} \bar{e}$ and invertible **cases** are basic invertible primitives in $\lambda_{\rightarrow}^{\text{PI}}$. They are enough to make our system reversible Turing complete [Bennett 1973] (Theorem 3.5); i.e., all bijections can be implemented in the language. For simplicity, we assume that patterns are non-overlapping both for unidirectional and invertible cases. We do not include **lift**, which imports external code into SPARCL, as it is by definition unsafe. Instead, we will discuss it separately in Section 3.7.

Different from conventional reversible/invertible programming languages, the constructs **unlift** (together with \triangleright and \triangleleft) and **pin** support communication between the unidirectional world and the invertible world. The **unlift** construct together with \triangleright and \triangleleft runs invertible computation in the unidirectional world. The **pin** operator is the key to partiality; it enables us to temporarily convert a value in the invertible world into a value in the unidirectional world.

3.3 Types

Types in $\lambda_{\rightarrow}^{\text{PI}}$ are defined as below.

$$A, B ::= \alpha \mid T \bar{A} \mid A \rightarrow_{\pi} B \mid A^{\bullet} \mid A \rightleftharpoons B$$

Here, α denotes a type variable, T denotes a type constructor, $A \rightarrow_{\pi} B$ is a function type annotated with the argument's multiplicity π , $(-)^{\bullet}$ marks invertibility, and $A \rightleftharpoons B$ is a bijection type.

Each type constructor T comes with a set of constructors C of type

$$C : A_1 \multimap A_2 \multimap \cdots \multimap A_n \multimap T \bar{\alpha}$$

with $\text{fv}(A_i) \in \{\bar{\alpha}\}$ for any i .⁹ Type variables α are only used for types of constructors in the language. For example, the standard multiplicative product \otimes and additive sum \oplus [Wadler 1993] are represented by the following constructors.

$$(-, -) : \alpha_1 \multimap \alpha_2 \multimap \alpha_1 \otimes \alpha_2 \quad \text{InL} : \alpha_1 \multimap (\alpha_1 \oplus \alpha_2) \quad \text{InR} : \alpha_2 \multimap (\alpha_1 \oplus \alpha_2)$$

⁹For simplicity, we assume a constructor can only have linear fields; extending our discussions to constructors with unrestricted field is straightforward.

We assume that the set of type constructors at least include \otimes and **Bool**, where **Bool** has the constructors **True** : **Bool** and **False** : **Bool**. Types can be recursive via constructors; for example, we can have a list type $\text{List } \alpha$ with the following constructors.

$$\text{Nil} : \text{List } \alpha \quad \text{Cons} : \alpha \multimap \text{List } \alpha \multimap \text{List } \alpha$$

We may write $\bar{A} \multimap B$ for $A_1 \multimap A_2 \multimap \dots \multimap A_n \multimap B$ (when n is zero, $\bar{A} \multimap B$ is B). We shall also instantiate constructors implicitly and write $C : \bar{A}' \multimap \top \bar{B}$ when there is a constructor $C : \bar{A} \multimap \top \bar{\alpha}$ and $A'_i = A_i[\bar{B}/\bar{\alpha}]$ for each i . Thus we assume all types in our discussions are closed.

Negative recursive types are allowed in our system, which, for example, enables us to define general recursions without primitive fixpoint operators. Specifically, via **F** with the constructor $\text{MkF} : (\text{F } \alpha \rightarrow \alpha) \multimap \text{F } \alpha$, we have a fixpoint operator as below.

$$\begin{aligned} \text{fix}_\pi &\triangleq \lambda_{\omega} f. \lambda_{\pi} a. (\lambda_{\omega} x. \lambda_{\pi} a. f (\text{out } x \ x) \ a) (\text{MkF } (\lambda_{\omega} x. \lambda_{\pi} a. f (\text{out } x \ x) \ a)) \ a \\ &\text{where } \text{out} \triangleq \lambda_1 x. \text{case}_1 \ x \ \text{of } \{ \text{MkF } \ t \rightarrow t \} \end{aligned}$$

Here, *out* has type $\text{F } C \rightarrow \text{F } C \rightarrow C$ for any C (in this case $C = A \rightarrow_\pi B$), and thus fix_π has type $((A \rightarrow_\pi B) \rightarrow (A \rightarrow_\pi B)) \rightarrow A \rightarrow_\pi B$.

The most special type in the language is A^\bullet , which is the invertible version of A . More specifically, the invertible type A^\bullet represents invertible code that are executed forwards and backwards at the second stage to compute and “uncompute” A -typed values. Values of type A^\bullet must be treated linearly, and can only be manipulated by invertible operations, such as lifted constructors, invertible pattern matching, and **pin**. To keep our type system simple, or more specifically single-kinded, we allow types like $(A \multimap B)^\bullet$ and $(A^\bullet)^\bullet$, while the category of (not-necessary-total) bijections are not closed and $\lambda_{\rightarrow}^{\text{PI}}$ has no third stage. These types do not pose any problem, as such components cannot be inspected in invertible computation by any means (except in **with** conditions, which are unidirectional, i.e., run at the first stage).

Note that we consider the primitive bijection types $A \rightleftharpoons B$ as separate from $(A \rightarrow B) \otimes (B \rightarrow A)$.

3.4 Typing Relation

A *typing environment* is a mapping from variables x to pairs of type A and its multiplicity π , meaning that x has type A and can be used π -many times. We write $x_1 :_{\pi_1} A_1, \dots, x_n :_{\pi_n} B_n$ instead of $\{x_1 \mapsto (A_1, \pi_1), \dots, x_n \mapsto (B_n, \pi_n)\}$ for readability, and write ε for the empty environment. Reflecting the two stages, we adopt a dual context system [Davies and Pfenning 2001], which has *unidirectional* and *invertible* environments, denoted by Γ and Θ respectively. This separation of the two is purely theoretical, for the purpose of facilitating reasoning when we interpret A^\bullet -typed expressions that are closed in unidirectional variables but may have free variables in Θ as bijections. In fact, our prototype implementation does not distinguish the two environments. For all invertible environments Θ , without the loss of generality we assume that the associated multiplicities must be 1, i.e., $\Theta(x) = (A_x, 1)$ for any $x \in \text{dom}(\Theta)$. Thus, we shall sometimes omit multiplicities for Θ . This assumption is actually an invariant in our system since any variables introduced in Θ must have multiplicity 1. We make this explicit in order to simplify the theoretical discussions. Moreover, we assume that the domains of Γ and Θ are disjoint.

Given two unidirectional typing environments Γ_1 and Γ_2 , we define the addition $\Gamma_1 + \Gamma_2$ as below.

$$(\Gamma_1 + \Gamma_2)(x) = \begin{cases} (A, \omega) & \text{if } \Gamma_1(x) = (A, _) \text{ and } \Gamma_2(x) = (A, _) \\ (A, \pi) & \text{if } \Gamma_i(x) = (A, \pi) \text{ and } x \notin \text{dom}(\Gamma_j) \text{ for some } i \neq j \in \{1, 2\} \end{cases}$$

If $\text{dom}(\Gamma_1)$ and $\text{dom}(\Gamma_2)$ are disjoint, we sometimes write Γ_1, Γ_2 instead of $\Gamma_1 + \Gamma_2$ to emphasize the disjointness. A similar addition applies to invertible environments. But as only multiplicity 1 is allowed in Θ , $\Theta_1 + \Theta_2 = \Theta$ implicitly implies $\text{dom}(\Theta_1) \cap \text{dom}(\Theta_2) = \emptyset$.

Typing Rules for Expressions $\boxed{\Gamma; \Theta \vdash e : A}$ **and Patterns** $\boxed{p : A \triangleright_{\pi} \Gamma}$

$$\begin{array}{c}
\frac{}{\omega\Gamma + x :_1 A; \omega\Theta \vdash x : A} \text{T-UVAR} \quad \frac{\Gamma, x :_{\pi} A; \Theta \vdash e : B}{\Gamma; \Theta \vdash \lambda_{\pi} x. e : A \rightarrow_{\pi} B} \text{T-ABS} \quad \frac{\Gamma_1; \Theta_1 \vdash e_1 : A \rightarrow_{\pi} B \quad \Gamma_2; \Theta_2 \vdash e_2 : A}{\Gamma_1 + \pi\Gamma_2; \Theta_1 + \pi\Theta_2 \vdash e_1 e_2 : B} \text{T-APP} \\
\\
\frac{n = |\bar{e}| = |\bar{A}| \quad C : \bar{A} \multimap T \bar{B} \quad \{\Gamma_i; \Theta_i \vdash e_i : A_i\}_i}{\omega\Gamma_0 + \Gamma_1 + \dots + \Gamma_n; \Theta_1 + \dots + \Theta_n \vdash C \bar{e} : T \bar{B}} \text{T-CON} \\
\\
\frac{\Gamma_0; \Theta_0 \vdash e_0 : A \quad \{p_i : A \triangleright_{\pi} \Gamma_i \quad \Gamma, \Gamma_i; \Theta \vdash e_i : B\}_i}{\pi\Gamma_0 + \Gamma; \pi\Theta_0 + \Theta \vdash \text{case}_{\pi} e_0 \text{ of } \{p \rightarrow e\} : B} \text{T-CASE} \\
\\
\frac{}{\omega\Gamma; x : A \vdash x : A^{\bullet}} \text{T-RVAR} \quad \frac{n = |\bar{e}| = |\bar{A}| \quad C : \bar{A} \multimap T \bar{B} \quad \{\Gamma_i; \Theta_i \vdash e_i : A_i^{\bullet}\}_i}{\omega\Gamma_0 + \Gamma_1 + \dots + \Gamma_n; \Theta_1 + \dots + \Theta_n \vdash C^{\bullet} \bar{e} : (T \bar{B})^{\bullet}} \text{T-RCON} \\
\\
\frac{\Gamma_0; \Theta_0 \vdash e_0 : A^{\bullet} \quad \{p_i : A \triangleright_1 \Theta_i \quad \Gamma; \Theta, \Theta_i \vdash e_i : B^{\bullet} \quad \Gamma'; \Theta' \vdash e'_i : B \rightarrow_{\omega} \text{Bool}\}_i}{\Gamma_0 + \Gamma + \omega\Gamma'; \Theta_0 + \Theta + \omega\Theta' \vdash \text{case } e_0 \text{ of } \{p^{\bullet} \rightarrow e \text{ with } e'\} : B^{\bullet}} \text{T-RCASE} \\
\\
\frac{\Gamma_1; \Theta_1 \vdash e_1 : A^{\bullet} \quad \Gamma_2; \Theta_2 \vdash e_2 : A \rightarrow_{\omega} B^{\bullet}}{\Gamma_1 + \Gamma_2; \Theta_1 + \Theta_2 \vdash \text{pin } e_1 e_2 : (A \otimes B)^{\bullet}} \text{T-PIN} \quad \frac{\Gamma; \Theta \vdash e : A^{\bullet} \rightarrow_1 B^{\bullet}}{\omega\Gamma; \omega\Theta \vdash \text{unlift } e : A \rightleftharpoons B} \text{T-UNLIFT} \\
\\
\frac{\Gamma_1; \Theta_1 \vdash e_1 : A \rightleftharpoons B \quad \Gamma_2; \Theta_2 \vdash e_2 : A}{\Gamma_1 + \omega\Gamma_2; \Theta_1 + \omega\Theta_2 \vdash e_1 \triangleright e_2 : B} \text{T-FAAPP} \quad \frac{\Gamma_1; \Theta_1 \vdash e_1 : A \rightleftharpoons B \quad \Gamma_2; \Theta_2 \vdash e_2 : B}{\Gamma_1 + \omega\Gamma_2; \Theta_1 + \omega\Theta_2 \vdash e_1 \triangleleft e_2 : A} \text{T-BAAPP} \\
\\
\frac{C : \bar{A} \multimap T \bar{B}}{C \bar{x} : T \bar{B} \triangleright_{\pi} x :_{\pi} A}
\end{array}$$

Fig. 2. Typing rules for expressions and patterns

We define multiplication of multiplicities as below.

$$1\pi = \pi 1 = \pi \quad \omega\pi = \pi\omega = \omega$$

For $\Gamma = x_1 :_{\pi_1} A_1, \dots, x_n :_{\pi_n} A_n$, we write $\pi\Gamma$ for the environment $x_1 :_{\pi\pi_1} A_1, \dots, x_n :_{\pi\pi_n} A_n$. A similar notation applies to invertible environments. Again, $\omega\Theta' = \Theta$ means that $\Theta' = \varepsilon$. Notice that it can hold that $\Gamma = \Gamma + \Gamma$ and $\Gamma = \omega\Gamma = 1\Gamma$ if $\Gamma(x) = (_, \omega)$ for all $x \in \text{dom}(\Gamma)$.

The typing relation $\Gamma; \Theta \vdash e : A$ reads that under environments Γ and Θ , expression e has type A (Fig. 2). The definition basically follows λ^q , [Bernardy et al. 2018] except having two environments for the two stages. Although multiplicities in Θ are always 1, some of the typing rules refers to $\omega\Theta$ (which implies $\Theta = \varepsilon$) in the conclusion parts, to emphasize that Γ and Θ are treated similarly by the rules. The idea underlying this type system is that, together with the operational semantics in Section 3.5, a term-in-context $\varepsilon; \Theta \vdash e : A^{\bullet}$ defines a piece of code representing a bijection between Θ and A , and hence $\varepsilon; \varepsilon \vdash e' : A \rightleftharpoons B$ defines a bijection between A and B (see Section 3.6). Our Agda implementation mentioned in Sections 1 and 2.7 follows this idea with some generalization.

Intuitively, the multiplicity of a variable represents the usage of a resource to be associated with the variable. Hence, multiplicities in Γ and Θ are synthesized rather than checked in typing. This viewpoint is useful for understanding T-APP and T-CASE; it is natural that, if an expression e is used π times, the multiplicities of variables in e are multiplied by π . Discarding variables, or weakening, is performed in the rules T-UVAR, T-RVAR, T-CON, and T-RCON which can be leaves in a derivation tree. Note that weakening is not allowed for Θ -variables as they are linear.

The typing rules for the invertible part would need additional explanation. In T-RVAR, x has type A^{\bullet} if the invertible typing environment is the singleton mapping $x : A$. One explanation for

this is that Θ represents the typing environment for the object (i.e., invertible) system. Another explanation is that we simply omit $(-)^{\bullet}$ as all variables in Θ must have types of the form A^{\bullet} . Rule T-RCON says that we can lift a constructor to the invertible world leveraging the injective nature of the constructor. Rule T-RCASE says that the invertible **case** is for pattern-matching on $(-)^{\bullet}$ -typed data; the pattern matching is done in the invertible world, and thus the bodies of the branches must also have $(-)^{\bullet}$ -types. Recall that the **with**-conditions (e'_i) are used for deciding which branch is used in backward computation. The type $B \rightarrow_{\omega} \text{Bool}$ indicates that they are conventional unrestricted functions, and $\omega\Gamma'$ and $\omega\Theta'$ in the conclusion part of the rule indicates that their uses are unconstrained. Notice that, since the linearity comes only from the use of $(-)^{\bullet}$ -typed values, there is little motivation to use linear variables to define conventional functions in $\lambda_{\rightarrow}^{\text{PI}}$. The operators **pin**, **unlift**, \triangleright , and \triangleleft are special in $\lambda_{\rightarrow}^{\text{PI}}$. The operator **pin** is simply a fully-applied version of the one in Section 2; so we do not repeat the explanation. Rules T-UNLIFT, T-FAPP, and T-BAPP are inherited from the types of **fwd** and **bwd** in Section 2. Recall that $\omega\Theta$ ensures $\Theta = \varepsilon$, and thus the arguments of **unlift** and constructed bijections must be closed in terms of invertible variables. It might look a little weird that $e_1 \triangleright e_2/e_1 \triangleleft e_2$ uses e_1 linearly; this is not a problem because Θ_1 in T-FAPP/T-BAPP must be empty for expressions that occur in evaluation (Lemma 3.2).

3.5 Operational Semantics

The semantics of $\lambda_{\rightarrow}^{\text{PI}}$ consists of three evaluation relations: *unidirectional*, *forward*, and *backward*. The unidirectional evaluation evaluates away the unidirectional constructs such as λ -abstractions and applications, and the forward and backward evaluation specifies bijections.

For example, let us consider an expression $e = (\lambda_{\omega} f. f (f y)) (\lambda_1 x. S^{\bullet} x)$. Due to λ -abstractions and function applications, it is not immediately clear how we can interpret the expression as a bijection. The unidirectional evaluation \Downarrow is used to evaluate these unidirectional constructs away to make way for the forward and backward evaluation to interpret the residual term. For the above expression, we have $e \Downarrow S^{\bullet} (S^{\bullet} y)$ where the residual $S^{\bullet} (S^{\bullet} y)$ is ready to be interpreted bijectively. The forward evaluation $\mu \vdash E \Rightarrow v$ evaluates a residual E under an environment μ to obtain a value v as usual. For example, we have $\{y \mapsto Z\} \vdash S^{\bullet} (S^{\bullet} y) \Rightarrow S (S Z)$. The backward evaluation $E \Leftarrow v \dashv \mu$ does the opposite; it inversely evaluates E to find an environment μ for a given value v , so that the corresponding forward evaluation of E returns the value for the environment. For example, we have $S^{\bullet} (S^{\bullet} y) \Leftarrow S (S Z) \dashv \{y \mapsto Z\}$.

This is the basic story, but computation can be more complicated in general. With **case** and **pin**, the forward \Rightarrow and backward \Leftarrow evaluation depend on the unidirectional evaluation \Downarrow ; and with \triangleright and \triangleleft , the unidirectional evaluation \Downarrow also depends on the forward \Rightarrow and backward \Leftarrow ones. Technically, the linear type system is also the key to the latter type of dependency, which is an important difference from related work in bidirectional programming [Matsuda and Wang 2018c].

3.5.1 Values and Residuals. We first define a set of *values* v and a set of *residuals* E as below.

$$\begin{aligned} \text{Values:} \quad v &::= \lambda_{\pi} x. e \mid C \bar{v} \mid E \mid \langle x. E \rangle \\ \text{Residuals:} \quad E &::= x \mid C^{\bullet} \bar{E} \mid \text{case } E_0 \text{ of } \overline{\{p^{\bullet} \rightarrow e \text{ with } \lambda_{\omega} x. e'\}} \mid \text{pin } E_1 (\lambda_{\omega} x_2. e_2) \end{aligned}$$

The residuals are $(-)^{\bullet}$ -typed expressions, which are subject to the forward and backward evaluations. The syntax of residuals makes it clear that branch bodies in invertible **cases** are not evaluated in the unidirectional evaluation; otherwise, recursive definitions involving them usually diverge. A variable is also a value. Indeed, our evaluation targets expressions/residuals that may be open in term of invertible variables. The value $\langle x. E \rangle$ represents a bijection. Intuitively, $\langle x. E \rangle$ is a single-holed residual E where the hole is represented by the variable x . The type system ensures that the x is the only variable in E so that E is ready to be interpreted as a bijection. Since $\langle x. E \rangle$ is not an

Unidirectional Evaluation $\boxed{e \Downarrow v}$

$$\begin{array}{c}
\frac{}{\lambda_{\pi}x.e \Downarrow \lambda_{\pi}x.e} \quad \frac{e_1 \Downarrow \lambda_{\pi}x.e \quad e_2 \Downarrow v_2 \quad e[v_2/x] \Downarrow v}{e_1 e_2 \Downarrow v} \quad \frac{\overline{e \Downarrow v}}{C \bar{e} \Downarrow C \bar{v}} \quad \frac{e_0 \Downarrow v_0 \quad p_i \mu = v_0 \quad e_i \mu \Downarrow v}{\text{case}_{\pi} e_0 \text{ of } \{p \rightarrow e\} \Downarrow v} \\
\frac{}{x \Downarrow x} \quad \frac{e \Downarrow \bar{E}}{C \bullet \bar{e} \Downarrow C \bullet \bar{E}} \quad \frac{e_0 \Downarrow E_0 \quad e' \Downarrow \overline{\lambda_{\omega}x.e''} \quad \alpha\text{-renaming to make } \text{fv}(p) \text{ fresh}}{\text{case } e_0 \text{ of } \{p \bullet \rightarrow e \text{ with } e'\} \Downarrow \text{case } E_0 \text{ of } \{p \bullet \rightarrow e \text{ with } \lambda_{\omega}x.e''\}} \\
\frac{e_1 \Downarrow E_1 \quad e_2 \Downarrow \lambda_{\omega}x_2.e'_2}{\text{pin } e_1 e_2 \Downarrow \text{pin } E_1 (\lambda_{\omega}x_2.e'_2)} \quad \frac{}{\langle x.E \rangle \Downarrow \langle x.E \rangle} \quad \frac{e \Downarrow \lambda_1x.e' \quad e'[y/x] \Downarrow E \quad y: \text{fresh}}{\text{unlift } e \Downarrow \langle y.E \rangle} \\
\frac{e_1 \Downarrow \langle x.E \rangle \quad e_2 \Downarrow v_2 \quad \{x \mapsto v_2\} \vdash E \Rightarrow v}{e_1 \triangleright e_2 \Downarrow v} \quad \frac{e_1 \Downarrow \langle x.E \rangle \quad e_2 \Downarrow v_2 \quad E \Leftarrow v_2 \vdash \{x \mapsto v\}}{e_1 \triangleleft e_2 \Downarrow v}
\end{array}$$

Forward Evaluation $\boxed{\mu \vdash E \Rightarrow v}$

$$\begin{array}{c}
\frac{}{\{x \mapsto v\} \vdash x \Rightarrow v} \quad \frac{\overline{\mu \vdash E \Rightarrow v}}{\uplus \bar{\mu} \vdash C \bullet \bar{E} \Rightarrow C \bar{v}} \quad \frac{\mu_1 \vdash E_1 \Rightarrow v_1 \quad e_2[v_1/x] \Downarrow E_2 \quad \mu_2 \vdash E_2 \Rightarrow v_2}{\mu_1 \uplus \mu_2 \vdash \text{pin } E_1 (\lambda_{\omega}x.e_2) \Rightarrow (v_1, v_2)} \\
\frac{\mu_0 \vdash E_0 \Rightarrow p_i \mu_i \quad \text{dom}(\mu_i) = \text{fv}(p_i) \quad e_i \Downarrow E_i \quad \mu \uplus \mu_i \vdash E_i \Rightarrow v \quad e'_i[v/x_i] \Downarrow \text{True} \quad \{e'_j[v/x_j] \Downarrow \text{False}\}_{j \neq i}}{\mu_0 \uplus \mu \vdash \text{case } E_0 \text{ of } \{p \bullet \rightarrow e \text{ with } \lambda_{\omega}x.e'\} \Rightarrow v}
\end{array}$$

Backward Evaluation $\boxed{E \Leftarrow v \vdash \mu}$

$$\begin{array}{c}
\frac{}{x \Leftarrow v \vdash \{x \mapsto v\}} \quad \frac{\overline{E \Leftarrow v \vdash \mu}}{C \bullet \bar{E} \Leftarrow C \bar{v} \vdash \uplus \bar{\mu}} \quad \frac{E_1 \Leftarrow v_1 \vdash \mu_1 \quad e_2[v_1/x] \Downarrow E_2 \quad E_2 \Leftarrow v_2 \vdash \mu_2}{\text{pin } E_1 (\lambda_{\omega}x.e_2) \Leftarrow (v_1, v_2) \vdash \mu_1 \uplus \mu_2} \\
\frac{e'_i[v/x_i] \Downarrow \text{True} \quad \{e'_j[v/x_j] \Downarrow \text{False}\}_{j \neq i} \quad e_i \Downarrow E_i \quad E_i \Leftarrow v \vdash \mu \uplus \mu_i \quad \text{dom}(\mu_i) = \text{fv}(p_i) \quad E_0 \Leftarrow p_i \mu_i \vdash \mu_0}{\text{case } E_0 \text{ of } \{p \bullet \rightarrow e \text{ with } \lambda_{\omega}x.e'\} \Leftarrow v \vdash \mu_0 \uplus \mu}
\end{array}$$

Fig. 3. Evaluation relations: unidirectional, forward and backward.

expression defined so far, we extend expressions to include this form as $e ::= \dots \mid \langle x.E \rangle$ together with the following typing rule:

$$\frac{\Gamma; \Theta, x : A \vdash E : B}{\omega\Gamma; \omega\Theta \vdash \langle x.E \rangle : A \Leftarrow B} \text{T-HOLED}$$

It is crucially important that x is added to the invertible environment.

3.5.2 Three Evaluation Relations: Unidirectional, Forward and Backward. The evaluation relations are shown in Fig. 3, which are defined by mutually-dependent evaluation rules.

The unidirectional evaluation is rather standard, except that it treats invertible primitives (such as lifted constructors, invertible **cases**, **lift**, and **pin**) as constructors. A subtlety is that we assume dynamic α -renaming of invertible **cases** to avoid variable capturing. The evaluation rules can evaluate open expressions by having $x \Downarrow x$; recall that residuals can contain variables. The **unlift** operator uses a fresh variable y in the evaluation to make a single-holed residual $\langle y.E \rangle$ as a representation of bijection. Such single-holed residuals can be used in the forward direction by $e_1 \triangleright e_2$ and in the backward direction by $e_1 \triangleleft e_2$, by triggering the corresponding evaluation.

The forward evaluation $\mu \vdash E \Rightarrow v$ states that under environment μ , a residual E evaluates to a value v , and the backward evaluation $E \Leftarrow v \vdash \mu$ inversely evaluates E to return the environment μ

from a value v : the forward and backward evaluation relations form a bijection. For variables and invertible constructors, both forward and backward evaluation rules are rather straightforward. The rules for invertible **case** expression are designed to ensure that every branch taken in one direction *may* and *must* be taken in the other direction too. This is why we check the **with** conditions even in the forward evaluation: the condition is considered as a post-condition that must exclusively hold after the evaluation of a branch. The **pin** operator changes the behavior of the backward computation of the second argument based on the result of the first argument; notice that v_1 , the parameter for the second argument, is obtained as the evaluation result of the first argument in the forward evaluation, and as the first component of the result pair in the backward evaluation. Notice that the unidirectional evaluation \Downarrow involved in the presented evaluation rules are performed in the same way in both evaluation, which is the key to bijectivity of E .

3.6 Metatheory

In this subsection, we present the key properties about $\lambda_{\Downarrow}^{\text{PI}}$.

3.6.1 Subject Reduction. First, we show a substitution lemma for $\lambda_{\Downarrow}^{\text{PI}}$. We only need to consider substitution for unidirectional variables because substitution for invertible variables never happens in evaluation; recall that we use environments (μ) in the forward and backward evaluation.

Lemma 3.1. $\Gamma, x :_{\pi} A; \Theta \vdash e : B$ and $\Gamma'; \Theta' \vdash e' : A$ implies $\Gamma + \pi\Gamma'; \Theta + \pi\Theta' \vdash e[e'/x] : B$. \square

Note that the substitution is only valid when $\Theta + \pi\Theta'$ satisfy the assumption that invertible variables have multiplicity 1. This assumption is guaranteed by typing.

Then, by Lemma 3.1, we have the subject reduction properties as follows:

Lemma 3.2 (subject reduction). The following properties hold:

- Suppose $\varepsilon; \Theta \vdash e : A$ and $e \Downarrow v$. Then, $\varepsilon; \Theta \vdash v : A$ holds.
- Suppose $\varepsilon; \Theta \vdash E : A^{\bullet}$ and $\mu \vdash E \Rightarrow v$. Then, $\text{dom}(\Theta) = \text{dom}(\mu)$ holds, and $\varepsilon; \varepsilon \vdash \mu(x) : \Theta(x)$ for all $x \in \text{dom}(\Theta)$ implies $\varepsilon; \varepsilon \vdash v : A$.
- Suppose $\varepsilon; \Theta \vdash E : A^{\bullet}$ and $E \Leftarrow v \dashv \mu$. Then, $\text{dom}(\Theta) = \text{dom}(\mu)$ holds, and $\varepsilon; \varepsilon \vdash v : A$ implies $\varepsilon; \varepsilon \vdash \mu(x) : \Theta(x)$ for all $x \in \text{dom}(\Theta)$.

PROOF. By (mutual) induction on the derivation steps of evaluation. \square

The statements correspond to the three evaluation relations in $\lambda_{\Downarrow}^{\text{PI}}$. Note that the unidirectional evaluation targets expressions that are closed in terms of unidirectional variables, but may be open in terms of invertible variables, a property that is reflected in the first statement above. The second and third statements are more standard, assuming closed expressions in terms of both unidirectional and invertible variables. This assumption is actually an invariant; even though open expressions and values are involved in the unidirectional evaluation, the forward and backward evaluations always take and return closed values.

3.6.2 Bijectivity. Roughly speaking, correctness means that every value of type $A \rightleftharpoons B$ forms a bijection. Values of type $A \rightleftharpoons B$ has the form $\langle x.E \rangle$, and, By Lemma 3.2 and T-HOLED, values that occur in the evaluation of a well-typed term can be typed as $\varepsilon; \varepsilon \vdash \langle x.E \rangle : A \rightleftharpoons B$, which implies $\varepsilon; x : A \vdash E : B^{\bullet}$. Since values $\langle x.E \rangle$ can only be used by \triangleright and \triangleleft , bijectivity is represented as: $\{x \mapsto v\} \vdash E \Rightarrow v'$ if and only if $E \Leftarrow v' \dashv \{x \mapsto v\}$.¹⁰

¹⁰Here, we consider syntactic (definitional) equality of values, but it is rather easy to extend the discussion to observational equivalence.

To do so, we prove the following more general correspondence between the forward and backward evaluation relations, which is rather straightforward as the rules of the two evaluations are designed to be symmetric.

Lemma 3.3 (bijectivity of residuals). $\mu \vdash E \Rightarrow v$ if and only if $E \Leftarrow v \dashv \mu$.

PROOF. Each direction is proved by induction on a derivation of the corresponding evaluation. Note that every unidirectional evaluation judgment $e' \Downarrow v'$ occurring in a derivation of one direction also appears in the corresponding derivation of the other direction, and hence we can treat the unidirectional evaluation as a block box in this proof. \square

Then, by Lemma 3.2, we have the following corollary stating that $\langle x.E \rangle : A \rightleftharpoons B$ actually implements a bijection between A -typed values and B -typed values.

Corollary 3.4 (bijectivity of bijection values). Suppose $\varepsilon; \varepsilon \vdash \langle x.E \rangle : A \rightleftharpoons B$. Then, for any v and u such that $\varepsilon; \varepsilon \vdash v : A$ and $\varepsilon; \varepsilon \vdash v' : B$, we have $\{x \mapsto v\} \vdash E \Rightarrow v'$ if and only if $E \Leftarrow v' \dashv \{x \mapsto v\}$. \square

3.6.3 Note on the Progress Property. Progress is another important property that, together with subjection reduction, proves the absence of certain errors during evaluation. However, a standard progress property is usually based on small-step semantics, and yet $\lambda_{\Downarrow}^{\text{PI}}$ has a big-step operational semantics, which was chosen for its advantage in clarifying the input-output relationship of the forward and backward evaluation, as demonstrated by Lemma 3.3. A standard small-step semantics, which defines one-step evaluation as a relation between terms, is not suitable in this regard. Abstract machines are also unsatisfactory, as they will obscure the correspondence between the forward and backward evaluations.

We instead establish progress by directly showing that the evaluations do not get stuck other than with branching-related errors. This is done as an Agda implementation (mentioned in Sections 1 and 2.7) of a definitional [Reynolds 1998] interpreter, which uses the (sized) delay monad [Abel and Chapman 2014; Capretta 2005] and manipulates intrinsically-typed (i.e., Church style) expressions, values and residuals. The interpreter uses sums, products, and iso-recursive types instead of constructors. Also, instead of substitution, value environments are used in the unidirectional evaluation to avoid the shifting of de Bruijn terms. When the interpreter encounters imprecise **with** conditions, it goes to a infinite loop, and pattern match failures cannot happen as exhaustive patterns are used. This interpreter is typechecked in Agda, which serves as a constructive proof that there are no other kind of errors. We note that, as a bonus track, the Agda implementation comes with a formal proof of Lemma 3.3.

3.6.4 Reversible Turing Completeness. Reversible Turing completeness [Bennett 1973] is an important property that general-purpose reversible languages are expected to have. Similar to the standard Turing completeness, being reversible Turing complete for a language means that all bijections can be expressed in the language [Bennett 1973].

It is unsurprising that $\lambda_{\Downarrow}^{\text{PI}}$ is reversible Turing complete, as it has recursion (via fix_{π} in Section 3.3) and reversible branching (i.e., invertible **case**).

Theorem 3.5. $\lambda_{\Downarrow}^{\text{PI}}$ is reversible Turing complete. \square

The proof is done by constructing a simulator for a given reversible Turing machine, which is omitted due to space limitations. We follow the construction in Yokoyama et al. [2011] except the last step, in which we use a general reversible looping operator as below.¹¹

$$\text{trace} : ((a \oplus x)^{\bullet} \multimap (b \oplus x)^{\bullet}) \rightarrow a^{\bullet} \multimap b^{\bullet}$$

¹¹The operator is named after the trace operator [Joyal et al. 1996] in the category of bijections [Abramsky et al. 2002].

As its type suggests, *trace h* applies *h* to $\text{InL } a$ repeatedly until it returns $\text{InL } b$; the function loops while *h* returns a value of the form $\text{InR } x$. Intuitively, this behavior corresponds to the reversible loop [Lutz 1986]. In functional programming, loops are naturally encoded as tail recursions, which, however, are known to be difficult to handle in the contexts of program inversion [Glück and Kawabe 2004; Matsuda et al. 2010; Mogensen 2006; Nishida and Vidal 2011]. In fact, our implementation uses a non-trivial reversible programming technique, namely Yokoyama et al. [2012]’s optimized version of Bennett [1973]’s encoding. The higher-orderness of $\lambda_{\rightarrow}^{\text{PI}}$ (and SPARCL) helps here, as the effort is made once and for all.

3.7 Extension with The lift Operator

One feature we have not discussed is the **lift** operator that creates primitive bijections from unidirectional programs, for example, *sub* as we have seen in Section 2.

Adding **lift** to $\lambda_{\rightarrow}^{\text{PI}}$ is rather easy. We extend expressions to include **lift** as $e ::= \dots \mid \mathbf{lift} \ e_1 \ e_2 \ e_3$ together with the following typing rule.

$$\frac{\Gamma_1; \Theta_1 \vdash e_1 : A \rightarrow_{\omega} B \quad \Gamma_2; \Theta_2 \vdash e_2 : B \rightarrow_{\omega} A \quad \Gamma_3; \Theta_3 \vdash e_3 : A^{\bullet}}{\omega\Gamma_1 + \omega\Gamma_2 + \Gamma_3; \omega\Theta_1 + \omega\Theta_2 + \Theta_3 \vdash \mathbf{lift} \ e_1 \ e_2 \ e_3 : B^{\bullet}} \text{T-LIFT}$$

Accordingly, we extend evaluation by adding residuals of the form **lift** $(\lambda_{\omega}x_1.e_1) (\lambda_{\omega}x_2.e_2) E_3$ together with the following forward and backward evaluation rules (we omit the obvious unidirectional evaluation rule for obtaining residuals of this form).

$$\frac{\mu \vdash E_3 \Rightarrow v_3 \quad e_1[v_3/x_1] \Downarrow v}{\mu \vdash \mathbf{lift} (\lambda_{\omega}x_1.e_1) (\lambda_{\omega}x_2.e_2) E_3 \Rightarrow v} \quad \frac{e_2[v/x_2] \Downarrow v_3 \quad E_3 \Leftarrow v_3 \dashv \mu}{\mathbf{lift} (\lambda_{\omega}x_1.e_1) (\lambda_{\omega}x_2.e_2) E_3 \Leftarrow v \dashv \mu}$$

The substitution lemma (Lemma 3.1) and the subject reduction properties (Lemma 3.2) are also lifted to **lift**.

However **lift** is by nature unsafe, which requires an additional condition to ensure correctness. Specifically, the bijectivity of $A \rightleftharpoons B$ -typed values is only guaranteed if **lift** is used for pairs of functions that actually form bijections. For example, the uses of **lift** to construct *sub* in Section 2 are indeed safe. In Section 4.2.1, we will see another interesting example showing the use of conditionally safe **lifts** (see *unsafeNew* in Section 4.2.1).

4 LARGER EXAMPLES

In this section, we demonstrate the utility of SPARCL with examples, in which partial invertibility supported by SPARCL is the key for programming. The two examples are rebuilding trees from preorder and inorder traversals [Mu and Bird 2003], and a simplified version of the Huffman coding [Salomon 2008].

4.1 Rebuilding Trees from a Pre-Order and an In-Order Traversals

It is well-known that we can rebuild a node-labeled binary tree from its preorder and inorder traversals, provided that all labels in the tree are distinct. That is, for binary trees of type

```
data Tree = L | N Int Tree Tree
```

the following Haskell function *pi* is bijective.

```
pi :: Tree -> ([Int], [Int])
pi t = (preorder t, inorder t)
preorder L      = []
preorder (N a l r) = a : preorder l ++ preorder r
```

$$\begin{aligned} \text{inorder } L &= [] \\ \text{inorder } (N \ a \ l \ r) &= \text{inorder } l ++ [a] ++ \text{inorder } r \end{aligned}$$

For example, for binary trees $t_1 = N \ 1 \ (N \ 2 \ (N \ 3 \ L \ L) \ L) \ L$, $t_2 = N \ 1 \ (N \ 2 \ L \ (N \ 3 \ L \ L)) \ L$, $t_3 = N \ 1 \ (N \ 2 \ L \ L) \ (N \ 3 \ L \ L)$, $t_4 = N \ 1 \ L \ (N \ 2 \ (N \ 3 \ L \ L) \ L)$, and $t_5 = N \ 1 \ L \ (N \ 2 \ L \ (N \ 3 \ L \ L))$ that share the preorder traversal $[1, 2, 3]$, the inorder traversals distinguish them: $\text{inorder } t_1 = [3, 2, 1]$, $\text{inorder } t_2 = [2, 3, 1]$, $\text{inorder } t_3 = [2, 1, 3]$, $\text{inorder } t_4 = [1, 3, 2]$, and $\text{inorder } t_5 = [1, 2, 3]$.

The uniqueness of labels is key to the bijectivity of pi . It is clear that pi^{-1} returns L for $([], [])$, so the non-trivial part is how pi^{-1} will do for a pair of non-empty lists. Let us write $(a : p, i)$ for the pair. Then, since i contains exactly one a , we can unambiguously split i as $i = i_1 ++ [a] ++ i_2$. Then, by $pi^{-1}(\text{take } (\text{length } i_1) \ p, i_1)$, we can recover the left child l , and, by $pi^{-1}(\text{drop } (\text{length } i_1) \ p, i_2)$, we can recover the right child r . After that, from a , l , and r , we can construct the original input as $N \ a \ l \ r$. Notice that this inverse computation already involves partial invertibility such as the splitting of the inorder traversal list based on a .

It is straightforward to implement the above procedure in SPARCL. However, such a program is inefficient due to the cost of splitting. Program calculation is an established technique for deriving efficient programs through equational reasoning [Gibbons 2002], and in this case of tree-rebuilding, it is known that a linear-time inverse exists and can be derived [Mu and Bird 2003].

In the following, we demonstrate that program calculation works well in the setting of SPARCL. Interestingly, thinking in terms of partial-invertibility not only produces a Sparcl program, but actually improves the calculation by removing some of the more-obscure steps. Our calculation presented below basically follows Mu and Bird [2003, Section 3], although the presentation is a bit different as we focus on partial invertibility, especially the separation of unidirectional and invertible computation.

Note that recently, Glück and Yokoyama [2019] gives a reversible version of tree rebuilding using (an extension of) R-WHILE [Glück and Yokoyama 2016], a reversible imperative language inspired by Janus [Lutz 1986; Yokoyama et al. 2008]. However, R-WHILE only supports a very limited form of partial invertibility (Section 5.1), and the difference between their definition and ours is similar to what is demonstrated by the *goSubs* and *goSubsF* examples in Fig. 1.

4.1.1 Calculation of the Original Definition. The first step is tupling [Chin 1993; Hu et al. 1997] which eliminates multiple data traversals. The elimination of multiple data traversals is known to be useful for program inversion [Eppstein 1985; Matsuda et al. 2012].

$$\begin{aligned} pi &:: \text{Tree} \rightarrow ([\text{Int}], [\text{Int}]) \\ pi \ L &= ([], []) \\ pi \ (N \ a \ l \ r) &= \text{let } (pr, ir) = pi \ r; (pl, il) = pi \ l \ \text{in } (a : pl ++ pr, il ++ [a] ++ ir) \end{aligned}$$

Mu and Bird [2003, Section 3] also use tupling as the first step in their derivation.

The next step is to eliminate $++$, a source of inefficiency. The standard technique is to use accumulation parameters [Kühnemann et al. 2001]. Specifically, we obtain piA satisfying $piA \ t \ py \ iy = \text{let } (p, i) = pi \ t \ \text{in } (p ++ py, i ++ iy)$ as below.

$$\begin{aligned} piA &:: \text{Tree} \rightarrow [\text{Int}] \rightarrow [\text{Int}] \rightarrow ([\text{Int}], [\text{Int}]) \\ piA \ L \ \ py \ iy &= (py, iy) \\ piA \ (N \ a \ l \ r) \ py \ iy &= \text{let } (pr, ir) = piA \ r \ py \ iy; (pl, il) = piA \ l \ pr \ (a : ir) \ \text{in } (a : pl, il) \end{aligned}$$

The invertibility of piA is still not clear because piA is called with two different forms of the accumulation parameter iy : one is the case where iy is empty (e.g., the initial call $pi \ x = piA \ x \ [] \ []$), and the other is the case where it is not (e.g., the recursion for the left child $piA \ l \ pr \ (a : ir)$). This distinction between the two is important because, unlike the former, an inverse for the

$ \begin{aligned} & \mathit{piR} : \mathbf{Tree}^\bullet \multimap (\mathbf{List\ Int} \otimes \mathbf{List\ Int})^\bullet \\ & \mathit{piR\ L}^\bullet = (\mathbf{Nil}^\bullet, \mathbf{Nil}^\bullet)^\bullet \\ & \qquad \qquad \qquad \mathbf{with\ null} \circ \mathit{fst} \\ & \mathit{piR\ (N\ a\ l\ r)}^\bullet = \\ & \quad \mathbf{let\ (pr, ir)}^\bullet = \mathit{piR\ r\ in} \\ & \quad \mathbf{let\ (a, (pl, il))}^\bullet = \\ & \quad \quad \mathbf{pin\ a\ (\lambda a'. \mathit{piASR\ a'\ l\ pr\ ir})\ in} \\ & \quad (\mathbf{Cons}^\bullet\ a\ pl, il)^\bullet \end{aligned} $	$ \begin{aligned} & \mathit{piASR} : \mathbf{Int} \rightarrow \mathbf{Tree}^\bullet \multimap (\mathbf{List\ Int})^\bullet \multimap (\mathbf{List\ Int} \otimes \mathbf{List\ Int})^\bullet \\ & \mathit{piASR\ h\ L}^\bullet \quad \mathit{py\ iy} = (\mathit{py}, \mathbf{Cons}^\bullet\ (\mathit{new\ eqInt\ h})\ iy)^\bullet \\ & \qquad \qquad \qquad \mathbf{with\ eqInt\ h} \circ \mathbf{head} \circ \mathit{snd} \\ & \mathit{piASR\ h\ (N\ a\ l\ r)}^\bullet \mathit{py\ iy} = \\ & \quad \mathbf{let\ (pr, ir)}^\bullet = \mathit{piASR\ h\ r\ py\ iy\ in} \\ & \quad \mathbf{let\ (a, (pl, il))}^\bullet = \mathbf{pin\ a\ (\lambda a'. \mathit{piASR\ a'\ l\ pr\ ir})\ in} \\ & \quad (\mathbf{Cons}^\bullet\ a\ pl, il)^\bullet \end{aligned} $
--	--

Fig. 4. Invertible pre- and in-order traversal in SPARCL

latter is responsible for searching for the appropriate place to separate the in-order-traversal list. Nevertheless, this separation can be achieved by deriving a specialized version pi of piA satisfying $\mathit{pi\ x} = \mathit{piA\ x\ []\ []}$ (we reuse the name as it implements the same function).

$$\begin{aligned}
\mathit{pi} &:: \mathbf{Tree} \rightarrow ([\mathbf{Int}], [\mathbf{Int}]) \\
\mathit{pi\ L} &= ([], []) \\
\mathit{pi\ (N\ a\ l\ r)} &= \mathbf{let\ (pr, ir) = \mathit{pi\ r};\ (pl, il) = \mathit{piA\ l\ pr\ (a : ir)}\ in\ (a : pl, il)
\end{aligned}$$

Having this new version of pi , we now have an invariant that iy of $\mathit{piA\ t\ py\ iy}$ is always non-empty; the other case is separated into a call to pi . Moreover, we can determine the head h of iy beforehand in both forward and backward computations; this is exactly the label we search for to split the in-order-traversal list. Indeed, if we know the head h of iy beforehand, we can distinguish the ranges of the two branches of piA : for the first branch $(\mathit{py}, \mathit{iy})$, as iy is returned as is, the head of the second component is the same as h , and for the second branch $(a : \mathit{pl}, \mathit{il})$, the head of the second component of the return value cannot be equal to h , i.e., the head of iy . Recall that $\mathit{piA\ t\ py\ iy} = \mathbf{let\ (p, i) = \mathit{pi\ t\ in\ (p\ ++\ py, i\ ++\ iy)}$; thus, ir in the definition of piA must have the form of $\dots ++ \mathit{iy}$, and then il must have the form of $\dots ++ [a] ++ \dots ++ \mathit{iy}$.

Thus, as the last step of our calculation, we clarify the unidirectional part, namely the head of the second component of the accumulation parameters of piA , by changing it to a separate parameter. Specifically, we prepare the function piAS satisfying $\mathit{piAS\ h\ t\ py\ iy} = \mathit{piA\ t\ py\ (h : iy)}$ as below.

$$\begin{aligned}
\mathit{piAS} &:: \mathbf{Int} \rightarrow \mathbf{Tree} \rightarrow [\mathbf{Int}] \rightarrow [\mathbf{Int}] \rightarrow ([\mathbf{Int}], [\mathbf{Int}]) \\
\mathit{piAS\ h\ L} \quad \mathit{py\ iy} &= (\mathit{py}, h : \mathit{iy}) \\
\mathit{piAS\ h\ (N\ a\ l\ r)} \mathit{py\ iy} &= \mathbf{let\ (pr, ir) = \mathit{piAS\ h\ r\ py\ iy};\ (pl, il) = \mathit{piAS\ a\ l\ pr\ ir\ in\ (a : pl, il)}
\end{aligned}$$

Also, we replace the function call of piA in pi appropriately.

$$\begin{aligned}
\mathit{pi} &:: \mathbf{Tree} \rightarrow ([\mathbf{Int}], [\mathbf{Int}]) \\
\mathit{pi\ L} &= ([], []) \\
\mathit{pi\ (N\ a\ l\ r)} &= \mathbf{let\ (pr, ir) = \mathit{pi\ r};\ (pl, il) = \mathit{piAS\ a\ l\ pr\ ir\ in\ (a : pl, il)}
\end{aligned}$$

4.1.2 Making Partial-Invertibility Explicit. An efficient implementation in SPARCL falls out from the above calculation (see Fig. 4): the only additions are the types and the use of \mathbf{pin} . Recall that $\mathbf{let\ p}^\bullet = e_1 \mathbf{in\ } e_2$ is syntactic sugar for $\mathbf{case\ } e_1 \mathbf{of\ \{p}^\bullet \rightarrow e_2 \mathbf{with\ } \lambda_.. \mathbf{True}\}}$. Recall also that the first match principle is assumed and the catch-all \mathbf{with} conditions for the second branches are omitted. The function new in the program lifts an A -typed value a to an A^\bullet -typed value, corresponding to a bijection between $()$ and $\{a\}$.

$$\begin{aligned}
\mathit{new} &:: (a \rightarrow a \rightarrow \mathbf{Bool}) \rightarrow a \rightarrow a^\bullet \\
\mathit{new\ eq\ c} &= \mathbf{lift\ (\lambda_..c)\ (\lambda c'. \mathbf{case\ eq\ c\ c'\ of\ \{True \rightarrow ()\})\ ()}^\bullet
\end{aligned}$$

Note that the arguments of `lift` in `new eq` form a bijection, provided that `eq` implements the equality on A .

The backward evaluation of `piR` has the same behavior as that which [Mu and Bird \[2003, Section 3\]](#) derived. The partial bijection that `piASR` defines indeed corresponds to `reb` in their calculation. Their `reb` function is introduced as a rather magical step; our calculation can be seen as a justification of their choice.

4.1.3 new and delete. In the above example, we used `new`, which can be used to introduce redundancy to the output. For example, it is common to include checksum information in encoded data. The `new` function is effective for this scenario, as demonstrated below.

```
checkSum : List Int• → List Int•
checkSum xs = let (xs, s) = pin xs (λxs'.new eqInt (sum xs')) in
               Cons• s xs
```

In the forward direction, `checkSum` computes the sum of the list and prepends it to the list. In the backward direction, it checks if the head of the input list is the sum of its tail: if the check succeeds, the backward computation of `checkSum` returns the tail, and (correctly) fails otherwise.

It is worth mentioning that the pattern `new eq` is a finer operation than reversible copying where the inverse is given by equivalence checking [[Glück and Kawabe 2003](#)]; reversible copying can be implemented as $\lambda x. \text{pin } x \text{ (new eq)} : A^{\bullet} \rightarrow (A \otimes A)^{\bullet}$, assuming appropriate $eq : A \rightarrow A \rightarrow \text{Bool}$.

The `new` function has the corresponding inverse `delete`, which can be used to remove redundancy from the input.

```
delete : (a → a → Bool) → a → a• → ()•
delete eq c a = lift (λc'.case eq c c' of {True → ()}) (λ...c) a
```

It is interesting to note that `new` and `delete` can be used to define a safe variant of `lift`.

```
safeLift : (a → a → Bool) → (b → b → Bool) → (a → b) → (b → a) → a• → b•
safeLift eqA eqB f g a = let (a, b)• = pin a (λa'.new eqB (f a')) in
                          let (b, ())• = pin b (λb'.delete eqA (g b') a) in
                          b
```

In the forward computation, the function applies f to the input, and tests whether g is an inverse of f by applying g to the output and checking if the result is the same as the original input by `eqA`. The backward computation does the opposite: it applies g and tests the result by using f and `eqB`. This function is called “safe”, as it guarantees correctness by the runtime check, provided that `eqA` and `eqB` implement the equality on the domains.

4.2 Huffman Coding

The Huffman coding is one of the most popular compression algorithms [[Salomon 2008](#)]. The idea of the algorithm is to assign short code to frequently occurring symbols. For example, consider that we have symbols a, b, c and d that occur in the text to be encoded with probability 0.6, 0.2, 0.1, and 0.1 respectively. If we assign code as $a : 0, b : 10, c : 110$ and $d : 111$, then a text `aabacabdaa` will be encoded into 16-bit code $\underset{a}{0} \underset{a}{0} \underset{b}{10} \underset{0}{0} \underset{110}{110} \underset{0}{0} \underset{10}{10} \underset{111}{111} \underset{0}{0} \underset{0}{0}$, which is smaller than the 20-bit code obtained under the naive encoding that assigns two bits for each symbol.

4.2.1 Two-Pass Huffman Coding. Assume that we have a data structure for a Huffman coding table, represented by type `Huff`. The table may be represented as an array (or arrays) or a tree, and in practice one may want to use different data structures for encoding and decoding (for example, an array for encoding, and a trie for decoding). In this case, `Huff` is a pair of two data structures,

$ \begin{aligned} & \mathit{huffCompress} : (\text{List Symbol})^\bullet \multimap (\text{Huff} \otimes \text{List Bit})^\bullet \\ & \mathit{huffCompress} \ s = \\ & \quad \mathbf{let} \ (s, h)^\bullet = \mathbf{pin} \ s \ (\lambda s'. \mathit{new} \ \mathit{eqHuff} \ (\mathit{makeHuff} \ s')) \ \mathbf{in} \\ & \quad \mathbf{pin} \ h \ (\lambda h'. \mathit{encode} \ h' \ s) \end{aligned} $	$ \begin{aligned} & \mathit{encode} : \text{Huff} \rightarrow (\text{List Symbol})^\bullet \multimap (\text{List Bit})^\bullet \\ & \mathit{encode} \ h \ \text{Nil}^\bullet = \text{Nil}^\bullet \ \mathbf{with} \ \mathit{null} \\ & \mathit{encode} \ h \ (\text{Cons} \ s \ ss)^\bullet = \mathit{encR} \ h \ s \ (\mathit{encode} \ h \ ss) \end{aligned} $
--	--

Fig. 5. Two-pass Huffman coding in SPARCL

where each one is used only in one direction. To handle such a situation, we treat it as an abstract type with the following functions.

$$\begin{aligned}
\mathit{makeHuff} & : \text{List Symbol} \rightarrow \text{Huff} \\
\mathit{enc} & : \text{Huff} \rightarrow \text{Symbol} \rightarrow \text{List Bit} \\
\mathit{dec} & : \text{Huff} \rightarrow \text{List Bit} \rightarrow \text{Symbol} \otimes \text{List Bit}
\end{aligned}$$

Here, enc and dec satisfy the properties $\mathit{dec} \ h \ (\mathit{enc} \ h \ s \ ++ \ ys) = (s, ys)$ and $\mathit{dec} \ h \ ys = (s, ys')$ implies $\mathit{enc} \ s \ ++ \ ys = ys'$, where $++$ is the list append function.

Then, by enc and dec , we can define an bijective version encR as below.

$$\begin{aligned}
\mathit{encR} & : \text{Huff} \rightarrow \text{Symbol}^\bullet \multimap (\text{List Bit})^\bullet \multimap (\text{List Bit})^\bullet \\
\mathit{encR} \ h \ s \ r & = \mathbf{lift} \ (\lambda(s, ys). \mathit{enc} \ h \ s \ ++ \ ys) \ (\lambda ys. \mathit{dec} \ h \ ys) \ (s, r)^\bullet
\end{aligned}$$

An encoder can be defined by firstly constructing a Huffman coding table and then encoding symbol by symbol. We can program this procedure in a natural way in SPARCL (Fig. 5) by using \mathbf{pin} . This is an example where multiple \mathbf{pins} are used to convert data. The input symbol list is firstly passed to $\mathit{makeHuff}$ under new to create a Huffman table h in the first \mathbf{pin} ; here the input symbol list is unidirectional (static), while the constructed Huffman table is invertible. Then, the input symbol list is encoded with the constructed Huffman table in the second \mathbf{pin} ; here the input symbol list is invertible, while the Huffman table is unidirectional (static). A subtlety here is the use of $\mathit{eqHuff} : \text{Huff} \rightarrow \text{Huff} \rightarrow \text{Bool}$ to test the equality of the Huffman encoding tables. This check ensures the property that $\mathbf{fwd} \ \mathit{huffCompress} \ (\mathbf{bwd} \ \mathit{huffCompress} \ (h, ys)) = (h, ys)$. This equation holds only when h is the table obtained by applying $\mathit{makeHuff}$ to the decoded text; indeed, eqHuff checks the condition. One could avoid this check by using the following $\mathit{unsafeNew}$ instead.

$$\begin{aligned}
\mathit{unsafeNew} & : a \rightarrow a^\bullet \\
\mathit{unsafeNew} \ a & = \mathbf{lift} \ (\lambda().a) \ (\lambda a'.()) \quad \text{-- assuming } a = a'
\end{aligned}$$

The use of $\mathit{unsafeNew} \ a$ is safe only when its backward execution always receives a . Replacing new with $\mathit{unsafeNew}$ violates this assumption, but for this case, the replacement just widens the domain of $\mathbf{bwd} \ \mathit{huffCompress}$, which is acceptable even though $\mathbf{fwd} \ \mathit{huffCompress}$ and $\mathbf{bwd} \ \mathit{huffCompress}$ do not form a bijection due to $\mathit{unsafeNew}$. But in general this outcome is unreliable, unless the condition above can be guaranteed.

4.2.2 Concrete Representation of Huffman Tree in SPARCL. In the above we have modelled the case where different data structures are used for encoding and decoding, which demands the use of abstract type and consequently the use of $\mathbf{lifting}$. In this section, we define encR directly in SPARCL, which is possible when the same data structure is used for encoding and decoding.

To do so, we first give a concrete representation of Huff .

$$\mathbf{data} \ \text{Huff} = \text{Lf Symbol} \mid \text{Br Huff Huff}$$

Here, $\text{Lf} \ s$ encodes s into the empty sequence, and $\text{Br} \ l \ r$ encodes s into $\text{Cons} \ 0 \ c$ if l encodes s to c , and $\text{Cons} \ 1 \ c$ if r encodes s to c . For example, $\text{Br} \ (\text{Lf} \ 'a') \ (\text{Br} \ (\text{Lf} \ 'b') \ (\text{Br} \ (\text{Lf} \ 'c') \ (\text{Lf} \ 'd'))))$ is the Huffman tree used to encode the example presented in the beginning of Section 4.2.

$$\begin{array}{l|l}
\text{huffCompress} : (\text{List Symbol})^\bullet \multimap (\text{Huff} \otimes \text{List Bit})^\bullet & \text{encode} : \text{Huff} \rightarrow (\text{List Symbol})^\bullet \multimap (\text{List Bit})^\bullet \\
\text{huffCompress} = \text{encode initHuff} & \text{encode } h \text{ Nil}^\bullet = \text{Nil}^\bullet \quad \text{with null} \\
& \text{encode } h (\text{Cons } s \text{ ss})^\bullet = \\
& \quad \text{let } (s, r)^\bullet = \text{pin } s (\lambda s'. \text{encode } (\text{updHuff } s' h) \text{ ss}) \text{ in} \\
& \quad \text{encR } h \text{ s } r
\end{array}$$

Fig. 6. Adaptive Huffman coding in SPARCL

Now let us define *encR* to be used in *encode* above. It is easier to define it via its inverse *decR*.

$$\begin{array}{l}
\text{decR} : \text{Huff} \rightarrow (\text{List Bit})^\bullet \multimap (\text{Symbol} \otimes \text{List Bit})^\bullet \\
\text{decR } (\text{Lf } s) \text{ ys} = (\text{new eqSym } s, \text{ys})^\bullet \\
\text{decR } (\text{Br } l \text{ r}) \text{ ys} = \text{case } \text{ys} \text{ of } (\text{Cons } 0 \text{ ys}')^\bullet \rightarrow \text{decR } l \text{ ys}' \text{ with } \lambda(s, _). \text{member } s \text{ l} \\
\quad (\text{Cons } 1 \text{ ys}')^\bullet \rightarrow \text{decR } r \text{ ys}' \\
\text{encR } h \text{ s } \text{ys} = \text{invert } (\text{decR } h) (s, \text{ys})^\bullet
\end{array}$$

Here, *member* : *Symbol* \rightarrow *Huff* \rightarrow *Bool* is a membership test function. Recall that *invert* implements inversion of a bijection (Section 2). One can find that searching *s* in *l* for every recursive call is inefficient, and this cost can be avoided by additional information on *Br* that makes a Huffman tree a search tree. Another solution is to use different data structures for encoding and decoding as we demonstrated in Section 4.2.1.

4.2.3 Adaptive Huffman Coding. In the above *huffCompress*, a Huffman coding table is fixed during compression which requires the preprocessing *makeHuff* to compute the table. This is sometimes suboptimal: for example, a one-pass method is preferred for streaming while a text could consist of several parts with very different frequency distributions of symbols.

Being adaptive means that we have the following two functions instead of *makeHuff*.

$$\text{initHuff} : \text{Huff} \quad \text{updHuff} : \text{Symbol} \rightarrow \text{Huff} \rightarrow \text{Huff}$$

Instead of constructing a Huffman coding table beforehand, the Huffman coding table is constructed and changed throughout compression here.

The updating process of the Huffman coding table is the same in both compression and decompression, which means that SPARCL is effective for writing an invertible and adaptive version of Huffman coding in a natural way (Fig. 6). This is another demonstration of the SPARCL's strength in partial invertibility. Programming the same bijection in a fully-invertible language gets a lot more complicated due to the irreversible nature of *updHuff*.

5 RELATED WORK

5.1 Program Inversion and Invertible/Reversible Computation

In the literature of program inversion (a program transformation technique to find f^{-1} for a given f), it is known that an inverse of a function may not arise from reversing all the execution steps of the original program. Partial inversion [Nishida et al. 2005] addresses the problem by classifying inputs/outputs into *known* and *unknown*, where known information is available also for inverses. This classification can be viewed as a binding-time analysis [Gomard and Jones 1991; Jones et al. 1993] where the known part is treated as static. The partial inversion is further extended so that the return values of inverses are treated as known as well [Almendros-Jiménez and Vidal 2006]; in this case, it can no longer be explained as a binding-time analysis. This extension introduces additional power, but makes inversion fragile as success depends on which function is inverted first. For example, the partial inversion for *goSubs* succeeds when it inverts $x - n$

first, but fails if it tried to invert *goSubs* x x s first. The design of SPARCL is inspired by these partial inversion methods: we use $(-)^{\bullet}$ -types to distinguish the known and unknown parts, and **pin** together with **case** to control orders. Semi inversion [Mogensen 2005] essentially converts a program to logic programs, which also allows the original and inverse programs to have common computations. Its extension [Mogensen 2008] can handle a limited form of function arguments. The PINS system allows users to specify control structures as they sometimes differ from the original program [Srivastava et al. 2011]. As we mentioned in Section 1, these program inversion methods may fail, and often for reasons that are not obvious to programmers.

Embedded languages can be seen as two-staged (a host and a guest), and there are several embedded invertible/reversible programming languages. A popular approach to implement such languages is based on combinators [Kennedy and Vytiniotis 2012; Mu et al. 2004b; Rendel and Ostermann 2010; Wang et al. 2013], in which users program by composing bijections through designated combinators. To the best of our knowledge, only [Kennedy and Vytiniotis 2012] has an operator like $\mathbf{pin} : A^{\bullet} \multimap (A \rightarrow B^{\bullet}) \multimap A \otimes B^{\bullet}$, which is key to partial invertibility. More specifically, Kennedy and Vytiniotis [2012] has an operator $\mathit{depGame} :: \mathit{Game} \ a \rightarrow (a \rightarrow \mathit{Game} \ b) \rightarrow \mathit{Game} \ (a, b)$. The types suggest that *Game* and $(-)^{\bullet}$ play a similar role; indeed they both represent invertibility but in different ways. In their system, *Game* a represents (total) bijections from bit sequences and a -typed values, while in our system A^{\bullet} represents a bijection whose range is A but domain is determined when **unlift** is applied. One consequence of this difference is that, in their domain-specific system, there is no restriction of using a value $v :: \mathit{Game} \ a$ linearly, because there is no problem of using an encoder/decoder pair for type a multiple times, even though nonlinear use of $v : A^{\bullet}$, especially discarding, leads to non-bijection. Another consequence of the difference is that their system is hardwired to bit sequences and therefore does not support deriving general bijections between a and b from $\mathit{Game} \ a \rightarrow \mathit{Game} \ b$, whereas we can obtain a (not-necessarily-total) bijections between A and B from any function of type $A^{\bullet} \multimap B^{\bullet}$ that does not contain linear free variables.

The **pin** operator can be seen as a functional generalization of reversible update statements [Axelsen et al. 2007] $x \oplus = e$ in reversible imperative languages [Frank 1997; Glück and Yokoyama 2016; Lutz 1986; Yokoyama et al. 2008], of which the inverse is given by $x \ominus = e$ with \ominus satisfying $(x \oplus y) \ominus y = x$ for any y ; examples of \oplus (and \ominus) include addition, subtraction, bitwise XOR, and replacement of `nil` [Glück and Yokoyama 2016] as a form of reversible copying [Glück and Kawabe 2003]. Having $(x \oplus y) \ominus y$ means that \oplus and \ominus are partially invertible, and indicates that they correspond to the second argument of **pin**. Whereas the operators such as \oplus and \ominus are fixed in those languages, in SPARCL, leveraging its higher-orderness, any function of an appropriate type can be used as the second argument of **pin**, which leads to concise function definitions as demonstrated in *goSub* in Section 2 and the examples in Section 4.

Most of the existing reversible programming languages [Baker 1992; Frank 1997; Lutz 1986; Mu et al. 2004b; Wang et al. 2013; Yokoyama et al. 2008, 2011] do not support function values, and higher-order reversible programming languages are uncommon. One notable exception is Abramsky [2005] that shows a subset of the linear λ -calculus concerning \multimap and $!$ (more precisely, a combinator logic that corresponds to the subset) can be interpreted as manipulations of (not-necessarily-total) bijections. However, it is known to be difficult to extend their system to primitives such as constructors and invertible pattern matching [Abramsky 2005, Section 7].

A few reversible functional programming languages also support a limited form of partial invertibility. RFunT,¹² a typed variant of RFun [Yokoyama et al. 2011] with Haskell-like syntax, allows a function to take additional parameters called ancilla parameters. The reversibility restriction

¹²<https://github.com/kirkedal/rfun-interp>

is relaxed for ancilla parameters, and they can be discarded and pattern-matched without requiring a way to determine branching from their results. However, these ancilla parameters are supposed to be translated into auxiliary inputs and outputs that stay the same before and after reversible computation, and mixing unidirectional computation is not their primary purpose. In fact, very limited operations are allowed for these ancilla data by the system. CoreFun also supports ancilla parameters [Jacobsen et al. 2018]. Their ancilla parameters are treated as static inputs to reversible functions, and arguments that appear at ancilla positions are free from the linearity restriction.¹³ The system is overly conservative: all the functions are (partially) reversible, and thus functions themselves used in the ancilla positions must obey the linearity restriction. More crucially, both RFunT and CoreFun are first-order languages (to be precise, they allow top-level function names to be used as values, but not partial application or λ -abstraction), which limits flexible programming. In contrast, A^\bullet is an ordinary type in SPARCL, and there is no syntactic restriction on expressions of type A^\bullet . This feature, combined with the higher-orderness, gives extra flexibility in mixing unidirectional and invertible programming. For example, SPARCL allows a function composition operator that can be used for both unidirectional (hence unrestricted) and invertible (hence linear) functions, using multiplicity polymorphism [Bernardy et al. 2018; Matsuda 2020].

5.2 Functional Quantum Programming Languages

In quantum programming many operation are reversible, and there are a few higher-order quantum programming languages [Rios and Selinger 2017; Selinger and Valiron 2006]. Among them, the type system of Proto-Quipper-M [Rios and Selinger 2017] is similar to $\lambda_{\perp}^{\text{PI}}$ in the sense that it also uses a linear-type system and distinguishes two sorts of variable environments as we do with Γ and Θ , although the semantic back-ends are different. They do not have any language construct that introduces new variables to the second sort of environments (a counterpart of our Θ), because their language does not have a counterpart to our invertible case.

It is also interesting to see that some quantum languages allow weakening (i.e., discarding) [Selinger and Valiron 2006] and some allow contraction (i.e., copying) [Altenkirch and Grattage 2005]. In these frameworks, weakening is allowed because one can throw away a quantum bit after measuring, and contraction is allowed because states can be shared through introducing entanglements. As our goal is to obtain a bijection as final product, weakening in general is not possible in our context. On the other hand, it is a design choice whether or not contraction is allowed. Since the inverse of copying can be given by equivalence checking and vice versa [Glück and Kawabe 2003]. However, careless uses of copying may result in unintended domain restriction. Moreover supporting such a feature requires hard-wired equivalence checks for all types of variables that can be in Θ (notice that multiple uses of a variable in Γ will be reduced to multiple uses of variables in Θ [Matsuda and Wang 2018c]). This requires the type system to distinguish types that can be in Θ from general ones, as types such as $A \multimap B$ do not have decidable equality. Moreover, the hard-wired equivalence checks would prevent users from using abstract types such as Huff in Section 4, for which the definition of equivalence can differ from that on their concrete representations.

5.3 Bidirectional Programming Languages

It is perhaps not surprising that many of the concerns in designing invertible/bijective/reversible languages are shared by the closely related field of bidirectional programming [Foster et al. 2007]. A bidirectional transformation is a generalization of a pair of inverses that allows a component to be non-bijective; for example, an (asymmetric) bidirectional transformation between a and b are given by two functions called $get : a \rightarrow b$ and $put : a \rightarrow b \rightarrow a$ [Foster et al. 2007]. Similarly to ours, in the

¹³A correction to Jacobsen et al. [2018] (personal communication with Michael-Kirkedal Thomsen, Jun 2020).

bidirectional language HOBiT [Matsuda and Wang 2018c], a bidirectional transformation between a and b is represented by a function from $\mathbf{B} a$ to $\mathbf{B} b$, and top-level functions of type $\mathbf{B} a \rightarrow \mathbf{B} b$ can be converted to a bidirectional transformation between a and b . Despite the similarity, there are unique challenges in invertible programming. Notably, the handling of partial-invertibility that this paper focuses on, and the introduction of the operator `pin` as a solution. Another difference is that SPARCL is based on a linear type system, which, as we have seen, perfectly supports the need for the intricate connections between unidirectional and inverse computation in addressing partial invertibility. One of the consequences of this difference in the underlying type system is that Matsuda and Wang [2018c] can only interpret *top-level* functions of type $\mathbf{B} a \rightarrow \mathbf{B} b$ as bidirectional transformations between a and b , yet we can interpret functions of type $A^\bullet \rightarrow B^\bullet$ *in any places* as bijections between A and B , as long as they have no linear free variables. Linear types also clarify the roles of values and prevent users from unintended failures caused by erroneous use of variables. For example, the type $A^\bullet \rightarrow (A \rightarrow B^\bullet) \rightarrow (A \otimes B)^\bullet$ of `pin` clarifies that the function argument of `pin` can safely discard or copy its input as the nonlinear uses do not affect the domain of the resulting bijection.

It is worth mentioning that, in addition to bidirectional transformations, HOBiT provides a way to lift bidirectional combinators (i.e., functions that take and return bidirectional transformations). However, the same is not obvious in SPARCL due to its linear type system, as the combinators need to take care of the manipulation of Θ environments such as splitting $\Theta = \Theta_1 + \Theta_2$. On the other hand, there is less motivation to lift combinators in the context of bijective/reversible programming especially for languages that are expressive enough to be reversible Turing complete [Bennett 1973].

The applicative-lens framework [Matsuda and Wang 2015a, 2018a], which is an embedded domain specific language in Haskell, provides a function *lift* that converts a bidirectional transformation ($a \rightarrow b, a \rightarrow b \rightarrow a$) to a function of type $L s a \rightarrow L s b$ where L is an abstract type parameterized by s . As in HOBiT, bidirectional transformations are represented as functions so that they can be composed by unidirectional functions; the name *applicative* in fact comes from the applicative (point-wise functional) programming style. (To be precise, L together with certain operations forms a lax monoidal functor [Mac Lane 1998, Section XI.2] as Applicative instances [McBride and Paterson 2008; Paterson 2012] but not endo to be an Applicative instance [Matsuda and Wang 2018a].) The type parameter s has a similar role to the s of the $ST s$ monad [Launchbury and Jones 1994], which enables the *unlifting* that converts a polymorphic function $\forall s. L s a \rightarrow L s b$ back to a bidirectional transformation ($a \rightarrow b, a \rightarrow b \rightarrow a$). That is, unlike HOBiT, functions that will be interpreted as bidirectional transformations are not limited to top-level ones. However, in exchange for this utility, the expressive power of the applicative lens is limited compared with HOBiT; for example, bidirectional cases are not supported in the framework, and resulting bidirectional transformations cannot propagate structural updates as a result.

As a remark, duplication (contraction) of values is also a known challenge in bidirectional transformation, for the purpose of supporting multiple views of the same data and synchronization among them [Hu et al. 2004]. However, having unrestricted duplication makes compositional reasoning of correctness very difficult; in fact most of the fundamental properties of bidirectional transformation, including well-behavedness [Foster et al. 2007] and its weaker variants [Hidaka et al. 2010; Mu et al. 2004a], are not preserved in the presence of unrestricted duplication [Matsuda and Wang 2015b].

5.4 Linear Type Systems

SPARCL is based on λ_{\perp}^q , a core system of Linear Haskell [Bernardy et al. 2018], with qualified typing [Jones 1995; Vytiniotis et al. 2011] for effective inference [Matsuda 2020]. An advantage of

this system is that the only place where we need to explicitly handle linearity is the manipulation of $(-)^{\bullet}$ -typed values; there is no need of any special annotations for the unidirectional parts, as demonstrated in the examples. This is different from Wadler [1993]’s linear type system, which would require a lot of ! annotations in the code. Linear Haskell is not the only approach that is able to avoid the scattering of !s. Mazurak et al. [2010] use kinds (\circ and $*$) to distinguish types that are treated in a linear way (\circ) from those that are not ($*$). Thanks to the subkinding $* \leq \circ$, no syntactic annotations are required to convert the unrestricted values to linear ones. Their system has two sort of function types: $\overset{\circ}{\rightarrow}$ for the functions that themselves are treated in the linear way and $\overset{*}{\rightarrow}$ for the functions that are unrestricted. As a result, a function can have multiple incomparable types; e.g., the K combinator can have four types [Morris 2016]. Universal types accompanied by kind abstraction [Tov and Pucella 2011] addresses the issue to some extent; it works well especially for K , but still gives the B combinator two incomparable types [Morris 2016]. Morris [2016] further extends these two systems to overcome the issue by using qualified types [Jones 1995], which can infer principal types thank to inequality constraints. Note that the implementation of SPARCL uses an inference system by Matsuda [2020], which, based on OUTSIDEIN(X) [Vytiniotis et al. 2011], also uses qualified typing with inequality constraints for $\lambda^{\underline{q}}$, inspired by Morris [2016].

6 CONCLUSION

We have designed SPARCL, a language for partially-invertible computation. The key idea of SPARCL is to use types to distinguish data that are subject to invertible computation and those that are not; specifically the type constructor $(-)^{\bullet}$ is used for marking the former. A linear type system is utilized for connecting the two worlds. We have presented the syntax, type system and semantics of SPARCL, and proved that invertible computations defined in SPARCL are in fact invertible (and hence bijective). To demonstrate the utility of our proposed language, we have proved its reversible Turing completeness, and presented non-trivial examples of tree rebuilding and Huffman coding.

There are several future directions of this research. One direction is to use finer type systems. Recall that we need to check **with** conditions even in the forward computation, which can be costly. We believe that refinement types and their inference [Rondon et al. 2008; Xi and Pfenning 1998] would be useful for addressing this issue. Currently, our prototype implementation is standalone, preventing users from writing functions in another language to be used in **lift**, and from using functions obtain by **fwd** and **bwd** in the other language. Although prototypical implementation of a compiler of SPARCL to Haskell is in progress, a seamless integration through an embedded implementation would be desirable [Matsuda and Wang 2018b]. Another direction is to extend our approach to bidirectional transformations [Foster et al. 2007] to create the notion of partially bidirectional programming. As discussed in Section 5, handling copying (i.e., contraction) is an important issue; we want to find the sweet spot of allowing flexible copying without compromising reasoning about correctness.

ACKNOWLEDGMENTS

We thank the IFIP 2.1 members for their critical but constructive comments on a preliminary version of this research, and Samantha Frohlich for her helpful suggestions and comments on the presentation of this paper. We also thank the anonymous reviewers of ICFP 2020 for their constructive comments. This work was partially supported by JSPS KAKENHI Grant Numbers 15H02681, 19K11892 and 20H04161, JSPS Bilateral Program, Grant Number JPJSBP120199913, the Kayamori Foundation of Informational Science Advancement, EPSRC Grant *EXHIBIT: Expressive High-Level Languages for Bidirectional Transformations* (EP/T008911/1), and Royal Society Grant *Bidirectional Compiler for Software Evolution* (IES\R3\170104).

REFERENCES

- Andreas Abel and James Chapman. 2014. Normalization by Evaluation in the Delay Monad: A Case Study for Coinduction via Copatterns and Sized Types. In *Proceedings 5th Workshop on Mathematically Structured Functional Programming, MSFP@ETAPS 2014, Grenoble, France, 12 April 2014 (EPTCS)*, Paul Levy and Neel Krishnaswami (Eds.), Vol. 153. 51–67. <https://doi.org/10.4204/EPTCS.153.4>
- Sergei M. Abramov, Robert Glück, and Yuri A. Klimov. 2006. An Universal Resolving Algorithm for Inverse Computation of Lazy Languages. In *Ershov Memorial Conference (Lecture Notes in Computer Science)*, Irina Virbitskaite and Andrei Voronkov (Eds.), Vol. 4378. Springer, 27–40.
- Samson Abramsky. 2005. A structural approach to reversible computation. *Theor. Comput. Sci.* 347, 3 (2005), 441–464. <https://doi.org/10.1016/j.tcs.2005.07.002>
- Samson Abramsky, Esfandiar Haghverdi, and Philip J. Scott. 2002. Geometry of Interaction and Linear Combinatory Algebras. *Mathematical Structures in Computer Science* 12, 5 (2002), 625–665. <https://doi.org/10.1017/S0960129502003730>
- Jesús Manuel Almendros-Jiménez and Germán Vidal. 2006. Automatic Partial Inversion of Inductively Sequential Functions. In *Implementation and Application of Functional Languages, 18th International Symposium, IFL 2006, Budapest, Hungary, September 4-6, 2006, Revised Selected Papers (Lecture Notes in Computer Science)*, Zoltán Horváth, Viktória Zsók, and Andrew Butterfield (Eds.), Vol. 4449. Springer, 253–270. https://doi.org/10.1007/978-3-540-74130-5_15
- Thorsten Altenkirch, James Chapman, and Tarmo Uustalu. 2010. Monads Need Not Be Endofunctors. In *Foundations of Software Science and Computational Structures, 13th International Conference, FOSSACS 2010, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2010, Paphos, Cyprus, March 20-28, 2010. Proceedings (Lecture Notes in Computer Science)*, C.-H. Luke Ong (Ed.), Vol. 6014. Springer, 297–311. https://doi.org/10.1007/978-3-642-12032-9_21
- Thorsten Altenkirch and Jonathan Grattage. 2005. A Functional Quantum Programming Language. In *20th IEEE Symposium on Logic in Computer Science (LICS 2005), 26-29 June 2005, Chicago, IL, USA, Proceedings*. IEEE Computer Society, 249–258. <https://doi.org/10.1109/LICS.2005.1>
- Sergio Antoy, Rachid Echahed, and Michael Hanus. 2000. A needed narrowing strategy. *J. ACM* 47, 4 (2000), 776–822. <https://doi.org/10.1145/347476.347484>
- Holger Bock Axelsen, Robert Glück, and Tetsuo Yokoyama. 2007. Reversible Machine Code and Its Abstract Processor Architecture. In *Computer Science - Theory and Applications, Second International Symposium on Computer Science in Russia, CSR 2007, Ekaterinburg, Russia, September 3-7, 2007, Proceedings (Lecture Notes in Computer Science)*, Volker Diekert, Mikhail V. Volkov, and Andrei Voronkov (Eds.), Vol. 4649. Springer, 56–69. https://doi.org/10.1007/978-3-540-74510-5_9
- Henry G. Baker. 1992. NREVERSAL of Fortune - The Thermodynamics of Garbage Collection. In *Memory Management, International Workshop IWMM 92, St. Malo, France, September 17-19, 1992, Proceedings (Lecture Notes in Computer Science)*, Yves Bekkers and Jacques Cohen (Eds.), Vol. 637. Springer, 507–524. <https://doi.org/10.1007/BFb0017210>
- Charles H. Bennett. 1973. Logical Reversibility of Computation. *IBM Journal of Research and Development* 17, 6 (Nov 1973), 525–532. <https://doi.org/10.1147/rd.176.0525>
- Jean-Philippe Bernardy, Mathieu Boespflug, Ryan R. Newton, Simon Peyton Jones, and Arnaud Spiwack. 2018. Linear Haskell: practical linearity in a higher-order polymorphic language. *PACMPL* 2, POPL (2018), 5:1–5:29. <https://doi.org/10.1145/3158093>
- Venanzio Capretta. 2005. General recursion via coinductive types. *Logical Methods in Computer Science* 1, 2 (2005). [https://doi.org/10.2168/LMCS-1\(2:1\)2005](https://doi.org/10.2168/LMCS-1(2:1)2005)
- Wei-Ngan Chin. 1993. Towards an Automated Tupling Strategy. In *Proceedings of the ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation, PEPM'93, Copenhagen, Denmark, June 14-16, 1993*, David A. Schmidt (Ed.), ACM, 119–132. <https://doi.org/10.1145/154630.154643>
- Rowan Davies and Frank Pfenning. 2001. A modal analysis of staged computation. *J. ACM* 48, 3 (2001), 555–604. <https://doi.org/10.1145/382780.382785>
- David Eppstein. 1985. A Heuristic Approach to Program Inversion. In *IJCAI*. 219–221.
- J. Nathan Foster, Michael B. Greenwald, Jonathan T. Moore, Benjamin C. Pierce, and Alan Schmitt. 2007. Combinators for bidirectional tree transformations: A linguistic approach to the view-update problem. *ACM Trans. Program. Lang. Syst.* 29, 3 (2007).
- Michael P. Frank. 1997. The R Programming Language and Compiler. MIT Reversible Computing Project Memo #M8, MIT AI Lab. Available on: <https://github.com/mikepfrank/Rlang-compiler/blob/master/docs/MIT-RCP-MemoM8-RProgLang.pdf>.
- Jeremy Gibbons. 2002. *Calculating Functional Programs*. Springer Berlin Heidelberg, Berlin, Heidelberg, 151–203. https://doi.org/10.1007/3-540-47797-7_5
- Robert Glück and Masahiko Kawabe. 2003. A Program Inverter for a Functional Language with Equality and Constructors. In *APLAS (Lecture Notes in Computer Science)*, Atsushi Ohori (Ed.), Vol. 2895. Springer, 246–264. https://doi.org/10.1007/978-3-540-40018-9_17

- Robert Glück and Masahiko Kawabe. 2004. Derivation of Deterministic Inverse Programs Based on LR Parsing. In *FLOPS (Lecture Notes in Computer Science)*, Yuki Yoshi Kameyama and Peter J. Stuckey (Eds.), Vol. 2998. Springer, 291–306.
- Robert Glück and Tetsuo Yokoyama. 2016. A Linear-Time Self-Interpreter of a Reversible Imperative Language. *Computer Software* 33, 3 (2016), 3_108–3_128. https://doi.org/10.11309/jssst.33.3_108
- Robert Glück and Tetsuo Yokoyama. 2019. Constructing a binary tree from its traversals by reversible recursion and iteration. *Inf. Process. Lett.* 147 (2019), 32–37. <https://doi.org/10.1016/j.ipl.2019.03.002>
- Carsten K. Gomard and Neil D. Jones. 1991. A Partial Evaluator for the Untyped lambda-Calculus. *J. Funct. Program.* 1, 1 (1991), 21–69. <https://doi.org/10.1017/S0956796800000058>
- Soichiro Hidaka, Zhenjiang Hu, Kazuhiro Inaba, Hiroyuki Kato, Kazutaka Matsuda, and Keisuke Nakano. 2010. Bidirectionalizing graph transformations. In *Proceeding of the 15th ACM SIGPLAN international conference on Functional programming, ICFP 2010, Baltimore, Maryland, USA, September 27-29, 2010*, Paul Hudak and Stephanie Weirich (Eds.). ACM, 205–216. <https://doi.org/10.1145/1863543.1863573>
- Zhenjiang Hu, Hideya Iwasaki, Masato Takeichi, and Akihiko Takano. 1997. Tupling Calculation Eliminates Multiple Data Traversals. In *Proceedings of the 1997 ACM SIGPLAN International Conference on Functional Programming (ICFP '97), Amsterdam, The Netherlands, June 9-11, 1997.*, Simon L. Peyton Jones, Mads Tofte, and A. Michael Berman (Eds.). ACM, 164–175. <https://doi.org/10.1145/258948.258964>
- Zhenjiang Hu, Shin-Cheng Mu, and Masato Takeichi. 2004. A programmable editor for developing structured documents based on bidirectional transformations. In *Proceedings of the 2004 ACM SIGPLAN Workshop on Partial Evaluation and Semantics-based Program Manipulation, 2004, Verona, Italy, August 24-25, 2004*, Nevin Heintze and Peter Sestoft (Eds.). ACM, 178–189. <https://doi.org/10.1145/1014007.1014025>
- Petur Andrias Højgaard Jacobsen, Robin Kaarsgaard, and Michael Kirkedal Thomsen. 2018. \mathsf{CoreFun} : A Typed Functional Reversible Core Language. In *Reversible Computation - 10th International Conference, RC 2018, Leicester, UK, September 12-14, 2018, Proceedings (Lecture Notes in Computer Science)*, Jarkko Kari and Irek Ulidowski (Eds.), Vol. 11106. Springer, 304–321. https://doi.org/10.1007/978-3-319-99498-7_21
- Roshan P. James and Amr Sabry. 2012. Information effects. In *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012, Philadelphia, Pennsylvania, USA, January 22-28, 2012*, John Field and Michael Hicks (Eds.). ACM, 73–84. <https://doi.org/10.1145/2103656.2103667>
- Mark P. Jones. 1995. *Qualified Types: Theory and Practice*. Cambridge University Press, New York, NY, USA.
- Neil D. Jones, Carsten K. Gomard, and Peter Sestoft. 1993. *Partial evaluation and automatic program generation*. Prentice Hall.
- André Joyal, Ross Street, and Dominic Verity. 1996. Traced Monoidal Categories. *Mathematical Proceedings of the Cambridge Philosophical Society* 119, 3 (Apr 1996), 447–468.
- Andrew J. Kennedy and Dimitrios Vytiniotis. 2012. Every bit counts: The binary representation of typed data and programs. *J. Funct. Program.* 22, 4-5 (2012), 529–573.
- Armin Kühnemann, Robert Glück, and Kazuhiko Kakehi. 2001. Relating Accumulative and Non-accumulative Functional Programs. In *RTA (Lecture Notes in Computer Science)*, Aart Middeldorp (Ed.), Vol. 2051. Springer, 154–168.
- Rolf Landauer. 1961. Irreversibility and Heat Generation in the Computing Process. *IBM Journal of Research and Development* 5, 3 (1961), 183–191. <https://doi.org/10.1147/rd.53.0183>
- John Launchbury and Simon L. Peyton Jones. 1994. Lazy Functional State Threads. In *Proceedings of the ACM SIGPLAN'94 Conference on Programming Language Design and Implementation (PLDI), Orlando, Florida, USA, June 20-24, 1994*, Vivek Sarkar, Barbara G. Ryder, and Mary Lou Soffa (Eds.). ACM, 24–35. <https://doi.org/10.1145/178243.178246>
- Christopher Lutz. 1986. Janus: a time-reversible language. (1986). *Letter to R. Landauer*. Available on: <http://tetsuo.jp/ref/janus.pdf>.
- Saunders Mac Lane. 1998. *Categories for the Working Mathematician* (second edition ed.). Graduate Texts in Mathematics, Vol. 5. Springer.
- Kazutaka Matsuda. 2020. Modular Inference of Linear Types for Multiplicity-Annotated Arrows. In *Programming Languages and Systems - 29th European Symposium on Programming, ESOP 2020, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, 2020, Proceedings (Lecture Notes in Computer Science)*, Peter Müller (Ed.), Vol. 12075. Springer, 456–483. https://doi.org/10.1007/978-3-030-44914-8_17 The full version is available on: <http://arxiv.org/abs/1911.00268v2>.
- Kazutaka Matsuda, Zhenjiang Hu, Keisuke Nakano, Makoto Hamana, and Masato Takeichi. 2007. Bidirectionalization transformation based on automatic derivation of view complement functions. In *Proceedings of the 12th ACM SIGPLAN International Conference on Functional Programming, ICFP 2007, Freiburg, Germany, October 1-3, 2007*, Ralf Hinze and Norman Ramsey (Eds.). ACM, 47–58. <https://doi.org/10.1145/1291151.1291162>
- Kazutaka Matsuda, Kazuhiro Inaba, and Keisuke Nakano. 2012. Polynomial-time inverse computation for accumulative functions with multiple data traversals. *Higher-Order and Symbolic Computation* 25, 1 (2012), 3–38. <https://doi.org/10.1007/s10990-013-9097-8>

- Kazutaka Matsuda, Shin-Cheng Mu, Zhenjiang Hu, and Masato Takeichi. 2010. A Grammar-Based Approach to Invertible Programs. In *ESOP (Lecture Notes in Computer Science)*, Andrew D. Gordon (Ed.), Vol. 6012. Springer, 448–467.
- Kazutaka Matsuda and Meng Wang. 2013. FliPpr: A Prettier Invertible Printing System. In *ESOP (Lecture Notes in Computer Science)*, Matthias Felleisen and Philippa Gardner (Eds.), Vol. 7792. Springer, 101–120. https://doi.org/10.1007/978-3-642-37036-6_6
- Kazutaka Matsuda and Meng Wang. 2015a. Applicative bidirectional programming with lenses. In *ICFP*, Kathleen Fisher and John H. Reppy (Eds.). ACM, 62–74. <https://doi.org/10.1145/2784731.2784750>
- Kazutaka Matsuda and Meng Wang. 2015b. "Bidirectionalization for free" for monomorphic transformations. *Sci. Comput. Program.* 111 (2015), 79–109. <https://doi.org/10.1016/j.scico.2014.07.008>
- Kazutaka Matsuda and Meng Wang. 2018a. Applicative bidirectional programming: Mixing lenses and semantic bidirectionalization. *J. Funct. Program.* 28 (2018), e15. <https://doi.org/10.1017/S0956796818000096>
- Kazutaka Matsuda and Meng Wang. 2018b. Embedding invertible languages with binders: a case of the FliPpr language. In *Proceedings of the 11th ACM SIGPLAN International Symposium on Haskell, Haskell@ICFP 2018, St. Louis, MO, USA, September 27-17, 2018*, Nicolas Wu (Ed.). ACM, 158–171. <https://doi.org/10.1145/3242744.3242758>
- Kazutaka Matsuda and Meng Wang. 2018c. HOBiT: Programming Lenses Without Using Lens Combinators. In *ESOP (Lecture Notes in Computer Science)*, Amal Ahmed (Ed.), Vol. 10801. Springer, 31–59. https://doi.org/10.1007/978-3-319-89884-1_2
- Karl Mazurak, Jianzhou Zhao, and Steve Zdancewic. 2010. Lightweight linear types in system fdegree. In *TLDI*. ACM, 77–88.
- Conor McBride and Ross Paterson. 2008. Applicative programming with effects. *J. Funct. Program.* 18, 1 (2008), 1–13. <https://doi.org/10.1017/S0956796807006326>
- Torben A. Mogensen. 2005. Semi-inversion of Guarded Equations. In *Generative Programming and Component Engineering, 4th International Conference, GPCE 2005, Tallinn, Estonia, September 29 - October 1, 2005, Proceedings (Lecture Notes in Computer Science)*, Robert Glück and Michael R. Lowry (Eds.), Vol. 3676. Springer, 189–204. https://doi.org/10.1007/11561347_14
- Torben A. Mogensen. 2006. Report on an Implementation of a Semi-inverter. In *Ershov Memorial Conference (Lecture Notes in Computer Science)*, Irina Virbitskaite and Andrei Voronkov (Eds.), Vol. 4378. Springer, 322–334.
- Torben A. Mogensen. 2008. Semi-inversion of functional parameters. In *Proceedings of the 2008 ACM SIGPLAN Symposium on Partial Evaluation and Semantics-based Program Manipulation, PEPM 2008, San Francisco, California, USA, January 7-8, 2008*, Robert Glück and Oege de Moor (Eds.). ACM, 21–29. <https://doi.org/10.1145/1328408.1328413>
- Eugenio Moggi. 1998. Functor Categories and Two-Level Languages. In *Foundations of Software Science and Computation Structure, First International Conference, FoSSaCS'98, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS'98, Lisbon, Portugal, March 28 - April 4, 1998, Proceedings (Lecture Notes in Computer Science)*, Maurice Nivat (Ed.), Vol. 1378. Springer, 211–225. <https://doi.org/10.1007/BFb0053552>
- J. Garrett Morris. 2016. The best of both worlds: linear functional programming without compromise. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18-22, 2016*, Jacques Garrigue, Gabriele Keller, and Eijiro Sumii (Eds.). ACM, 448–461. <https://doi.org/10.1145/2951913.2951925>
- Shin-Cheng Mu and Richard S. Bird. 2003. Rebuilding a Tree from Its Traversals: A Case Study of Program Inversion. In *Programming Languages and Systems, First Asian Symposium, APLAS 2003, Beijing, China, November 27-29, 2003, Proceedings (Lecture Notes in Computer Science)*, Atsushi Ohori (Ed.), Vol. 2895. Springer, 265–282. https://doi.org/10.1007/978-3-540-40018-9_18
- Shin-Cheng Mu, Zhenjiang Hu, and Masato Takeichi. 2004a. An Algebraic Approach to Bi-directional Updating. In *Programming Languages and Systems: Second Asian Symposium, APLAS 2004, Taipei, Taiwan, November 4-6, 2004, Proceedings (Lecture Notes in Computer Science)*, Wei-Ngan Chin (Ed.), Vol. 3302. Springer, 2–20. https://doi.org/10.1007/978-3-540-30477-7_2
- Shin-Cheng Mu, Zhenjiang Hu, and Masato Takeichi. 2004b. An Injective Language for Reversible Computation. In *Mathematics of Program Construction, 7th International Conference, MPC 2004, Stirling, Scotland, UK, July 12-14, 2004, Proceedings (Lecture Notes in Computer Science)*, Dexter Kozen and Carron Shankland (Eds.), Vol. 3125. Springer, 289–313. https://doi.org/10.1007/978-3-540-27764-4_16
- Flemming Nielson and Hanne Riis Nielson. 1992. *Two-Level Functional Languages*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511526572>
- Naoki Nishida, Masahiko Sakai, and Toshiki Sakabe. 2005. Partial Inversion of Constructor Term Rewriting Systems. In *Term Rewriting and Applications, 16th International Conference, RTA 2005, Nara, Japan, April 19-21, 2005, Proceedings (Lecture Notes in Computer Science)*, Jürgen Giesl (Ed.), Vol. 3467. Springer, 264–278. https://doi.org/10.1007/978-3-540-32033-3_20
- Naoki Nishida and Germán Vidal. 2011. Program Inversion for Tail Recursive Functions. In *RTA (LIPICs)*, Manfred Schmidt-Schauß (Ed.), Vol. 10. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 283–298.
- Ross Paterson. 2012. Constructing Applicative Functors. In *MPC (Lecture Notes in Computer Science)*, Jeremy Gibbons and Pablo Nogueira (Eds.), Vol. 7342. Springer, 300–323. https://doi.org/10.1007/978-3-642-31113-0_15

- Tillmann Rendel and Klaus Ostermann. 2010. Invertible syntax descriptions: unifying parsing and pretty printing. In *Haskell*, Jeremy Gibbons (Ed.). ACM, 1–12.
- John C. Reynolds. 1998. Definitional Interpreters for Higher-Order Programming Languages. *Higher-Order and Symbolic Computation* 11, 4 (1998), 363–397. <https://doi.org/10.1023/A:1010027404223>
- Francisco Rios and Peter Selinger. 2017. A categorical model for a quantum circuit description language. In *Proceedings 14th International Conference on Quantum Physics and Logic, QPL 2017, Nijmegen, The Netherlands, 3-7 July 2017. (EPTCS)*, Bob Coecke and Aleks Kissinger (Eds.), Vol. 266. 164–178. <https://doi.org/10.4204/EPTCS.266.11>
- Patrick Maxim Rondon, Ming Kawaguchi, and Ranjit Jhala. 2008. Liquid types. In *Proceedings of the ACM SIGPLAN 2008 Conference on Programming Language Design and Implementation, Tucson, AZ, USA, June 7-13, 2008*, Rajiv Gupta and Saman P. Amarasinghe (Eds.). ACM, 159–169. <https://doi.org/10.1145/1375581.1375602>
- David Salomon. 2008. *A Concise Introduction to Data Compression*. Springer. <https://doi.org/10.1007/978-1-84800-072-8>
- Peter Selinger and Benoît Valiron. 2006. A lambda calculus for quantum computation with classical control. *Mathematical Structures in Computer Science* 16, 3 (2006), 527–552. <https://doi.org/10.1017/S0960129506005238>
- Saurabh Srivastava, Sumit Gulwani, Swarat Chaudhuri, and Jeffrey S. Foster. 2011. Path-based inductive synthesis for program inversion. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2011, San Jose, CA, USA, June 4-8, 2011*, Mary W. Hall and David A. Padua (Eds.). ACM, 492–503. <https://doi.org/10.1145/1993498.1993557>
- Jesse A. Tov and Riccardo Pucella. 2011. Practical affine types. In *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*, Thomas Ball and Mooly Sagiv (Eds.). ACM, 447–458. <https://doi.org/10.1145/1926385.1926436>
- Dimitrios Vytiniotis, Simon L. Peyton Jones, Tom Schrijvers, and Martin Sulzmann. 2011. OutsideIn(X) Modular type inference with local assumptions. *J. Funct. Program.* 21, 4-5 (2011), 333–412. <https://doi.org/10.1017/S0956796811000098>
- Philip Wadler. 1993. A Taste of Linear Logic. In *Mathematical Foundations of Computer Science 1993, 18th International Symposium, MFCS'93, Gdansk, Poland, August 30 - September 3, 1993, Proceedings (Lecture Notes in Computer Science)*, Andrzej M. Borzyszkowski and Stefan Sokolowski (Eds.), Vol. 711. Springer, 185–210. https://doi.org/10.1007/3-540-57182-5_12
- David Walker. 2004. Substructural Type Systems. In *Advanced Topics in Types and Programming Languages*, Benjamin C. Pierce (Ed.). MIT Press, 3–43.
- Meng Wang, Jeremy Gibbons, Kazutaka Matsuda, and Zhenjiang Hu. 2013. Refactoring pattern matching. *Science of Computer Programming* 78, 11 (2013), 2216 – 2242. <https://doi.org/10.1016/j.scico.2012.07.014> Special section on Mathematics of Program Construction (MPC 2010) and Special section on methodological development of interactive systems from Interaccion 2011.
- Hongwei Xi and Frank Pfenning. 1998. Eliminating Array Bound Checking Through Dependent Types. In *Proceedings of the ACM SIGPLAN '98 Conference on Programming Language Design and Implementation (PLDI), Montreal, Canada, June 17-19, 1998*, Jack W. Davidson, Keith D. Cooper, and A. Michael Berman (Eds.). ACM, 249–257. <https://doi.org/10.1145/277650.277732>
- Tetsuo Yokoyama, Holger Bock Axelsen, and Robert Glück. 2008. Principles of a reversible programming language. In *Proceedings of the 5th Conference on Computing Frontiers, 2008, Ischia, Italy, May 5-7, 2008*, Alex Ramirez, Gianfranco Bilardi, and Michael Gschwind (Eds.). ACM, 43–54. <https://doi.org/10.1145/1366230.1366239>
- Tetsuo Yokoyama, Holger Bock Axelsen, and Robert Glück. 2011. Towards a Reversible Functional Language. In *RC (Lecture Notes in Computer Science)*, Alexis De Vos and Robert Wille (Eds.), Vol. 7165. Springer, 14–29. https://doi.org/10.1007/978-3-642-29517-1_2
- Tetsuo Yokoyama, Holger Bock Axelsen, and Robert Glück. 2012. Optimizing Reversible Simulation of Injective Functions. *Multiple-Valued Logic and Soft Computing* 18, 1 (2012), 5–24. <http://www.oldcitypublishing.com/journals/mvlsc-home/mvlsc-issue-contents/mvlsc-volume-18-number-1-2012/mvlsc-18-1-p-5-24/>