

Monotonicity Types for Distributed Dataflow

Kevin Clancy*

Heather Miller†

ABSTRACT

The Lasp [2] programming language provides combinator functions such as Union and Intersection for combining set CRDTs. When designing a CRDT combinator, care must be taken to ensure that the combinator is monotone separately in each of its arguments, so that applying it to a tuple of increasing input streams yields an increasing output stream. We consider designing a type system which can prove the monotonicity of CRDT combinators.

ACM Reference format:

Kevin Clancy and Heather Miller. 2017. Monotonicity Types for Distributed Dataflow. In *Proceedings of ACM Conference, Washington, DC, USA, July 2017 (Conference'17)*, 10 pages. DOI: 10.1145/nmnnnnn.nnnnnnn

1 INTRODUCTION

The Lasp language provides a convenient framework for programming distributed systems in which edge computations are performed with minimal need to synchronize with servers. The Lasp programmer defines one or more increasing streams of state-based CRDT values. A stream of CRDT type T provides a *monotonic read* operation, which takes a threshold value of type T and yields control until the stream produces a value that is greater than or equal to the threshold with respect to a semilattice order defined over the values of T.

Another fundamental tool provided by Lasp is the ability to combine several CRDT streams. For example, from two increasing streams whose values belong to a set CRDT, we can create a new increasing stream whose values are equal to the unions of the values at the heads of the two input streams. Such a *CRDT combinator* is executed by a process acting as follows. The process first waits for a change to a value at the head of any of its input streams. When a change is detected, the process applies a function to the tuple of values taken from the heads of each input stream, writing the result of this application to an output stream. In the set union example, when one of the two input sets grows, the two sets at the heads of the input streams are used as arguments to the set union function.

The values of the output stream are joined into a third CRDT instance. Lest joining subsequent values of the output stream becomes idempotent, care must be taken so that they increase along

*Northeastern University

†Northeastern University

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference'17, Washington, DC, USA

© 2017 ACM. 978-x-xxxx-xxxx-x/YY/MM...\$15.00

DOI: 10.1145/nmnnnnn.nnnnnnn

the semilattice order. In order to guarantee this, we require that CRDTs be combined using *monotonic m-ary functions*. An *m-ary function* f mapping an *m*-tuple of CRDT values to a CRDT value is a *monotonic m-ary function* whenever

$$s \leq s' \implies f(\dots, s, \dots) \leq f(\dots, s', \dots)$$

for each of the *m* argument positions of f . In words, a monotonic *m-ary function* is a function that is monotone in each of its arguments separately when all other arguments are held fixed. When applied to an *m*-tuple of increasing CRDT streams, a monotonic *m-ary function* produces an increasing output stream. This property is crucial for Lasp; an output stream should increase along its CRDT's semilattice order, just as a stream generated from a CRDT replica does. This way we can perform monotonic reads on it and provide it as input to other CRDT combinators.

2 MONOTONICITY TYPING

Lasp provides a fixed API of CRDTs and CRDT combinators. Included is an *Observed-Remove Set* CRDT, which models arbitrary set addition and removal monotonically. It also provides combinators for producing the union, intersection, and cartesian product of two observed-remove sets. Programmers may want to construct their own CRDTs along with combinators acting upon them. In such cases, we believe it would be helpful to have a type system which ensures that the implementations of CRDTs and CRDT combinators respect CRDT semantics. In particular, letting \leq_T denote the semilattice order of CRDT T and letting $s.u(a)$ denote the invocation of a method u on CRDT instance s with argument a , such a system would prove the following properties:

- A CRDT's update methods are *inflative*; that is, if s is an instance of CRDT T then for all update methods u and arguments a , $s \leq_T s.u(a)$.
- CRDT combinators are *monotone in each argument separately*; that is, if $f : T_1 \times \dots \times T_m \rightarrow T$ is an *m-ary CRDT combinator* then for all $i \in 1..m$ and $s_i, s'_i \in T_i$ we have that $s_i \leq_{T_i} s'_i$ implies $f(\dots, s_i, \dots) \leq_T f(\dots, s'_i, \dots)$.

Properties such as inflativeness and monotonicity are not typically tracked by type systems. But because these properties propagate systematically across function composition, we expect that a simple, compositional, deductive system similar to a type system can be employed.

In the context of logic programming, Datafun[1] provides a concrete example of a type system which tracks monotonicity. Datafun imposes a top-down style of monotonicity reasoning which may prove counterintuitive. To create a single-argument monotone function, the programmer gives the function abstraction a special syntax $(\lambda x.t)$, writing the bound variable x in bold to indicate that it is monotone. Within the abstraction body t , x may only occur monotonically. For a term-in-context $\Gamma \vdash t$, the variable x occurs

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58

monotonically when for any pair of closing substitutions γ and γ' of Γ with $\gamma(x) \leq \gamma'(x)$ we have $\gamma t \leq \gamma' t$.

It's not clear that programmers can easily determine which positions in a term are monotone. A sound type system is not effective if the programmer does not intuitively understand the rules which govern it, because writing a program then entails a grueling process of trial and error.

Another arguable weakness of Datafun is that monotonicity is "hard coded" into Datafun's type system. The typing context is partitioned into two parts: one for variables which may only be used monotonically and the other for variables without usage restrictions. There are separate application rules for regular and monotone functions. For general purpose programming, it may be preferable to use a type system in which monotonicity is just one of the many tracked properties (including inflativeness) which propagate systematically across function composition.

Finally, Datafun requires the programmer to consider partial orderings over function types. This may be less intuitive than our approach, which avoids the need for partially ordering function types by providing primitive multi-argument functions.

3 TOWARD A CRDT TYPE SYSTEM

3.1 Notational preliminaries

We use colored overlines to represent vectors of syntactic objects. When such a vector is indexed by a set, the index set is indicated by a superscript. When vectors of the same color are juxtaposed, we assume they have the same length or index set. When vectors of different colors are juxtaposed, we assume they have differing lengths or index sets. When omitting a superscript, the index set should be clear from context. The superscript $i \in 1..n$ accompanying an overline indicates that the index set consists of all integers from 1 to some implicitly introduced natural number n . As an example, $\overline{x_i}^{i \in 1..n}$ denotes the vector of variables x_1, x_2, \dots, x_n .

3.2 Type Syntax

The catalog of state-based CRDTs described in [4] suggests that a CRDT's semilattice order can be derived in a straightforward manner from the CRDT's type structure. As a starting point, we need base types such as Natural, with the standard ordering on natural numbers, 0 as a bottom element, and maximum as the join operator. We would also include semilattice base types for sets (parameterized by element type) and booleans. A type such as String, whose values do not inherently form a bounded join-semilattice, could be lifted to one by ordering its values discretely and adding top and bottom elements. Compound types such as records, arrays, and maps would be ordered coordinatewise by default. For example, we would define the order \leq_T of the record type $T = \{l : T_l\}$ as follows: $\{l = v_l\} \leq_T \{l = u_l\} \iff \overline{v_l} \leq_{T_l} \overline{u_l}$. The state based Last-Writer-Wins Register would require a special pair type that is ordered lexicographically rather than coordinatewise.

3.3 Example

Consider implementing an intersection combinator for the 2P-Set CRDT described in [4]. A potential implementation of such a combinator is given in Figure 1. In this program the CRDT type $2PSet$

```
1 type 2Pset = {A : IntSet, R : IntSet}
2
3 fun 2PIntersect(a : 2Pset, b : 2Pset) : 2Pset[↑a, ↑b] =
4   { A = intersect(a.A, b.A), R = union(a.R, b.R) }
```

Figure 1: An intersection combinator for the 2P-Set




Figure 2: A Hasse diagram of the partial order on qualifiers

is defined as a record, ordered coordinatewise, with labeled components A (for added elements) and R (for removed elements) each belonging to a standard $IntSet$ datatype. The $2PIntersect$ function's return type annotation is augmented with the *qualifiers* $\uparrow a$ and $\uparrow b$, meaning that it is monotone separately in the formal arguments named a and b . In addition to the \uparrow qualifier, which denotes monotonicity, we also include the qualifier \downarrow for antitonicity, $=$ for arguments which are equal to their function's result, $?$ for arguments whose relation to their function's result is unknown, and \sim for superfluous arguments which do not affect their function's result. In the spirit of refinement types [6], these qualifiers are organized into a preordered set (in this case, a poset) which induces a subtyping relation. The partial order among qualifiers is shown in the Hasse diagram of Figure 2.

We prove $2PIntersect$'s monotonicity with respect to its formal arguments a and b in a bottom-up manner, by qualifying each of the subphrases of its body with respect to a and b . We will demonstrate this process in an informal manner. As a starting point, a has type $2Pset[= a, \sim b]$ and b has type $2Pset[\sim a, = b]$. $a.A$ and $a.R$ then have type $2Pset[\uparrow a, \sim b]$, while $b.A$ and $b.R$ have type $2Pset[\sim a, \uparrow b]$ due to the coordinatewise ordering of record types.

We compute the type of "intersect($a.A$, $b.A$)" by viewing it as a function composition. $intersect$'s functional characteristics are given by the type $(x : IntSet, y : IntSet) \Rightarrow IntSet[\uparrow x, \uparrow y]$, provided by a built-in API. The terms $a.A$ and $b.A$ defining $intersect$'s actual arguments are both functions of $2PIntersect$'s formal arguments a and b , and so their types are qualified with respect to these formal arguments. Denoting $a.A$ and $b.A$ with functional notations $f_x(a, b)$ and $f_y(a, b)$ respectively, and denoting $intersect$ with the functional notation $g(x, y)$, our goal is then to compute qualifiers for the function $h(a, b) = g(f_x(a, b), f_y(a, b))$.

Computing these qualifiers is a two-step process. We start by computing qualifiers for the function

$$\hat{h}(a_x, b_x, a_y, b_y) = g(f_x(a_x, b_x), f_y(a_y, b_y))$$

We compute the qualifier for a_x by combining f_x 's qualifier for a (\uparrow) with g 's qualifier for x (\uparrow) using a qualifier composition operator \circ described in Figure 3, yielding the qualifier \uparrow . Qualifiers for b_x , a_y , and b_y are obtained similarly. Finally, to obtain h 's qualifier for a , we combine \hat{h} 's qualifiers for a_x (\uparrow) and a_y (\sim) using a *qualifier*

o		?	↑	↓	=	~
?		?	?	?	?	~
↑		?	↑	↓	↑	~
↓		?	↓	↑	↓	~
=		?	↑	↓	=	~
~		~	~	~	~	~

Figure 3: Qualifier composition \circ

+		?	↑	↓	=	~
?		?	?	?	=	?
↑		?	↑	?	=	↑
↓		?	?	↓	=	↓
=		=	=	=	=	=
~		?	↑	↓	=	~

Figure 4: Qualifier contraction $+$

contraction operator $+$, described in Figure 4, yielding \uparrow . h 's qualifier for b is obtained similarly.

4 FORMALIZATION

We present a simple calculus containing novel features which could be used for tracking monotonicity. Along with this, we also present an outline of a soundness proof. We are currently in the process of proving the stated theorems.

4.1 Introduction

The computed type of a program term conveys the form of the value to which it reduces. For example, $2 + 3$ is not an *Int* but instead a term, consisting of two integers and an operator, which reduces to the *Int* 5. Likewise $\text{intersect}(a.A, b.A)$ is not a *IntSet* but a term which reduces to one. Recall that we computed the type $\text{IntSet}[\uparrow a, \uparrow b]$ for $\text{intersect}(a.A, b.A)$. This is supposed to mean not only that the expression $\text{intersect}(a.A, b.A)$ reduces to an *IntSet* when a and b are substituted with *2PSets*, but also that substituting a with a larger *2PSet* will result in an expression which reduces to a larger *IntSet* (and likewise for b).

Since monotonicity is about the values resulting from multiple instantiations of the variables of a term rather than just a single instantiation, it falls outside the standard static typing paradigm.

To make sense of this we define a language in Figure 5 which is similar to a standard lambda calculus but with an additional function form called the *sfun*. An *sfun* $(\tilde{\lambda}(x : B). i)$ is a first-order, multi-argument function intended for defining CRDT combinators such as Section 3.3's *2PIntersect*. While a term t which falls outside the scope of any *sfun* abstraction is said to belong to the *terminal language*, the body i of an *sfun* belongs to a superset of the terminal language called the *lifted language*, and is typed under a separate *lifted typing relation*.

We define a standard reduction relation, referred to as the *terminal reduction relation*, on our terminal language. Terminal reduction is statically characterized by a *terminal typing relation*. The lifted language's reduction relation is for normalizing the body of an *sfun* to a set-of-pairs representation that is by Theorem 4.2 extensionally

x, y, z		Variables
Type Syntax		
$q ::= \uparrow \mid \downarrow \mid \sim \mid = \mid ?$		Qualifiers
$\Xi ::= \emptyset \mid \Xi, q x$		Qualifier maps
$A, B ::= \text{Int} \mid \text{Bool}$		Base types
$\hat{A}, \hat{B} ::= B \mid B[\Xi]$		Lifted base types
$S, T, U ::= B \mid S \rightarrow T \mid \overline{(x_i : B_i^{i \in 1..n})} \Rightarrow A[\Xi]$		Types
$\hat{S}, \hat{T}, \hat{U} ::= S \mid \hat{B} \mid \hat{S} \rightarrow \hat{T}$		Lifted types
Environments		
$\Gamma ::= \emptyset \mid \Gamma, x : T$		Type environments
$\Omega, \Phi ::= \emptyset \mid \Omega, x : B$		Ambient environments
$\Delta ::= \emptyset \mid \Delta, x : \hat{T}$		Lifted type environments
$\gamma, \omega, \phi, \delta$		Environment valuations
Term Syntax		
$c ::= \text{true} \mid \text{false} \mid 0 \mid -1 \mid 1 \mid -2 \mid 2 \dots$		Base values
$\hat{c} ::= c \mid \langle \overline{\omega} \mapsto c \rangle$		Lifted base values
$v, w ::= c \mid (\lambda x : S. t) \mid (\tilde{\lambda}(x : B). i)$		Values
$\hat{v}, \hat{w} ::= v \mid \hat{c} \mid (\lambda x : \hat{S}. \hat{t})$		Lifted values
$s, t, u ::= v \mid x \mid t \mid t[\hat{t}]$		Terms
$\hat{s}, \hat{t}, \hat{u} ::= t \mid \hat{v} \mid \hat{t} \mid \hat{t}[\hat{t}]$		Lifted terms
Evaluation contexts		
$E ::= [] \mid E t \mid v E \mid E[\hat{t}] \mid v[\overline{v} E \hat{t}]$		Evaluation contexts
$\hat{E} ::= [] \mid \hat{E} \hat{t} \mid \hat{v} \hat{E} \mid \hat{E}[\hat{t}] \mid \hat{v}[\overline{v} \hat{E} \hat{t}]$		Lifted evaluation contexts

Figure 5: Syntax and environments

equivalent to the *sfun* itself; while lifted reduction is not computationally tractable, it is statically characterized by a *lifted typing relation* which can prove this set of pairs monotone, indirectly proving that the *sfun* is monotone under terminal reduction.

The lifted language generalizes the terminal language's base type constants to functions of the enclosing *sfun*'s formal parameters. In the lifted language, a terminal language constant such as the integer 1 is reinterpreted as a constant-valued function which maps any valuation of the enclosing *sfun*'s formal parameters to the value 1. The lifted language includes another form of base type value called the *ambient map*; an ambient map $\langle \overline{\omega} \mapsto c \rangle$ is a function—represented extensionally as a set of pairs rather than syntactically as a lambda abstraction—from the enclosing *sfun*'s domain to values of a specific terminal base type. Collectively, constants and ambient maps are referred to as *base values*, a class of syntactic objects defined in Figure 5 as productions of the metavariable \hat{c} .

To account for this enrichment of base values at the type level, we extend each base type into a poset of refinements. Specifically, this poset is a product of the qualifier set of Figure 2, associating one product component to each formal argument of the enclosing *sfun*, ordered componentwise. For example, in addition to type *Int* (which is the type of integer-valued constants and ambient maps which are constant with respect to the formals of the enclosing *sfun*) our lifted language's type syntax also includes the type $\text{Int}[\uparrow a, \uparrow b]$ of integer-valued constants and ambient maps which are monotone

with respect to formal arguments a and b . As an example of the preorder on these types, consider that $\text{Int}[\sim a, = b] \leq \text{Int}[\uparrow a, \uparrow b]$ because $\sim \leq \uparrow$ and $= \leq \uparrow$.

Whereas the terminal reduction relation models a series of function applications, lifted reduction interleaves function application with a special form of function composition. In this form of function composition, the outputs of n ambient maps are forwarded as inputs to an n -ary sfun abstraction. In the example of Section 3.3, the expression $h(a, b) = g(f_x(a, b), f_y(a, b))$ represents one such composition, where f_x and f_y are the ambient maps and g is the binary sfun into which their results are forwarded. Such a composition results in an ambient map, which we'll discuss in more detail in Section 4.8.

From now on, we will often refer to the formals of an sfun as *ambient variables*, because each base value is implicitly a function of these variables.

4.2 Term syntax

We use the metavariables s, t , and u for terms, and use \hat{s}, \hat{t} , and \hat{u} for lifted terms.

We've already introduced the syntactic form $\langle \omega \mapsto c \rangle$ for ambient maps and $(\tilde{\lambda}(x : B). t)$ for sfuns. Also note that the lifted language provides a new abstraction form $(\lambda x : \hat{S}. \hat{t})$ for functions whose formal argument type and body are lifted.

A function application is written as the juxtaposition of two terms, whereas the application of an sfun s to a vector \vec{t} of arguments is written $s [\vec{t}]$.

4.3 Type syntax

We use the metavariables S, T , and U for terminal types and \hat{S}, \hat{T} , and \hat{U} for lifted types. As with terms, the lifted type syntax is a superset of the terminal type syntax.

A lifted base type $B[\Xi]$ refines a terminal base type B with a *qualifier map* Ξ , which associates each ambient variable in scope with a qualifier q . We'll use the abbreviation $B[= x]$ for $B[\Xi]$ when Ξ maps x to $=$ and all other ambient variables in context to \sim .

Terminal types include base types B and function types $S \rightarrow T$. They also include sfun types $(x_i : B_i^{i \in 1..n}) \Rightarrow A[\Xi]$ where $x_i : B_i$ is a vector of typed formal arguments and $A[\Xi]$ is a lifted base type which qualifies each of the formal arguments \vec{x}_i . The lifted language does not provide an additional constructor for sfun types because we consider an sfun abstraction as encapsulated from the ambient variables of its context.

4.4 Type environments and valuations

A type environment Γ is a mapping from variables to terminal types. When $x \in \text{dom}(\Gamma)$ we write $\Gamma(x)$ to denote the type to which Γ maps x . A lifted type environment Δ is like a type environment, but maps variables to lifted types. An ambient environment Ω , which represents the formals of an sfun, is a mapping from variables to terminal base types. The notations $\Omega(x)$ and $\Delta(x)$ are analogous to $\Gamma(x)$.

A *valuation* γ of a type environment Γ maps each variable $x \in \text{dom}(\Gamma)$ to a value of type $\Gamma(x)$. We write $\gamma(x)$ for the value to which γ maps x . The symbols ω and ϕ are used for valuations of ambient environments and δ for valuations of lifted type environments. The

$$\begin{array}{c} \text{RED-CONTEXT} \\ \frac{t \rightarrow t'}{E[t] \rightarrow E[t']} \\ \\ \text{RED-SFUN-APP} \\ \frac{\omega = [x \mapsto v] \quad \delta \in \mathcal{G}_{x:B}[\overline{x : B[= x]}]}{(\tilde{\lambda}(x : B). t) [\vec{v}] \rightarrow \|\omega\| \delta t} \end{array} \quad \begin{array}{c} \text{RED-APP} \\ \frac{}{(\lambda x : S. t) v \rightarrow [x/v]t} \end{array}$$

Figure 6: Terminal reduction

$$C[\text{Bool}] \doteq \{\text{true}, \text{false}\}$$

$$C[\text{Int}] \doteq \{0, 1, -1, 2, -2, \dots\}$$

$$\mathcal{V}[\overline{x : B} \Rightarrow A[\Xi]] \doteq \{v \mid \forall c \in \mathcal{V}[B]. v [c] \in \mathcal{T}[A]\}$$

$$\mathcal{V}[x : S \rightarrow T] \doteq \{v \mid \forall w \in \mathcal{V}[S]. v w \in \mathcal{T}[T]\}$$

$$\mathcal{V}[B] \doteq \{c \mid c \in C[B]\}$$

$$\mathcal{T}[T] \doteq \{t \mid t \Downarrow v \in \mathcal{V}[T]\}$$

$$\mathcal{G}[\Gamma, x : T] \doteq \{\gamma[x \mapsto v] \mid \gamma \in \mathcal{G}[\Gamma] \wedge v \in \mathcal{V}[T]\}$$

$$\mathcal{G}[\emptyset] \doteq \{\emptyset\}$$

$$O[\Omega, x : B] \doteq \{\omega[x \mapsto c] \mid \omega \in O[\Omega] \text{ and } c \in \mathcal{V}[B]\}$$

$$O[\emptyset] \doteq \{\emptyset\}$$

Figure 7: Terminal logical relations

notations $\omega(x)$ and $\delta(x)$ are analogous to $\gamma(x)$. Substitution of a valuation γ into a term t is written γt . While substitutions of the form γt and δt are defined in a standard manner (and thus elided), we will not typically substitute ambient valuations ω this way. This is because the ambient variables of Ω live in a namespace distinct from that of the variables in Γ and Δ . The former can occur only in the domains of ambient maps and qualified base types, whereas the latter occur only as terms. Accordingly, we will use a special form of substitution for ambient valuations, described in Section 4.7 and Figure 8.

4.5 Terminal reduction

Under terminal reduction, described in Figure 6, an application is reduced by substituting the value on the right-hand side into the body of the abstraction on the left-hand side. Sfun application must not only substitute the supplied arguments into the body of the sfun, but also convert the body of the sfun from a lifted term to a terminal one by dropping all qualifier maps Ξ from type ascriptions. This is accomplished using an ambient substitution.

We will write $t \Uparrow$ to indicate that a term t diverges under terminal reduction, and write $t \Downarrow v$ to indicate that it normalizes to a value v .

4.6 Terminal logical relations

In Figure 7 we define a collection of predicates on closed terms which is indexed by the set of terminal types. We also define a collection environment-indexed predicates on substitutions. These definitions follow the *logical relations* technique (an accessible introduction to

$$\begin{array}{ll}
\|\omega\| t & = t \\
\|\omega\| \langle \dots \omega \mapsto c_\omega \dots \rangle & = c_\omega \\
\|\omega\| (\lambda x : \dot{S}. \dot{t}) & = (\lambda x : \|\omega\| \dot{S}. \|\omega\| \dot{t}) \\
\|\omega\| (\dot{t} \dot{i}) & = (\|\omega\| \dot{t}) (\|\omega\| \dot{i}) \\
\|\omega\| \dot{i} [\dot{t}] & = \|\omega\| \dot{i} [\|\omega\| \dot{t}] \\
\|\omega\| S & = S \\
\|\omega\| \dot{S} \rightarrow \dot{T} & = \|\omega\| \dot{S} \rightarrow \|\omega\| \dot{T} \\
\|\omega\| B[\Xi] & = B
\end{array}$$

Figure 8: Ambient substitution

which can be found in chapter 12 of [3]) for reasoning about typed languages.

A typing judgment for a terminal term conveys information about the behavior of that term under terminal reduction; the terminal logical relations capture this behavior. More concretely, we will prove that if $\Gamma \vdash t : T$ then for all $\gamma \in \mathcal{G}[\Gamma]$ we have $\gamma t \in \mathcal{T}[T]$. As a consequence, all well-typed terminal terms normalize under terminal reduction. The toy language which we present lacks recursion, and so normalization is not hard to achieve. We will discuss a technique in Section 5.4 for extending our language with a controlled form of recursion that precludes the possibility of divergence.

4.7 Ambient substitution

Let Ω be an ambient environment. An ambient map $\langle \overline{\phi} \mapsto c_\phi \rangle_{\phi \in \mathcal{O}[\Omega]}$ of domain $\mathcal{O}[\Omega]$ represents a constant whose value is undetermined but can be resolved given some valuation $\omega \in \mathcal{O}[\Omega]$. The *ambient substitution* $\|\omega\| \langle \overline{\phi} \mapsto c_\phi \rangle$ performs such a resolution, yielding c_ω . More generally, a lifted term \dot{t} whose ambient maps share the domain $\mathcal{O}[\Omega]$ represents a terminal term with several undetermined constants in various locations. These constants can be collectively resolved by a single $\omega \in \mathcal{O}[\Omega]$. Such a resolution is written $\|\omega\| \dot{t}$ and defined inductively in Figure 8.

4.8 Lifted reduction

While ambient substitution resolves the undetermined constants of a lifted term, the lifted reduction relation of Figure 9 allows us to reduce a lifted term before performing any such resolution. The LRED-SFUN-APP rule reduces an sfun application by brute force, performing each possible ambient substitution, normalizing under terminal reduction, and finally collecting the results into an ambient map. For each $\omega \in \mathcal{O}[\Omega]$, it has one premise of the form $(\tilde{\lambda}(x : B). \dot{i}) [\|\omega\| \dot{c}] \Downarrow c_\omega$.

Under an ambient environment Ω , a term of the form $(\tilde{\lambda}(x : B). \dot{i}) [\dot{c}]$ can be thought of as a composition of the same sort as the example

$$h(a, b) = g(f_x(a, b), f_y(a, b))$$

described in Section 3.3, given the following identifications:

- (1) The ambient environment Ω corresponds to the variables a and b .
- (2) The base values \dot{c} correspond to the functions $f_x(a, b)$ and $f_y(a, b)$.
- (3) The sfun $(\tilde{\lambda}(x : B). \dot{i})$ corresponds to the function g .

$$\begin{array}{c}
\text{LRED-CONTEXT} \\
\frac{\Omega \vdash \dot{i} \rightarrow \dot{i}'}{\Omega \vdash \dot{E}[\dot{i}] \rightarrow \dot{E}[\dot{i}']} \\
\text{LRED-APP} \\
\frac{}{\Omega \vdash (\lambda x : \dot{S}. \dot{t}) \dot{v} \rightarrow [x/\dot{v}]\dot{t}} \\
\text{LRED-SFUN-APP} \\
\frac{(\tilde{\lambda}(x : B). \dot{t}) [\|\omega\| \dot{c}] \Downarrow c_\omega}{\Omega \vdash (\tilde{\lambda}(x : B). \dot{t}) [\dot{c}] \rightarrow \langle \overline{\omega} \mapsto c_\omega \rangle}
\end{array}$$

Figure 9: Lifted reduction

We write $\Omega \vdash \dot{t} \uparrow$ to indicate divergence and $\Omega \vdash \dot{t} \Downarrow \dot{v}$ to indicate convergence to a lifted value \dot{v} under lifted reduction. Theorem 4.2 tells us that lifted reduction properly characterizes terminal reduction.

LEMMA 4.1. *If $\Omega \vdash \dot{i} \rightarrow \dot{i}'$ then for all $\omega \in \mathcal{O}[\Omega]$, $\|\omega\| \dot{i} \rightarrow^* \|\omega\| \dot{i}'$*

PROOF. By induction on derivations of $\Omega \vdash \dot{i} \rightarrow \dot{i}'$. \square

THEOREM 4.2. *If $\Omega \vdash \dot{t} \Downarrow \dot{v}$ then for all $\omega \in \mathcal{O}[\Omega]$, $\|\omega\| \dot{t} \Downarrow \|\omega\| \dot{v}$.*

PROOF. A simple corollary of Lemma 4.1. \square

4.9 Lifted logical relations

The lifted logical relations in Figure 10 are analogous to the terminal ones, except that types are interpreted through the lens of lifted reduction and ambient substitution. A new class \mathcal{K} of logical relations provides semantic interpretations in terms of ambient substitution for the various qualifiers. $\mathcal{X}[\Omega]$ denotes the set of qualifier maps on ambient environment Ω . Because lifted reduction is performed under some ambient environment Ω , we parameterize our relations by ambient environments, which appear as the subscripts beneath $\mathcal{K}, \mathcal{V}, \mathcal{T}$, and \mathcal{G} .

Because the lifted types are a superset of the terminal types, we must take care to define lifted logical relations for terminal base types as well as qualified ones. Under an ambient environment Ω , a terminal base type B is interpreted equivalently to $B[\Xi]$ where Ξ assigns every ambient variable in $\text{dom}(\Omega)$ to the qualifier \sim . This reflects that inside an sfun body a value of terminal type typically originates from outside the sfun and therefore should not depend on the arguments to which the sfun is applied.

4.10 Subtyping

Defined in Figure 11, a terminal subtyping judgment $S <: T$ means that the set of values of type S is contained in the set of values of type T . Terminal subtyping judgments are derived using standard rules for base and function types. The S-SFUN rule for sfun subtyping requires matching argument types and defers to lifted subtyping to ensure that the qualifiers of the result types respect the qualifier ordering.

Defined in Figure 12, a lifted subtyping judgment $\Omega \vdash \dot{S} <: \dot{T}$ means that under ambient environment Ω , the set of lifted values of type \dot{S} is contained in the set of lifted values of type \dot{T} . Four rules are required to handle base types, since both the left and right side can be either qualified or terminal. Because the lifted logical relations interpret terminal base types B equivalently to $B[\Xi]$ where

1	$\mathcal{X}[\Omega, x : B]$	\doteq	$\{\Xi, q \ x \mid \Xi \in \mathcal{X}[\Omega]\}$
2	$\mathcal{X}[\emptyset]$	\doteq	$\{\emptyset\}$
3			
4	$\mathcal{K}_\Omega[\uparrow x]$	\doteq	$\{\dot{c} \mid \forall \omega, \omega' \in \mathcal{O}[\Omega]. (\omega(x) \leq \omega'(x) \wedge \forall y \in \text{dom}(\Omega) - \{x\}. \omega(y) = \omega'(y)) \implies \ \omega\ \dot{c} \leq \ \omega'\ \dot{c}\}$
5	$\mathcal{K}_\Omega[\downarrow x]$	\doteq	$\{\dot{c} \mid \forall \omega, \omega' \in \mathcal{O}[\Omega]. (\omega(x) \leq \omega'(x) \wedge \forall y \in \text{dom}(\Omega) - \{x\}. \omega(y) = \omega'(y)) \implies \ \omega\ \dot{c} \geq \ \omega'\ \dot{c}\}$
6	$\mathcal{K}_\Omega[\sim x]$	\doteq	$\{\dot{c} \mid \forall \omega, \omega' \in \mathcal{O}[\Omega]. (\forall y \in \text{dom}(\Omega) - \{x\}. \omega(y) = \omega'(y)) \implies \ \omega\ \dot{c} = \ \omega'\ \dot{c}\}$
7	$\mathcal{K}_\Omega[= x]$	\doteq	$\{\dot{c} \mid \forall \omega \in \mathcal{O}[\Omega]. \ \omega\ \dot{c} = \omega(x)\}$
8	$\mathcal{K}_\Omega[? x]$	\doteq	$\{\dot{c} \mid \text{true}\}$
9			
10	$\mathcal{V}_\Omega[\overline{(x_i : B_i)^{i \in 1..n}} \Rightarrow A[\Xi]]$	\doteq	$\{\dot{v} \mid \overline{\forall \Xi_i \in \mathcal{X}[\Omega]. \forall \dot{c}_i \in \mathcal{V}_\Omega[B_i[\Xi]]}. \dot{v} [\dot{c}_i] \in \mathcal{T}_\Omega[A[\overline{(\sum_{i=1}^n \Xi_i(z) \circ \Xi(x_i)) z^{z \in \text{dom}(\Omega)}}]]}\}$
11	$\mathcal{V}_\Omega[\dot{S} \rightarrow \dot{T}]$	\doteq	$\{\dot{v} \mid \forall \dot{w} \in \mathcal{V}_\Omega[\dot{S}]. \dot{v} \dot{w} \in \mathcal{T}_\Omega[\dot{T}]\}$
12	$\mathcal{V}_\Omega[B[\Xi]]$	\doteq	$\{\dot{c} \mid (\forall \omega \in \mathcal{O}[\Omega]. \ \omega\ \dot{c} \in \mathcal{V}[B]) \wedge (\dot{c} \in \bigcap_{x \in \text{dom}(\Omega)} \mathcal{K}_\Omega[\Xi(x) x])\}$
13	$\mathcal{V}_\Omega[B]$	\doteq	$\{\dot{c} \mid (\forall \omega \in \mathcal{O}[\Omega]. \ \omega\ \dot{c} \in \mathcal{V}[B]) \wedge (\dot{c} \in \bigcap_{x \in \text{dom}(\Omega)} \mathcal{K}_\Omega[\sim x])\}$
14			
15	$\mathcal{T}_\Omega[\dot{T}]$	\doteq	$\{\dot{t} \mid \Omega \vdash \dot{t} \downarrow \dot{v} \in \mathcal{V}_\Omega[\dot{T}]\}$
16			
17	$\mathcal{G}_\Omega[\Delta, x : \dot{T}]$	\doteq	$\{\delta[x \mapsto \dot{v}] \mid \delta \in \mathcal{G}_\Omega[\Delta] \wedge \dot{v} \in \mathcal{V}_\Omega[\dot{T}]\}$
18	$\mathcal{G}_\Omega[\emptyset]$	\doteq	$\{\emptyset\}$
19			
20			

Figure 10: Lifted logical relations

S-SFUN	S-FUN	
$\frac{x : B \vdash A_1[\Xi_1] <: A_2[\Xi_2]}{(x : B) \Rightarrow A_1[\Xi_1] <: (x : B) \Rightarrow A_2[\Xi_2]}$	$\frac{T_1 <: S_1 \quad S_2 <: T_2}{S_1 \rightarrow S_2 <: T_1 \rightarrow T_2}$	$ty(0), ty(-1), ty(1), \dots = Int$
		$ty(true), ty(false) = Bool$

Figure 13: Constant types metafunction

Figure 11: Terminal subtyping

S-BASE	
$\frac{}{B <: B}$	
LS-SFUN	T-CONST
$\frac{x : B \vdash \dot{A}_1[\Xi_1] <: \dot{A}_2[\Xi_2]}{\Omega \vdash (x : B) \Rightarrow A_1[\Xi_1] <: (x : B) \Rightarrow A_2[\Xi_2]}$	$\frac{}{\Gamma \vdash c : ty(c)}$
LS-FUN	T-VAR
$\frac{\Omega \vdash \dot{T}_1 <: \dot{S}_1 \quad \dot{S}_2 <: \dot{T}_2}{\Omega \vdash \dot{S}_1 \rightarrow \dot{S}_2 <: \dot{T}_1 \rightarrow \dot{T}_2}$	$\frac{x \in \text{dom}(\Gamma)}{\Gamma \vdash x : \Gamma(x)}$
LS-BASE-TT	T-APP
$\frac{}{\Omega \vdash B <: B}$	$\frac{\Gamma \vdash t : S \rightarrow T \quad \Gamma \vdash s : S}{\Gamma \vdash t s : T}$
LS-BASE-TL	T-SFUN-APP
$\frac{\sim \leq q_x}{\Omega \vdash B <: B[\overline{q_x} x^{x \in \text{dom}(\Omega)}]}$	$\frac{\Gamma \vdash u : (x_i : A_i) \Rightarrow B[\Xi] \quad \Gamma \vdash s_i : A_i}{\Gamma \vdash u [\overline{s_i}^{i \in 1..n}] : B}$
LS-BASE-LT	T-FUN
$\frac{q_x \leq \sim}{\Omega \vdash B[\overline{q_x} x^{x \in \text{dom}(\Omega)}] <: B}$	$\frac{\Gamma, x : S \vdash t : T}{\Gamma \vdash (\lambda x : S. t) : S \rightarrow T}$
LS-BASE-LL	T-SFUN
$\frac{p_x \leq q_x}{\Omega \vdash B[\overline{p_x} x^{x \in \text{dom}(\Omega)}] <: B[\overline{q_x} x]}$	$\frac{\Gamma; x_i : B_i; x_i : B_i [= x_i] \vdash \dot{u} : A[\Xi]}{\Gamma \vdash (\overline{\lambda(x_i : B_i)^{i \in 1..n}}. \dot{u}) : (x_i : B_i) \Rightarrow A[\Xi]}$
	T-SUB
	$\frac{\Gamma \vdash t : S \quad S <: T}{\Gamma \vdash t : T}$

Figure 14: Terminal typing

The following theorems ensure that every subtyping derivation implies a containment among logical relations.

THEOREM 4.3 (FUNDAMENTAL TERMINAL SUBTYPING THEOREM). *If $S <: T$ then $\mathcal{V}[S] \subseteq \mathcal{V}[T]$.*

PROOF. By induction on the derivation of $S <: T$. □

THEOREM 4.4 (FUNDAMENTAL LIFTED SUBTYPING THEOREM). *If $\Omega \vdash \dot{S} <: \dot{T}$ then $\mathcal{V}_\Omega[\dot{S}] \subseteq \mathcal{V}_\Omega[\dot{T}]$.*

PROOF. By induction on the derivation of $\Omega \vdash \dot{S} <: \dot{T}$. □

Ξ qualifies every ambient variable as \sim , B and $B[\Xi]$ are isomorphic with respect to subtyping; i.e., $\Omega \vdash B <: B[\Xi]$ and $\Omega \vdash B[\Xi] <: B$.

Figure 12: Lifted subtyping

$\frac{}{\Gamma; \Omega; \Delta \vdash c : ty(c)}$	$\frac{}{\Gamma; \Omega; \Delta \vdash x : \Gamma(x)} \quad x \in dom(\Gamma)$	$\frac{}{\Gamma; \Omega; \Delta \vdash x : \Delta(x)} \quad x \in dom(\Delta)$
$\frac{LT-APP \quad \Gamma; \Omega; \Delta \vdash \dot{s} : \dot{S} \rightarrow \dot{T} \quad \Gamma; \Omega; \Delta \vdash \dot{s} : \dot{S}}{\Gamma; \Omega; \Delta \vdash \dot{s} \dot{s} : \dot{T}}$		
$\frac{LT-SFUN-APP \quad \Gamma; \Omega; \Delta \vdash \dot{u} : (x_i : B_i) \Rightarrow A[p_i x_i] \quad \Gamma; \Omega; \Delta \vdash \dot{s}_i : B_i[\overline{q_{iz} z}]}{\Gamma; \Omega; \Delta \vdash \dot{u} [\dot{s}_i^{i \in 1..n}] : A[\sum_{i=1}^n q_{iz} \circ p_i] z^{z \in dom(\Omega)}}$		
$\frac{LT-FUN \quad \Gamma; \Omega; \Delta, x : \dot{S} \vdash \dot{t} : \dot{T} \quad \Omega \vdash \dot{S}}{\Gamma; \Omega; \Delta \vdash (\lambda x : \dot{S}. \dot{t}) : \dot{S} \rightarrow \dot{T}}$	$\frac{T-SUB \quad \Gamma; \Omega; \Delta \vdash \dot{t} : \dot{S} \quad \Omega \vdash \dot{S} <: \dot{T}}{\Gamma; \Omega; \Delta \vdash \dot{t} : \dot{T}}$	

Figure 15: Lifted typing

4.11 Typing

The terminal typing relation is described in Figure 14. The only interesting rule is T-SFUN for sfun abstractions, which types the body of the sfun under the lifted typing relation with new ambient and lifted environments formed from the sfun's formal arguments. Recall from Section 4.4 that ambient variables live in a namespace separate from those of the terminal and lifted type environments; T-SFUN leverages this by adding identically-named variables to both ambient and lifted environments.

The lifted typing relation is described in Figure 15. Its context includes both a terminal type environment Γ and a lifted type environment Δ . Because every terminal value is also a lifted value, substituting terminal values for the variables of Γ in \dot{t} does not pose a problem. This typing relation does not include a rule for sfun abstractions; we consider programs with nested sfun abstractions ill-typed. However, the body of an sfun abstraction may still refer to other sfun abstractions through variables. As a consequence, a well-typed term may step to an ill-typed term when an sfun abstraction is substituted through the boundary of another sfun abstraction. Lack of type preservation should not pose a problem to our soundness proof outline, because it uses logical relations.

The only novel lifted typing rule is LT-SFUN-APP, for typing an sfun application $\dot{u} [\dot{s}_i^{i \in 1..n}]$. The first premise requires \dot{u} to have an n -ary sfun type $(x_i : A_i) \Rightarrow B[p_i x_i]$. The other n premises require that each argument \dot{s}_i has a lifted base type of the form $A_i[\overline{q_{iz} z}^{z \in dom(\Omega)}]$, where $\overline{q_{iz} z}$ is a qualifier map synthesized from type checking \dot{s}_i . The sfun application $\dot{u} [\dot{s}_i]$ then has the type $B[\sum_{i=1}^n q_{iz} \circ p_i]$, where \circ is the operator defined in Figure 3 and the "summation" $\sum_{i=1}^n \dots$ uses the qualifier contraction operator \vdash defined in Figure 4.

LT-FUN requires a lifted type well-formedness derivation of $\Omega \vdash \dot{S}$, which ensures that all qualifiers in \dot{S} are contained in the ambient environment Ω . The lifted type well-formedness rules have been elided because they are not interesting.

The following fundamental typing theorems imply that our language is sound.

THEOREM 4.5 (FUNDAMENTAL TERMINAL TYPING THEOREM). *For all ambient environments Ω , if $\Gamma \vdash t : T$ then for all $\gamma \in \mathcal{G}[\Gamma] \cap \mathcal{G}_\Omega[\Gamma]$, $\gamma t \in \mathcal{T}[\Gamma] \cap \mathcal{T}_\Omega[\Gamma]$.*

PROOF. By induction on the derivation of $\Gamma \vdash t : T$. \square

THEOREM 4.6 (FUNDAMENTAL LIFTED TYPING THEOREM). *If $\Gamma; \Omega; \Delta \vdash \dot{t} : \dot{T}$ then for all $\gamma \in \mathcal{G}_\Omega[\Gamma]$ and all $\delta \in \mathcal{G}_\Omega[\Delta]$, $\gamma \delta \dot{t} \in \mathcal{T}_\Omega[\dot{T}]$.*

PROOF. By induction on the derivation of $\Gamma; \Omega; \Delta \vdash \dot{t} : \dot{T}$. \square

The intersections in Theorem 4.5 are necessary to allow the substitution of terminal values through sfun abstractions.

4.12 Why logical relations?

Why reason about the soundness of our language using logical relations rather than the more popular approach of progress and preservation? Under lifted reduction, type preservation across sfun applications would prove problematic. Explained briefly, we care about type preservation across normalization rather than single-step reduction.

More specifically, imagine trying to prove type preservation for a well typed term of the form $(\tilde{\lambda}(x_i : B_i^{i \in 1..n}). \dot{t}) [\dot{c}_i]$. By LRED-SFUN-APP, such a term reduces in a single step to an ambient map, but deducing a lifted base type for this ambient map requires characterizing $(\tilde{\lambda}(x_i : B_i). \dot{t})$ extensionally. Because $(\tilde{\lambda}(x_i : B_i). \dot{t})$ is well-typed we might invoke an inversion lemma to get $\emptyset; \Omega; \emptyset \vdash (\tilde{\lambda}(x_i : B_i). \dot{t}) : (x_i : B_i) \Rightarrow A[\Xi]$ and $\emptyset; x_i : B_i; x_i : B_i [= x_i] \vdash \dot{t} : \dot{A}$. A substitution lemma may tell us that for $\delta \in \mathcal{G}_{x_i : B_i}[\llbracket x_i : B_i [= x_i] \rrbracket]$ we have $\emptyset; x_i : B_i; \emptyset \vdash \delta \dot{t} : \dot{A}$. Not knowing the derivation height of this latter judgment, we could not apply the inductive hypothesis to it. Even if we could, we would only conclude that $x_i : B_i \vdash \delta \dot{t} \rightarrow \dot{t}'$ for some \dot{t}' where $\emptyset; x_i : B_i; \emptyset \vdash \dot{t}' : \dot{A}$, whereas we actually need type preservation across normalization.

5 ADDITIONAL FEATURES

The language of Section 4 has been kept simple to avoid distracting from its novel feature: the sfun. Calling it a toy language would be charitable however, as it lacks recursion, conditionals, and the other features discussed below.

5.1 Sfun primitives

Sfuns abstractions only become useful when we have primitive sfuns with which to define them. When working with integers, for example, we'd like an addition operator with type

$$(x : Int, y : Int) \Rightarrow Int[\uparrow x, \uparrow y]$$

and a subtraction operator with type

$$(x : Int, y : Int) \Rightarrow Int[\uparrow x, \downarrow y]$$

Likewise, for working with sets we'd like sfun primitives including union, intersection, and set difference.

5.2 Divergence dilemma

The standard approach to recursion using fixpoint combinators would introduce the possibility that well typed terms diverge. Logical relations tend to accommodate recursion by defining

$$\mathcal{T}[[T]] \doteq \{t \mid t \uparrow \text{ or } t \Downarrow v \text{ where } v \in \mathcal{V}[[T]]\}$$

Such definitions leverage the idea that if t fails to normalize then we can soundly claim it belongs to any type T , since the set of values it might normalize to is vacuously contained in $\mathcal{V}[[T]]$. However, conditionals would introduce the possibility that some applications of an sfun normalize while other applications of the same sfun diverge. In this case, given that lifted reduction simulates the terminal reduction of all possible applications simultaneously, it would be reasonable for the sfun body to diverge under lifted reduction. This is unsound, as it allows typing derivations associating any qualifier map to the sfun's result type, even ones which contradict its set of normalizing applications.

5.3 Dependent refinement types

We would like to solve the aforementioned dilemma using dependent refinement types, drawing inspiration from LiquidHaskell [5]. In this approach a refined base type has the form $\{v : B \mid t\}$, where B is an unrefined base type such as *Bool* or *Int*, and t is a term of type *Bool*. Such a type describes the set of constants c of base type B such that $[c/v]t \Downarrow \text{true}$. Standard function types give way to dependent function types, which have the form $x : S \rightarrow T$. The variable x , which is scoped over type T , refers to the value to which a function of type $x : S \rightarrow T$ is applied.

Because types may contain variables, subtyping becomes parameterized over a typing context, which is significant to the following subtyping rule for refined base types.

$$\frac{\text{SUB-BASE} \quad \forall \gamma \in \mathcal{G}[[\Gamma, v : B]]. \gamma s \Downarrow \text{true} \implies \gamma t \Downarrow \text{true}}{\Gamma \vdash \{v : B \mid s\} <: \{v : B \mid t\}}$$

While this rule's premise is undecidable, it can be conservatively approximated by an SMT solver, which is the approach taken by LiquidHaskell.

In a dependent refinement type system, the type assigned to a variable by a type environment may depend on the variables which precede it, so $\mathcal{G}[-]$ would be redefined with:

$$\mathcal{G}[[\Gamma, x : T]] \doteq \{\gamma[x \mapsto v] \mid \gamma \in \mathcal{G}[[\Gamma]] \text{ and } v \in \mathcal{V}[[\gamma T]]\}$$

Our fundamental terminal subtyping and typing theorems would also change:

THEOREM 5.1 (DEPENDENT TERMINAL SUBTYPING THEOREM). *If $\Gamma \vdash S <: T$ then for all $\gamma \in \mathcal{G}[[\Gamma]]$, $\mathcal{V}[[\gamma S]] \subseteq \mathcal{V}[[\gamma T]]$.*

THEOREM 5.2 (DEPENDENT TERMINAL TYPING THEOREM). *If $\Gamma \vdash t : T$ then for all ambient environments Ω with $FV(\Omega) = \emptyset$ and all $\gamma \in \mathcal{G}[[\Gamma]] \cap \mathcal{G}_\Omega[[\Gamma]]$, we have $\gamma t \Downarrow v$ and $v \in \mathcal{V}[[\gamma T]] \cap \mathcal{V}_\Omega[[\gamma T]]$.*

Our lifted typing contexts would have dependencies as well. An entry of the ambient environment would depend on entries of the terminal environment. An entry of the lifted environment would depend both on entries of the terminal environment and preceding entries of the lifted environment.

```

1 fun f(metric : Nat, n : Int,
2   i : {v : Int | v ≤ n ∧ (metric = n-i)}) =
3   if i = n then
4     i
5   else
6     let metric' = n - (i+1) in
7     i + f metric' n (i+1)

```

Figure 16: Decreasing, well founded termination metric

The fundamental lifted subtyping and typing theorems would be revised:

THEOREM 5.3 (DEPENDENT LIFTED SUBTYPING THEOREM).

If $\Gamma; \Omega; \Delta \vdash S <: \dot{T}$ then for all $\gamma \in \mathcal{G}[[\Gamma]] \cap \mathcal{G}_{\gamma\Omega}[[\Gamma]]$ and $\delta \in \mathcal{G}_{\gamma\Omega}[[\gamma\Delta]]$ we have $\mathcal{V}_{\gamma\Omega}[[\gamma\delta S]] \subseteq \mathcal{V}_{\gamma\Omega}[[\gamma\delta \dot{T}]]$.

THEOREM 5.4 (DEPENDENT LIFTED TYPING THEOREM).

If $\Gamma; \Omega; \Delta \vdash t : \dot{T}$ then for all $\gamma \in \mathcal{G}[[\Gamma]] \cap \mathcal{G}_{\gamma\Omega}[[\Gamma]]$ and $\delta \in \mathcal{G}_{\gamma\Omega}[[\gamma\Delta]]$ we have $\gamma\delta t \in \mathcal{T}_{\gamma\Omega}[[\gamma\delta \dot{T}]]$.

5.4 Terminating fixpoints

A dependent refinement type system can be extended with a *terminating fixpoint* construct, which allows recursion given a decreasing, well founded termination metric. Using *Nat* as a shorthand for $\{v : \text{Int} \mid 0 \leq v\}$, Figure 16 demonstrates a function with such a termination metric. The function f , which given integers n and i with $i \leq n$ computes the sum of integers from i to n , has a termination metric as its leading argument. A typechecker can then prove this function terminating by checking its body under a weakened context in which f has the type:

$$m' : \{v : \text{Nat} \mid v < \text{metric}\} \\ \rightarrow n : \text{Int} \rightarrow i : \{v : \text{Int} \mid (v \leq n) \wedge (m' = n - i)\}$$

Using this technique LiquidHaskell is able to prove 96% of all recursive functions terminating [5].

LiquidHaskell treats the terminating fixpoint combinator *tfix* as a built-in constant, but to address its subtleties we treat it as a first class component of the language, which we summarize in Figure 17. We add a type constructor $\langle x < t \rangle T$ with the following introduction form

$$(\text{tfix } t \ x \ T \ y \ s)$$

This introduction form represents a suspended recursive computation where t is the current termination metric, x is a variable for the subsequent termination metric, T is the type of the recursive term (which should depend on x), y is a recursive self-reference, and s is a term (which should include occurrences of x and y) representing the computation itself. The elimination form $u \langle s \rangle$ resumes the recursive computation u using s as the subsequent value of its termination metric.

5.5 Another example

As a final example, Figure 18 describes a CRDT combinator for adding two GCounter CRDTs. We define the GCounter type in terms of a language primitive called *IntArray*, an array of integers ordered componentwise. The sfun *getAt*, for getting the element at index i of array a , is monotone with respect to a because *IntArray* is ordered

$$\begin{array}{c}
\text{RED-TFIX} \\
(\mathbf{tfix} \ v \ x \ T \ y \ s) \langle v' \rangle \rightarrow [v'/x][(\mathbf{tfix} \ v' \ x \ T \ y \ s)/y]s \\
\\
\text{T-TFIX-ABS} \\
\frac{\Gamma \vdash t : \mathit{Nat} \quad \Gamma, x : \{v : \mathit{Nat} \mid v < t\} \vdash T \quad \Gamma, x : \{v : \mathit{Nat} \mid v < t\}, y : \langle x' < x \rangle [x'/x]T \vdash s : T}{\Gamma \vdash (\mathbf{tfix} \ t \ x \ T \ y \ s) : \langle x < t \rangle T} \\
\\
\text{T-TFIX-APP} \\
\frac{\Gamma \vdash u : \langle x < t \rangle U \quad \Gamma \vdash s : \{v : \mathit{Nat} \mid v < t\}}{\Gamma \vdash u \langle s \rangle : [s/x]U}
\end{array}$$

Figure 17: Terminating fixpoint

componentwise. Its relation to i is unknown (thus qualified with ?) since we make no assumptions about the way distinct elements of an *IntArray* are related. The sfun *setAt* sets element i of *IntArray* a to the integer v . It is monotone with respect to a and v due to *IntArray*'s componentwise ordering, and unknown with respect to i due to the lack of information about how distinct elements of the array are related. The *len* function gets the length of an *IntArray* and the *makeArray* function creates a new array with all elements initialized to 0.

The nested function *sumCell* uses terminating recursion to iterate through the indices of the two arrays.

Because the type $\{v : \mathit{IntArray} \mid \mathit{len}(v) = n\}$ is isomorphic to

$$\{v : \mathit{IntArray} \mid \mathit{len}(v) = n\}[\sim x, \sim y]$$

the type of the *IntArray* returned from *makeArray* on line 19 is subsumed by

$$\{v : \mathit{IntArray} \mid \mathit{len}(v) = n\}[\uparrow x, \uparrow y]$$

Assuming that each call to *setAt* does not make a new copy of the entire array but instead returns a reference to the existing one, updating arrays in this manner is unsound. An array update invalidates the types of existing aliases to that array. With such a need for strong updates, it may be worth investigating the addition of linear or affine types.

6 CONCLUSION

While the proposed approach may have practical benefits, we are still attempting to provide it with a suitable operational model and soundness proof. Nonetheless, the daunting number of possibilities for combining CRDTs in a Lasp-style language provides strong motivation for the type checking technique that we have outlined.

```

1 1  getAt :: (a : IntArray, i : {v:Nat | v < len(a)}) ⇒ Int[↑ a, ? i]
2 2  setAt :: (a : IntArray, i : {v:Nat | v < len(a)}, v : Int) ⇒ IntArray[↑ a, ? i, ↑ v]
3 3  len :: IntArray -> Int
4 4  makeArray :: n:Nat -> { v : IntArray | len(v) = n }
5 5  + :: (x:Int, y:Int) ⇒ Int[↑ x, ↑ y]
6 6
7 7  type GCounter = { array : IntArray }
8 8
9 9  sfun sumCounters(x : GCounter, y : {v:GCounter| len(v) = len(x)}) : GCounter[↑ x, ↑ y] =
10 10    let xArray : IntArray[↑ x, ~ y] = x.array
11 11    let yArray : IntArray[~ x, ↑ y] = y.array
12 12    //termination metric: (len acc) - ind
13 13    fun sumCell(ind : Nat, acc : IntArray[↑ x, ↑ y]) : IntArray[↑ x, ↑ y] =
14 14      if ind = (len acc) then
15 15        acc
16 16      else
17 17        let acc' = setAt acc ind ((getAt xArray ind) + (getAt yArray ind))
18 18        sumCell (ind+1) acc'
19 19    let newArray : IntArray[↑ x, ↑ y] = makeArray (len x)
20 20    GCounter { array = sumCell 0 newArray }

```

Figure 18: A CRDT combinator computing the sum of two GCounters

REFERENCES

- [1] Michael Arntzenius and Neelakantan R Krishnaswami. Datafun: a functional datalog. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming*, pages 214–227. ACM, 2016.
- [2] Christopher Meiklejohn and Peter Van Roy. Lasp: A language for distributed, coordination-free programming. In *Proceedings of the 17th International Symposium on Principles and Practice of Declarative Programming*, pages 184–195. ACM, 2015.
- [3] Benjamin C Pierce. *Types and programming languages*. MIT press, 2002.
- [4] Marc Shapiro, Nuno Preguiça, Carlos Baquero, and Marek Zawirski. *A comprehensive study of convergent and commutative replicated data types*. PhD thesis, Inria-Centre Paris-Rocquencourt; INRIA, 2011.
- [5] Niki Vazou, Eric L Seidel, Ranjit Jhala, Dimitrios Vytiniotis, and Simon Peyton-Jones. Refinement types for haskell. In *ACM SIGPLAN Notices*, volume 49, pages 269–282. ACM, 2014.
- [6] Noam Zeilberger. Principles of type refinement. <http://noamz.org/oplss16/refinements-notes.pdf> [Online; accessed 29-April-2017], 2016.