

“TOWARDS” HOMOMORPHIC COMPUTATION FOR DISTRIBUTED COMPUTING

Christopher S. Meiklejohn
Kevin Clancy
Heather Miller

Université catholique de Louvain
Instituto Superior Técnico
Northeastern University

UCL

Université
catholique
de Louvain



Northeastern University

WHY COORDINATION-FREE PROGRAMMING?

1. Cost of coordination

- Increasing latency in geo-replicated applications

2. CALM result and beyond [CIDR 2011]

- Convergence guaranteed with a combination of lattice-based programming and monotone logic
- Regardless of network anomalies: message duplication and/or message reordering

3. Key insights:

- $v1 \sqsubseteq v2$ means that $v1$ approximates $v2$ ($v2$ contains everything in $v1$ and possibly more)
- “Once something has happened, it continues to have happened...”
 - Bloom’s protocol programming, LVars “Freeze After Writing”

PROGRAMMING MODEL

Data Structures: Lattices
Programming Model: Lattice Combinators

Logic and Lattices for Distributed Programming

Neil Conway
UC Berkeley
nrc@cs.berkeley.edu

William R. Marczak
UC Berkeley
wrm@cs.berkeley.edu

Peter Alvaro
UC Berkeley
palvaro@cs.berkeley.edu

Joseph M. Hellerstein
UC Berkeley
hellerstein@cs.berkeley.edu

David Maier
Portland State University
maier@cs.pdx.edu

ABSTRACT

In recent years there has been interest in achieving application-level consistency criteria without the latency and availability costs of strongly consistent storage infrastructure. A stan-

ard technique for achieving consistency is to use a replication protocol that reorders operations; this approach is used by the CALM replication protocol. In logic programming, operations are automatically verified without coordination.

In this paper, we generalize Bloom's logic programming analysis to support efficient evaluation of lattice-based code using well-known strategies from logic programming. Finally, we use Bloom^L to develop several practical distributed programs, including a key-value store similar to Amazon Dynamo, and show how Bloom^L encourages the safe composition of small, easy-to-analyze lattices into larger programs.

ditionally, a replication protocol need only register commutative, associative, idempotent merge functions [19]. This approach has roots in research in databases and systems [12, 14, 17, 27, 39] as well as groupware [20]. Shapiro, et al. recently proposed a formalism for conflict-free replicated data types (CRDTs) that combines these ideas into the algebraic framework of CRDTs [5].

CRDTs address two problems: (a) the programmer must specify lattice properties for their data types (commutativity, idempotence), and (b) the programmer must specify operations for individual data objects, and how they are combined in general. As an example of this, we consider the following problem:

Multi-tolerant courseware study teams. It uses two sets of students for Teams. The application reads a version of the course catalog and inserts the derived element <Alice,Bob> into Teams. Concurrently, Bob is removed from Students by another application replica. The use of CRDTs ensures that all replicas will eventually agree that Bob is absent from Students, but this is not enough: application-level state is inconsistent unless the derived values in Teams are updated consistently to reflect Bob's re-

Bloom^L, SoCC 2012
Logic Programming

Lasp: A Language for Distributed, Coordination-Free Programming

Christopher Meiklejohn
Basho Technologies, Inc.
cmiklejohn@basho.com

Peter Van Roy
Université catholique de Louvain
peter.vanroy@uclouvain.be

Abstract

We propose Lasp, a new programming model designed to simplify large-scale distributed programming. Lasp combines ideas from deterministic dataflow programming together with conflict-free replicated data types (CRDTs). This provides support for computations where not all participants are online together at a given moment. The initial design presented here provides a framework for composing CRDTs, which lets us build a more tolerant distributed applications with non-monotonic data. Given reasonable message delays, node failures, we prove that Lasp can be considered as a functional programming language with functional reasoning and programming techniques. Lasp is implemented as an Erlang library built on the distributed systems framework. We have developed a large-scale application, the advertisement counter, as part of the SyncFree research project. We plan to extend Lasp into a general-purpose language in which it can be used as little as possible.

Categories and Subject Descriptors D.1.3 [Programming Languages]: Concurrent Programming; E.1 [Data Structures]: Distributed data structures

Keywords Eventual Consistency, Commutative Operations, Erlang

the burden is placed on the programmer of these applications to ensure that concurrent operations performed on replicated data have both a deterministic and desirable outcome.

For example, consider the case where a user's gaming profile is replicated between two mobile devices. Concurrent operations, which can be thought of as operations performed during the period where both clients are online, can be performed. Communication, can modify the profile.

Lasp, PPDP 2015
Functional Programming

updates. The current profile is used to compute the next profile.

al. for that are types, (b), provided; distributed state.

tic res-erty is highly desirable in a distributed system because it no longer places the resolution logic in the hands of the programmer; programmers are able to use replicated data types that function as if they were their sequential counterparts. However, it has been shown that arbi-

PROGRAMMING MODEL: EXAMPLE

A :: Set

B :: Set

B = map $\lambda x.x$ A

Whenever A changes, B is recomputed from A.

Combinators can be user defined.

A :: OR-Set

B :: OR-Set

B = fold $\lambda x \lambda y. x + y \perp$ A

Custom, non-trivial lattices.

Higher-order programming with functions.

CHALLENGES: STATE-OF-THE-ART

1 Lasp's proof wasn't comprehensive enough to catch composition failures.

attices or place the onus on the developer

2 No formal specification of either systems execution model.

If you follow the rules, you have "correct" programs.

If you don't... ?

3. Incrementality of computations

Without monotone checks, program behavior is nondeterministic.
(ie. function reads clock, function is antitone, etc.)

functions where function application distributes

putation as a homomorphism and correctly

CHALLENGES: LATTICES

Bloom^L

User-provided ADTs

Combinators are user-provided

Combinators must be labeled as either monotone or as a homomorphism

Laspl

Any CRDT implementing a “CRDT” interface is supported
- ordering relation, merge, inflation, etc.

Combinators are built in for one CRDT:
the Observed-Remove Set

Higher-order combinator, fold, provided
but requires user to ensure monotonicity

CHALLENGES: COMBINATORS

Bloom^L

Must be labeled as either monotone or homomorphic

Functions are unchecked as to whether they are monotone, or homomorphic

Lasp

Built in combinators designed to provide monotonicity

Higher-order programming with fold requires unchecked user-implemented function is both monotone and has inverse function

CHALLENGES: INCREMENTALITY

Bloom^L

Unchecked annotations on homomorphic functions, and are incrementally evaluated

Lasp

(Some) built-in combinators are homomorphic, but are not evaluated as so

CONTRIBUTIONS

1. Computational delta objects

- Generalization of CRDTs to objects have a unified format enabling easier computation
- Changes expressed as “deltas” that can be derived through object decomposition
- Combinators are implemented in terms of monotone functions on deltas

2. Monotonicity typing

- Instantiation of Petricek’s structural coeffects system to detect monotonicity violations

3. Operational semantics for an incremental calculus with lattices

- Expresses computations as functions from lattice to lattice
- Incremental evaluation when lattice computations form homomorphisms
- Generalized semantics for Bloom^L and Lasp: previously not formalized

DATA TYPES & COMBINATORS: BLOOM^L

Specified by the user

Must have:

- Values that form a lattice
- Least-upper-bound function
- Combinators

Homomorphism

Annotations are **unchecked**

```
1 class Bud::SetLattice < Bud::Lattice
2   wrapper_name :lset
3
4   def initialize(x=[])
5     # Input validation removed for brevity
6     @v = x.uniq # Remove duplicates from i
7   end
8
9   def merge(i)
10    self.class.new(@v | i.reveal)
11  end
12
13  morph :intersect do |i|
14    self.class.new(@v & i.reveal)
15  end
16
17  morph :contains? do |i|
18    Bud::BoolLattice.new(@v.member? i)
19  end
20
21  monotone :size do
22    Bud::MaxLattice.new(@v.size)
23  end
24 end
```

Least-upper-bound

Monotone Function

Figure 4: Example implementation of the lset lattice.

DATA TYPES & COMBINATORS: LASP

Implemented through a “CRDT” interface

Combinators are built into runtime:

- Functional: map, filter
- Set-theoretic: union, product, intersection
- Higher-order: fold

Homomorphism

Fold must be monotonic, **unchecked**

Monotone Function

3.5 Functional Programming

We now look at the semantics for functional programming primitives that are lifted to operate over CRDTs: **map**, **filter**, and **fold**. We formalize them as follows:

Definition 3.3. The **map** function defines a process that never terminates, which reads elements of the input stream s and creates elements in the output stream t . For each element, the value v is separated from the metadata, the function f is applied to the value, and new metadata is attached to the resulting value $f(v)$. If two or more values map to the same $f(v)$ (for instance, if the function provided to map is surjective), the metadata is combined into one triple for all values of v .

$$F(s_i, f) = \{f(v) \mid (v, -, -) \in s_i\}$$

$$A(s_i, f, w) = \bigcup \{a \mid (v, a, -) \in s_i \wedge w = f(v)\}$$

$$R(s_i, f, w) = \bigcup \{r \mid (v, -, r) \in s_i \wedge w = f(v)\}$$

$$\text{map}'(s_i, f) = \{(w, A(s_i, f, w), R(s_i, f, w)) \mid w \in F(s_i, f)\}$$

$$\text{map}(s, f) = t = [\text{map}'(s_i, f) \mid s_i \in s]$$

(9)

$$\text{product}'(s_i, u_j) = \{(v, v'), a \times a', a \times r' \cup r \times a' \mid (v, a, r) \in s_i, (v', a', r') \in u_j\}$$

$$\text{product}(s, u) = t = [\text{product}'(s_i, u_j) \mid s_i \in s, u_j \in u]$$

(12)

RECONCILING THE “DELTA”

Challenges:

1. What is a useful generalization of the data types in Lasp and Bloom^L?
 - Generalization to **delta lattices**, and a special case for non-compositional types called **multichains**
2. How do we encode a notion of incrementality into the programming model?
 - Each delta lattice can be decomposed into a set of **maximal deltas**

Key Insights:

1. Monotone functions from lattice A to lattice B
 - $(A_\delta \rightarrow B)$ on deltas decomposed from A, combined using an arbitrary function $(B \rightarrow B \rightarrow B)$
2. Homomorphic functions from lattice A to lattice B
 - Monotone function where the combining function is **join** (least-upper-bound)

DELTA LATTICES & MULTICHAINS

Delta lattices

- Bounded join-semilattices, S
- $J(S)$, the set of join irreducible elements taken from S , are finitely join-dense
- Maximal chains from $J(S)$ are mutually exclusive

Multichains $M(L, T)$

- Special case of delta lattices
- $L \rightarrow T$
- L : a set of unordered labels
- T : a totally ordered set

Any element from the delta lattice can be decomposed into a join of deltas

Deltas are maximal, removing redundancy

Multichain deltas are $(L, \max(T))$

COMPUTATIONAL DELTA LATTICES + OBJECTS

Computational delta lattices (S, A_δ, F, G)

- S is a delta lattice
- A_δ is a poset
- F is a function from S to A_δ only defined where S is a join-irreducible element take from S
- G is a function from A_δ to S where $G; F$ for some S is the identity function

Computational delta objects (A, M, Q)

- A is a computational delta lattice
- M is a set of monotone mutation functions that return deltas that will be joined into the store
- Q is a set of query functions over A

For type safety on monotone functions on A_δ

Convert from join-irreducible singleton set to π_1
ie. $\{1,2\} \rightarrow \{1\} \cup \{2\} \rightarrow 1 \text{ or } 2$

A practical instantiation of a CRDT-like lattice-based data type

COMPUTATIONAL DELTA OBJECTS: BLOOM^L

LOrd[T]:

Multichain:

$(1, T)$

Mutation functions:

$\text{set}(t, m) = (1, t)$

Query functions:

$\text{value}(m) = \pi_2 m$

LBool = LOrd[2]

LMax = LOrd[NMax]

LMin = LOrd[NMin]

Ordered register

2: ordered using the natural's
ordering

Naturals ordered with bottom at
zero or infinity

COMPUTATIONAL DELTA OBJECTS: CRDTS

G-Counter:

Multichain:

$(\mathbb{R}, \text{NMax})$

Mutation functions:

$\text{increment}(r, m) = (r, m(r) + 1)$

Query functions:

$\text{value}(m) = \text{fold } (\pi_2, +) a$

Grow-only counter is replica to
count mappings

G-Set[A]:

Multichain:

$(\mathbb{A}, 2)$

Mutation functions:

$\text{add}(a, m) = (a, 1)$

Query functions:

$\text{value}(m) = \text{map } \pi_1 m$

Grow-only set is element to 2

COMPUTATIONAL DELTA OBJECTS: LASP

Observed-Remove Set[A]:

Multichain:

$((A \times T \times D), 2)$

Mutation functions:

$\text{add}(a, t, m) = ((a, t, \perp), 1)$

$\text{rem}(a, t, m) = ((a, t, \top), 1)$

Query functions:

$\text{value}(m) = \text{map } \pi_1 \{(a, t) \mid (a, t, \perp) \in m\} \setminus \{(a, t) \mid (a, t, \top) \in m\}$

IMPLEMENTING COMBINATORS: CHALLENGES

1. Combinators must be monotone, to ensure **convergence**
 - Alvaro et al., “Consistency Analysis in Bloom: a CALM and Collected Approach”, CIDR 2011
2. Combinators must be from lattice to lattice, to ensure **composition**
 - Conway et al., “Logic and Lattices for Distributed Programming”, SoCC 2012
 - Meiklejohn et al., “Lasp: A Language for Distributed, Coordination-Free Programming”, PPDP 2015
3. Combinators must be a lattice homomorphism, to enable **incremental evaluation**
 - Conway et al., “Logic and Lattices for Distributed Programming”, SoCC 2012

COEFFECTS

Petricek's structural coefffects system

Extend context with a “coeffect”

- Vector, entry per item in the typing context
- “Scalar” value with an ordering relation
- Composition and a contraction operation

Important points:

- Contraction used to when variables occur multiple times in the body (join)
- Composition used for function application (sum)

Combine scalars for multiple occurrences in expression

$$\boxed{\Gamma @ R \vdash e : \tau}$$

(CONST)

$$\frac{}{\cdot @ \langle \rangle \vdash c : ty(c)}$$

(VAR)

$$\frac{}{x : \tau @ \langle use \rangle \vdash x : \tau}$$

(ABS)

$$\frac{\Gamma, x : \sigma @ \langle \dots, s \rangle \vdash e : \tau}{\Gamma @ \langle \dots \rangle \vdash \lambda x. e : \sigma \xrightarrow{s} \tau}$$

(APP)

$$\frac{\Gamma_1 @ R \vdash e_1 : \sigma \xrightarrow{t} \tau \quad \Gamma_2 @ S \vdash e_2 : \sigma}{\Gamma_1, \Gamma_2 @ R \times t \otimes S \vdash e_1 e_2 : \tau}$$

(CTXT)

$$\frac{\Gamma @ R \vdash e : \tau \quad \Gamma' @ R' \rightsquigarrow \Gamma @ R, \theta}{\Gamma' @ R' \vdash \theta e : \tau}$$

$$\boxed{\Gamma', R' \rightsquigarrow \Gamma, R, \theta}$$

(WEAK)

$$\dots, ign \rightsquigarrow \Gamma @ \langle \dots \rangle, \emptyset$$

$$\dots \rightsquigarrow \Gamma_1, x : \tau, y : \sigma, \Gamma_2 @ \langle \dots t, s \dots \rangle, \emptyset$$

$$\Gamma_1, y : \tau, z : \tau, \Gamma_2 @ \langle \dots s, t \dots \rangle, [y, z \mapsto x]$$

(SUB)

$$\Gamma_1, x : \tau, \Gamma_2 @ \langle \dots s' \dots \rangle \rightsquigarrow \Gamma_1, x : \tau, \Gamma_2 @ \langle \dots s \dots \rangle, \emptyset \quad (s \leq s')$$

Application joins and composes disjoint typing contexts and coefffects

MONOTONICITY TYPING

Instantiation of structural coefficients system

Labels represent:

- Arbitrary (?)
- Monotone (+)
- Antitone (-)
- Unknown (~)

Key insights:

- Variable occurrence is monotone
- Function application “composes” monotonicity
- Multiple occurrences are joined with contraction

\otimes	?	+	-	~
?	?	?	?	~

Monotone and monotone compose to monotone (+ with + yields +)

\oplus	?	+	-	~
?	?	?	?	?
+	?	+	?	+
-	?	?	-	-
~	?	+	-	~

Figure 3: Scalar composition \otimes and contraction \oplus for monotonicity

Antitone to antitone compose to monotone (- with - yields +)

~

Figure 4: Hasse diagram for the partial order \leq on monotonicity scalars

LIFTED MONOTONE FUNCTIONS

Given we want a monotone function from delta lattice $A \rightarrow$ delta lattice $B...$

We write:

- Mapping from A delta to B. $f: (A_\delta \rightarrow B)$
- Summing / combining function for B's. $g: (B \rightarrow B \rightarrow B)$

Ca

- V
- C

Apply the transformation and combine

as all combining functions **may not be commutative**

se

Execution:

- `fold (g ◦ f) ⊥B deltas (A)`

Decompose A into deltas

LIFTED HOMOMORPHIC FUNCTIONS

Given we want a homomorphic function from delta lattice $A \rightarrow$ delta lattice $B...$

We write:

- Mapping from A delta to B. $f: (\mathbb{A}_\delta \rightarrow$

Deltas taken from A

Execution:

- `fold ($\sqcup \circ f$) \perp_B deltas (A)`

Apply the transformation and combine using the join

MONOTONE FUNCTIONS: BLOOM^L

size on LSet

- $\text{size} :: \text{LSet} \rightarrow \text{LMax}$
- $f : \lambda((a, 1)). \{(1, 1)\}$
- $g : +$

Produce a set of deltas for B

sum on LSet

- $\text{sum} :: \text{LSet} \rightarrow \text{LMax}$
- $f : \lambda((a, 1)). \{(1, a)\}$
- $g : +$

Combining function is applied
coordinate-wise

We've completed a full specification of Bloom^L monotone functions, without considering the dictionary.

HOMOMORPHIC FUNCTIONS: BLOOM^L

$+(y)$ on LMax

- $+ :: N \rightarrow LMax \rightarrow LMax$
- $f : \lambda((1, X)). \{(1, x + y)\}$

Produce a set of deltas for B

$-(y)$ on LMax

- $- :: N \rightarrow LMax \rightarrow LMax$
- $f : \lambda((1, X)). \{(1, x - y)\}$

$\text{map}(g)$ on LSet

- $\text{map} :: (A \rightarrow B) \rightarrow LSet \rightarrow LSet$
- $f : \lambda((a, 1)). \{(g(a), 1)\}$

$\text{filter}(g)$ on LSet

- $\text{filter} :: (A \rightarrow \text{Bool}) \rightarrow LSet \rightarrow LSet$
- $f : \lambda((a, 1)). \{\}$

We've completed a full specification of Bloom^L homomorphisms, without considering the dictionary.

OPERATIONAL SEMANTICS

$$\mu; \pi; \rho; \phi \hookrightarrow \mu'; \pi'; \rho'; \phi'$$

R-MON-E

$$\frac{e \rightarrow e'}{\mu; \pi; \rho; \phi_1[l_1, l_2, v, D, e] \phi_2 \hookrightarrow \mu; \pi; \rho; \phi_1[l_1, l_2, v, D, e'] \phi_2}$$

R-MON-V-NEXT

$$\frac{}{\mu; \pi; \rho; \phi_1[l, v, d :: D, a] \phi_2 \hookrightarrow \mu; \pi; \rho; \phi_1[l, v, D, (v \ d) \ a] \phi_2}$$

R-MON-V-FIN

$$\frac{\text{getMonUpdates}(l, \perp \boxtimes a, \pi) = \phi' \quad \text{updateHom}(l, \mu(l) \boxtimes a, \rho) = \rho'}{\mu; \pi; \rho; \phi_1[l, v, \emptyset, a] \phi_2 \hookrightarrow \mu[l \mapsto a]; \pi; \rho'; \phi_1 \phi_2 \phi'}$$

R-HOM-E

$$\frac{e \rightarrow e'}{\mu; \pi; \rho_1[l_1, l_2, v, D, e] \rho_2; \phi \hookrightarrow \mu; \pi; \rho_1[l_1, l_2, v, D, e'] \rho_2; \phi}$$

R-HOM-V-NEXT

$$\frac{\text{getMonUpdates}(l_2, \perp \boxtimes (\mu(l_2) \sqcup a), \pi) = \phi' \quad \text{updateHom}(l_2, \mu(l_2) \boxtimes a, \rho_1[l_1, l_2, v, D, v \ d] \rho_2) = \rho'}{\mu; \pi; \rho_1[l_1, l_2, v, d :: D, a] \rho_2; \phi \hookrightarrow \mu[l_2 \mapsto \mu(l_2) \sqcup a]; \pi; \rho'; \phi}$$

R-HOM-V-FIN

$$\frac{\text{getMonUpdates}(l_2, \perp \boxtimes (\mu(l_2) \sqcup a), \pi) = \phi' \quad \text{updateHom}(l_2, \mu(l_2) \boxtimes a, \rho_1[l_1, l_2, v, \emptyset, \perp] \rho_2) = \rho' \quad a \neq \perp}{\mu; \pi; \rho_1[l_1, l_2, v, \emptyset, a] \rho_2; \phi \hookrightarrow \mu[l_2 \mapsto \mu(l_2) \sqcup a]; \pi; \rho'; \phi \phi'}$$

$a \in A$ (Lattice elements)

$d \in A_\delta$ (Deltas)

$m \in M$ (Delta mutators)

Monotone functions: decompose lattice into deltas and apply/combine

$\mu ::= \cdot \mid \mu[k \mapsto a]$ (Stores)

$\pi ::= \cdot \mid \pi[l_1, l_2, \lambda x.e]$ (Monotone link environment)

$\rho ::= \cdot \mid \rho[l_1, l_2, \lambda x.e, D, e]$ (Homomorphic link environment)

$\phi ::= \cdot \mid \phi[l, v, \langle d_1, \dots, d_n \rangle, e]$ (Pending monotone link updates)

Homomorphisms: apply one delta at a time

TA
, v) is defined
 $v \rightarrow \delta(c, v)$

FUTURE WORK

1. Typed Structures

- Multichains are only sufficient for non-compositional data types
- Typed records, lift homomorphisms to operate on multichains embedded in a typed record

2. Binary Operations w/o Fixed Arguments

- ex. Lasp cartesian product vs. Bloom^L cartesian product

3. Fixed Point Combinator

- Links from locations to themselves – but, how expressive can this be?

4. Type System

- ex. verify that monotone functions on deltas have the correct type signature

5. Proofs

- Progress and preservation
- Correspondence between monolithic and incremental evaluation

STRUCTURAL COEFFECTS: DATAFLOW EXAMPLE

Dataflow language over streams

“pre” operation for accessing a previous item in the stream

Coeffects are naturals where:

- Ordering is standard order on naturals
- Composition is +
- Contraction is “max”

$$\text{APP} \frac{\text{VAR} \frac{}{pre : B \xrightarrow{1} B@ \langle 0 \rangle} \vdash pre : B \xrightarrow{1} B} \quad \frac{}{x : B@ \langle 0 \rangle} \vdash x : B \text{VAR}}{pre : B \xrightarrow{1} B, x : B@ \langle 0, 1 \rangle} \vdash pre x : B$$

Typing context and coeffects are joined / composed.

OPERATIONAL SEMANTICS: SYNTAX/EXPRESSIONS

Homomorphic and monotone links
between store locations

$a \in A$ (Lattice elements)
 $d \in A_\delta$ (Deltas)
 $m \in M$ (Delta mutators)
 $q \in Q$ (Query functions)
 $x \in X$ = Set of variables
 $c \in C$ = Set of constants
 $e ::= \lambda x.e \mid e_1 e_2 \mid c \mid a \mid d \mid l \mid \mathbf{hlink} \ e_1 \ e_2 \ e_3 \mid \mathbf{mlink} \ e_1 \ e_2 \ e_3 \mid \mathbf{new} \mid \mathbf{update} \ e_1 \ e_2$
 $v ::= c \mid a \mid d \mid m \mid q \mid \lambda x.e \mid l \mid ()$

$\mathbf{hlink} \ e_1 \ e_2 \ e_3 \mid \mathbf{mlink} \ e_1 \ e_2 \ e_3 \mid \mathbf{new} \mid \mathbf{update} \ e_1 \ e_2$

Interaction with the store

$l \in \mathcal{L}$ = Finite set of locations
 $\mu ::= \cdot \mid \mu[k \mapsto a]$ (Stores)
 $\pi ::= \cdot \mid \pi[l_1, l_2, \lambda x.e]$ (Monotone link environment)
 $\rho ::= \cdot \mid \rho[l_1, l_2, \lambda x.e, D, e]$ (Homomorphic link environment)
 $\phi ::= \cdot \mid \phi[l, v, \langle d_1, \dots, d_n \rangle, e]$ (Pending monotone link updates)

$e \rightarrow e'$

$\frac{\text{APP-L} \quad e_1 \rightarrow e'_1}{e_1 e_2 \rightarrow e'_1 e_2}$	$\frac{\text{APP-R} \quad e_2 \rightarrow e'_2}{v_1 e_2 \rightarrow v_1 e'_2}$	$\frac{\text{BETA}}{(\lambda x.e) v \rightarrow e[x := v]}$	$\frac{\text{DELTA} \quad \delta(c, v) \text{ is defined}}{c v \rightarrow \delta(c, v)}$
--	--	---	---

Standard expression reduction rules
extended with constants and
functional constants

OPERATIONAL SEMANTICS: STORES

R-NEW

$$\frac{}{\mu; \pi; \rho; \phi; \mathbf{new} \rightarrow \mu[l \rightarrow \perp]; \pi; \rho; \phi()} \quad l \notin \text{dom}(\mu)$$

Create locations in the store

R-UPDATE-1

$$e_1 \rightarrow e'_1$$

Monotonic links forward location
state decomposed into deltas

update e'_1 e_2

update l e'_2

R-UPDATE-V

$$\frac{a = \mu(l) \quad \text{getMonUpdates}(l, \perp \boxtimes (a' \sqcup a)) = \phi' \quad \text{updateHom}(l, a \boxtimes a', \rho) = \rho'}{\mu; \pi; \rho; \phi; \mathbf{update} \ l \ a' \rightarrow \mu[l \mapsto a' \sqcup a]; \pi; \rho'; \phi'()}$$

Updates store, and propagates
deltas forward using links

Homomorphic links propagate
forward only the deltas
representing the change

OPERATIONAL SEMANTICS: LINKS

$\pi; \rho \rightarrow \pi'; \rho'$

HLINK-1

$$\frac{e_1 \rightarrow e'_1}{\pi; \rho; \mathbf{hlink} \ l_1 \ l_2 \ e_3 \rightarrow \pi; \rho; \mathbf{hlink} \ l_1 \ l_2 \ e'_3}$$

HLINK-2

$$\frac{e_2 \rightarrow e'_2}{\pi; \rho; \mathbf{hlink} \ l_1 \ e_2 \ e_3 \rightarrow \pi; \rho; \mathbf{hlink} \ l_1 \ e'_2 \ e_3}$$

HLINK-3

$$\frac{e_3 \rightarrow e'_3}{\pi; \rho; \mathbf{hlink} \ l_1 \ l_2 \ e_3 \rightarrow \pi; \rho; \mathbf{hlink} \ l_1 \ l_2 \ e'_3}$$

HLINK

$$\frac{}{\pi; \rho; \mathbf{hlink} \ l_1 \ l_2 \ \lambda x.e \rightarrow \pi; \rho[l_1, l_2, \lambda x.e, \emptyset, \perp]; ()}$$

MLINK-1

$$\frac{e_1 \rightarrow e'_1}{\pi; \rho; \mathbf{mlink} \ e_1 \ e_2 \ e_3 \rightarrow \pi; \rho; \mathbf{mlink} \ e'_1 \ e_2 \ e_3}$$

MLINK-2

$$\frac{e_2 \rightarrow e'_2}{\pi; \rho; \mathbf{mlink} \ l_1 \ e_2 \ e_3 \rightarrow \pi; \rho; \mathbf{mlink} \ l_1 \ e'_2 \ e_3}$$

MLINK-3

$$\frac{e_3 \rightarrow e'_3}{\pi; \rho; \mathbf{mlink} \ l_1 \ l_2 \ e_3 \rightarrow \pi; \rho; \mathbf{mlink} \ l_1 \ l_2 \ e'_3}$$

MLINK

$$\frac{}{\pi; \rho; \mathbf{mlink} \ l_1 \ l_2 \ \lambda x.e \rightarrow \pi[l_1, l_2, \lambda x.e]; \rho; ()}$$

Homomorphic links store pending deltas and a monotone function between store locations.

Monotone function track a monotone function between store locations

OPERATIONAL SEMANTICS: MLINKS

$$\mu; \pi; \rho; \phi \hookrightarrow \mu'; \pi'; \rho'; \phi'$$

$$\text{R-MON-E} \quad \frac{e \rightarrow e'}{\mu; \pi; \rho; \phi_1[l_1, l_2, v, D, e]\phi_2 \hookrightarrow \mu; \pi; \rho; \phi_1[l_1, l_2, v, D, e']\phi_2}$$

R-MON-V-NEXT

$$\frac{}{\mu; \pi; \rho; \phi_1[l, v, d :: D, a]\phi_2 \hookrightarrow \mu; \pi; \rho; \phi_1[l, v, D, (v \ d) \ a]\phi_2}$$

R-MON-V-FIN

$$\frac{\text{getMonUpdates}(l, \perp \boxtimes a, \pi) = \phi' \quad \text{updateHom}(l, \mu(l) \boxtimes a, \rho) = \rho'}{\mu; \pi; \rho; \phi_1[l, v, \emptyset, a]\phi_2 \hookrightarrow \mu[l \mapsto a]; \pi; \rho'; \phi_1\phi_2\phi'}$$

R-HOM-E

$$\frac{e \rightarrow e'}{\mu; \pi; \rho_1[l_1, l_2, v, D, e]\rho_2; \phi \hookrightarrow \mu; \pi; \rho_1[l_1, l_2, v, D, e']\rho_2; \phi}$$

R-HOM-V-NEXT

$$\frac{\text{getMonUpdates}(l_2, \perp \boxtimes (\mu(l_2) \sqcup a), \pi) = \phi' \quad \text{updateHom}(l_2, \mu(l_2) \boxtimes a, \rho_1[l_1, l_2, v, D, v \ d]\rho_2) = \rho'}{\mu; \pi; \rho_1[l_1, l_2, v, d :: D, a]\rho_2; \phi \hookrightarrow \mu[l_2 \mapsto \mu(l_2) \sqcup a]; \pi; \rho'; \phi}$$

R-HOM-V-FIN

$$\frac{\text{getMonUpdates}(l_2, \perp \boxtimes (\mu(l_2) \sqcup a), \pi) = \phi' \quad \text{updateHom}(l_2, \mu(l_2) \boxtimes a, \rho_1[l_1, l_2, v, \emptyset, \perp]\rho_2) = \rho' \quad a \neq \perp}{\mu; \pi; \rho_1[l_1, l_2, v, \emptyset, a]\rho_2; \phi \hookrightarrow \mu[l_2 \mapsto \mu(l_2) \sqcup a]; \pi; \rho'; \phi\phi'}$$

Computing pending deltas for each monotonic link and combine the results

Once completed for all deltas, replace the value in the store and forward on new deltas to out-links

OPERATIONAL SEMANTICS

$$\mu; \pi; \rho; \phi \hookrightarrow \mu'; \pi'; \rho'; \phi'$$

$$\text{R-MON-E} \quad \frac{e \rightarrow e'}{\mu; \pi; \rho; \phi_1[l_1, l_2, v, D, e]\phi_2 \hookrightarrow \mu; \pi; \rho; \phi_1[l_1, l_2, v, D, e']\phi_2}$$

$$\text{R-MON-V-NEXT} \quad \frac{}{\mu; \pi; \rho; \phi_1[l, v, d :: D, a]\phi_2 \hookrightarrow \mu; \pi; \rho; \phi_1[l, v, D, (v \ d) \ a]\phi_2}$$

$$\text{R-MON-V-FIN} \quad \frac{\text{getMonUpdates}(l, \perp \boxtimes a, \pi) = \phi' \quad \text{updateHom}(l, \mu(l) \boxtimes a, \rho) = \rho'}{\mu; \pi; \rho; \phi_1[l, v, \emptyset, a]\phi_2 \hookrightarrow \mu[l \mapsto a]; \pi; \rho'; \phi_1\phi_2\phi'}$$

$$\text{R-HOM-E} \quad \frac{e \rightarrow e'}{\mu; \pi; \rho_1[l_1, l_2, v, D, e]\rho_2; \phi \hookrightarrow \mu; \pi; \rho_1[l_1, l_2, v, D, e']\rho_2; \phi}$$

$$\text{R-HOM-V-NEXT} \quad \frac{\text{getMonUpdates}(l_2, \perp \boxtimes (\mu(l_2) \sqcup a), \pi) = \phi' \quad \text{updateHom}(l_2, \mu(l_2) \boxtimes a, \rho_1[l_1, l_2, v, D, v \ d]\rho_2) = \rho'}{\mu; \pi; \rho_1[l_1, l_2, v, d :: D, a]\rho_2; \phi \hookrightarrow \mu[l_2 \mapsto \mu(l_2) \sqcup a]; \pi; \rho'; \phi}$$

$$\text{R-HOM-V-FIN} \quad \frac{\text{getMonUpdates}(l_2, \perp \boxtimes (\mu(l_2) \sqcup a), \pi) = \phi' \quad \text{updateHom}(l_2, \mu(l_2) \boxtimes a, \rho_1[l_1, l_2, v, \emptyset, \perp]\rho_2) = \rho' \quad a \neq \perp}{\mu; \pi; \rho_1[l_1, l_2, v, \emptyset, a]\rho_2; \phi \hookrightarrow \mu[l_2 \mapsto \mu(l_2) \sqcup a]; \pi; \rho'; \phi\phi'}$$

Compute a single delta from the pending deltas in for each homomorphic link

Once a single delta is computed, join the value with the current value in the store and propagate deltas to out-links