

Seminaïve Evaluation for a Higher-Order Functional Language

ANONYMOUS AUTHOR(S)

One of the workhorse techniques for implementing bottom-up Datalog engines is seminaïve evaluation [Bancilhon 1986]. This optimization improves the performance of Datalog's most distinctive feature: recursively defined predicates. These are computed iteratively, and under a naïve evaluation strategy, each iteration recomputes all previous values. Seminaïve evaluation computes a safe approximation of the *difference* between iterations. This can *asymptotically* improve the performance of Datalog queries.

Seminaïve evaluation is defined partly as a program transformation and partly as a modified iteration strategy, and takes advantage of the first-order nature of Datalog code. This paper extends the seminaïve transformation to higher-order programs written in the Datafun language, which extends Datalog with features like first-class relations, higher-order functions, and datatypes like sum types.

Additional Key Words and Phrases: Datafun, Datalog, functional languages, seminaïve evaluation, incremental computation

1 Introduction

Datalog [Ceri et al. 1989], along with the π -calculus and λ -calculus, is one of the jewel languages of theoretical computer science, connecting programming language theory, database theory, and complexity theory. In terms of programming languages, Datalog can be understood as a fully declarative subset of Prolog which is guaranteed to terminate and so can be evaluated in both top-down and bottom-up fashion. In terms of database theory, it is equivalent to the extension of relational algebra with a fixed point operator. In terms of complexity theory, stratified Datalog over ordered databases characterizes polytime computation [Dantsin et al. 2001].

In addition to its theoretical elegance, over the past twenty years Datalog has seen a surprisingly wide array of uses across a variety of practical domains, both in research and in industry. Whaley and Lam [Whaley 2007; Whaley and Lam 2004] implemented pointer analysis algorithms in Datalog, and found that they could reduce their analyses from thousands of lines of C code to *tens* of lines of Datalog code, while retaining competitive performance. The DOOP pointer analysis framework [Smaragdakis and Balatsouras 2015], built using the Soufflé Datalog engine [Jordan et al. 2016], shows that this approach can handle almost all of industrial languages like Java, even analysing features like reflection [Fourtounis and Smaragdakis 2019]. Semmler has developed the Datalog-based .QL language [de Moor et al. 2007; Schäfer and de Moor 2010] for analysing source code (which was used to analyze the code for NASA's Curiosity Mars rover), and LogicBlox has developed the LogiQL [Aref et al. 2015] language for business analytics and retail prediction. The Boom project at Berkeley has developed the Bloom language for distributed programming [Alvaro et al. 2011], and the Datomic cloud database [Hickey et al. 2012] uses Datalog (embedded in Clojure) as its query language. Microsoft's SecPAL language [Becker et al. 2010] uses Datalog as the foundation of its decentralised authorization specification language. In each case, when the problem formulated in Datalog, the specification became directly implementable, while remaining dramatically shorter and clearer than alternatives implemented in more conventional languages.

However, there are two flies in the ointment. First, even though each of these applications is supported by the skeleton of Datalog, they all had to extend it significantly beyond the theoretical core calculus. For example, core Datalog does not even support arithmetic, since its semantics only speaks of finite sets supporting equality of their elements. Moreover, arithmetic is not a trivial extension, since it can greatly complicate the semantics (for example, proving that termination

2018. 2475-1421/2018/1-ART1 \$15.00
<https://doi.org/>

continues to hold). So despite the fact that kernel Datalog has a very clean semantics, its metatheory seemingly needs to be laboriously re-established once again for each extension.

A natural approach to solving this problem is to find a language in which to write the extensions, which preserves the semantic guarantees that Datalog offers. Two such proposals are the Flix language [Madsen et al. 2016] and the Datafun language [Arntzenius and Krishnaswami 2016]. Conveniently for our exposition, these two languages embody two alternative design strategies.

Flix adopts the route of treating Datalog as an embedded domain-specific language [Leijen and Meijer 1999]. That is, Flix is a fairly conventional (albeit well-designed) functional programming language roughly comparable to ML or Haskell, extended with types representing Datalog predicates and programs. The evaluation of a Flix program builds a Datalog program, which is then handed off to a custom Datalog engine (albeit extended to support arbitrary semilattices). This stratification means that (a) standard Datalog implementation techniques can be used mostly off-the-shelf, but that (b) its functional programming side is mostly a very powerful macro system for Datalog. The only way Flix code runs at Datalog execution time is via the definition of new semilattices (which is functional Flix code implementing a semilattice interface), and for this Flix offers program-verification style correctness checking (including SMT-based tooling).

Like Flix, Datafun is a functional programming language, but unlike Flix, its type discipline supports tracking *monotonicity* of functions. Datalog-style recursively defined relations are given via an explicit fixed point operator; monotonicity ensures uniqueness of this fixed point, playing a role similar to Datalog’s stratification condition. Tracking monotonicity permits a much tighter integration between the functional and logic programming styles, but it comes at a cost: many of Datalog’s standard implementation techniques, developed in the context of a first-order logic language, are not obviously applicable in a higher-order functional setting.

Second, actually making Datalog perform well enough to use in practice calls for very sophisticated program analysis and compiler engineering. (This is a familiar experience from the database community, where query planners must encode a startling amount of knowledge to optimize relatively simple SQL queries.) A wide variety of techniques for optimizing Datalog programs have been studied in the literature, such as using binary decision diagrams to represent relations [Whaley 2007], exploiting the structure of well-behaved subsets (e.g., CFL-reachability problems correspond to the “chain program” fragment of Datalog [Afrati and Papadimitriou 1993]), and combining top-down and bottom-up evaluation via the “magic sets” algorithm [Bancilhon et al. 1986].

Today, one of the workhorse techniques for implementing bottom-up Datalog engines is *semi-naïve evaluation* [Bancilhon 1986]. This optimization improves the performance of Datalog’s most distinctive feature: recursively defined predicates. These can be understood as the fixed point of a set-valued function f . The naïve way to compute this is to iterate the sequence $\emptyset, f(\emptyset), f^2(\emptyset), \dots$ until $f^i(\emptyset) = f^{i+1}(\emptyset)$. However, each iteration will recompute all previous values. Semi-naïve evaluation instead computes a safe approximation of the *difference* between iterations. This optimization is critical, as it can asymptotically improve the performance of Datalog queries.

Contributions. The semi-naïve evaluation algorithm is defined partly as a program transformation on sets of Datalog rules, and partly as a modification of the fixed point computation algorithm. The central contribution of this paper is to give an extended version of this transformation which works on higher-order programs written in the Datafun language.

- We reformulate Datafun in terms of a kernel calculus based on the modal logic S4. Instead of giving a calculus with distinct monotonic and discrete function types, as in the original Datafun paper, we make discreteness into a comonad. In addition to regularizing the calculus and slightly improving its expressiveness, the explicit comonadic structure lets us impose a modal constraint on recursion reminiscent of Hoffman’s work on safe recursion [Hofmann 1997].

This brings the semantics of Datafun more closely in line with Datalog’s, and substantially simplifies the program transformation we present.

- We define a program transformation to *incrementalize* well-typed Datafun programs. The translation is a compositional type-and-syntax-directed transformation, and works uniformly at all types including function types. We build on the *change structure* approach to static program incrementalization introduced by Cai et al. [2014], extending it to support sum types, set types, comonads, and (well-founded) fixed points.
- We establish the correctness of our transformation using a novel logical relation which simultaneously defines the changes connecting old and updated programs, as well as the optimized version of both. The fundamental lemma shows that our transformation is semantics-preserving: any closed program of first-order type has the same meaning after optimization.
- We then discuss our implementation of a compiler from Datafun to Haskell, in both naïve and seminaïve form. This lets us empirically demonstrate the asymptotic speedups predicted by the theory. We additionally discuss the (surprisingly modest) set of program optimizations we found helpful for putting the optimization into practice.

2 Datalog and Datafun by Example

2.1 Datalog

Datalog’s syntax is a subset of Prolog’s. Programs are collections of predicate declarations:

```
parent(earendil, elrond).
parent(elrond, arwen).
parent(arwen, eldarion).

ancestor(X, Z) ← parent(X, Z).
ancestor(X, Z) ← ancestor(X, Y), parent(Y, Z).
```

This defines two binary relations, `parent` and `ancestor`. Lowercase terms like `elrond` and `arwen` are symbols *a la* Lisp, and uppercase terms like `X` and `Y` are variables.

The `parent` relation declares a set of ground *facts*; we assert that `earendil` is the parent of `elrond`, `elrond` the parent of `arwen`, and so on. The `ancestor` relation is declared as a pair of *rules*, the first saying `X` is an ancestor of `Z` if `X` is `Z`’s parent, and the second saying `X` is an ancestor of `Z` if there is an intermediate person `Y` such that `X` is `Y`’s ancestor and `Y` is `Z`’s parent.

Semantically, a predicate denotes the set of tuples that satisfy it. Recursive definitions can be interpreted by viewing the whole program as a relation transformer, and taking the least fixed point of that function. Datalog imposes syntactic restrictions which ensure the relation transformer the rules define is monotone, guaranteeing a unique least fixed point.

2.2 Datafun

The idea behind Datafun is that since the semantics of a Datalog program is a monotone set-valued operator, its natural home is the category **Poset** of partial orders and monotone functions. Since **Poset** is bicartesian closed, it can interpret the simply-typed λ -calculus, which gives us a notation for writing monotone and *higher-order* functions. This lets us *abstract* over Datalog rules, something not possible in Datalog itself! In the remainder of this section we reconstruct Datafun hewing closely to this semantic intuition.

148	types	$A, B ::= 1 \mid A \times B \mid A + B \mid A \rightarrow B \mid \square A \mid \{A\}_{eq}$
149	eqtypes	$A, B_{eq} ::= \{A\}_{eq} \mid 1 \mid A_{eq} \times B_{eq} \mid A_{eq} + B_{eq}$
150	semilattices	$L, M ::= \{A\}_{eq} \mid 1 \mid L \times M$
151	finite eqtypes	$A, B_{fin} ::= 1 \mid A_{fin} \times B_{fin} \mid A_{fin} + B_{fin} \mid \{A\}_{fin}$
152	fixtypes	$L, M_{fix} ::= \{A\}_{fin} \mid 1 \mid L_{fix} \times M_{fix}$
153	terms	$e, f, g ::= x \mid \mathbf{x} \mid \lambda x. e \mid e f \mid \langle \rangle \mid \langle e, f \rangle \mid \pi_i e$
154		$\text{in}_i e \mid \mathbf{case} e \mathbf{of} (\text{in}_i x_i \rightarrow f_i)_{i \in \{1,2\}}$
155		$[e] \mid \mathbf{let} [x] = e \mathbf{in} f \mid e = f \mid \mathbf{empty?} e \mid \mathbf{split} e$
156		$\perp \mid e \vee f \mid \{e_i\}_i \mid \mathbf{for} (x \in e) f \mid \mathbf{fix} e$
157		
158		
159		
160		
161		
162		
163		
164		
165		
166		
167		
168		
169		
170		
171		
172		
173		
174		
175		
176		
177		
178		
179		
180		
181		
182		
183		
184		
185		
186		
187		
188		
189		
190		
191		
192		
193		
194		
195		
196		

FIGURE 1. Datafun syntax

Datafun begins as the simply-typed λ -calculus with functions ($\lambda x. e$ and $e f$), sums ($\text{in}_i e$ and $\mathbf{case} e \mathbf{of} \dots$), and products ($\langle e, f \rangle$ and $\pi_i e$). To this, we add a type of finite sets¹ $\{A\}_{eq}$, introduced with set literals $\{e_0, \dots, e_n\}$, and eliminated using Moggi’s monadic bind syntax, $\mathbf{for} (x \in e_1) e_2$, signifying the union over all $x \in e_1$ of e_2 .

As long as all primitives are monotone, every definable function is also monotone. This is necessary for taking fixed points, but may seem too restrictive. There are many essential non-monotone operations, such as equality tests $e = f$. For example, $\{\} = \{\}$ is true, but if the first argument is increased to $\{1\}$ then it becomes false, a *decrease* (in Datafun, $false < true$).

How can we express non-monotone operations while preserving the property that all functions are monotone? We square this circle by introducing the *discreteness* type constructor $\square A$. The elements of $\square A$ are *exactly the same* as the elements of A , but the order on $\square A$ is discrete, $x \leq y : \square A$ iff $x = y$. Monotonicity of a function $\square A \rightarrow B$ is vacuous: $x = y$ always implies $f(x) \leq f(y)$! This lets us encode ordinary, possibly non-monotone, functions $A \rightarrow B$ as monotone functions $\square A \rightarrow B$. Semantically, \square is a comonad, and accordingly the syntax we use for this is a variant of Pfenning and Davies [2001]’s syntax for constructive S4 modal logic. We make discrete terms with the introduction form $[e]$ and eliminate them with a pattern matching eliminator $\mathbf{let} [x] = e \mathbf{in} f$. Discrete variables are colored and italicised x , while monotone variables are uncolored and upright x . Colored terms e are restricted to only refer to discrete variables. So in the equality test $e = f$, the terms e and f must be discrete.

Finally, Datafun includes fixed points, $\mathbf{fix} e$. The *fix* combinator takes a function $\square(L_{fix} \rightarrow L_{fix})$ and returns its least fixed point. We impose two restrictions on the fixed point operator to ensure well-definedness and termination. First, we require that recursion occur at *semilattice types*, L . A join-semilattice is a partial order with a least element \perp and a least-upper-bound (“join”) operation \vee . Finite sets (with the empty set as least element, and union as join) are an example, as are tuples of semilattices. As long as the lattice has no infinite ascending chains, recursion from the bottom element is guaranteed to find the least fixed point.

Second, we require that the recursive function be boxed, $\square(L_{fix} \rightarrow L_{fix})$. Since boxed expressions can only refer to discrete values, and fixed point functions themselves must be monotone, this has the effect of preventing semantically nested fixed points. We discuss this in more detail in §9. Note

¹To implement set types, their elements must support decidable equality. In our core calculus, we use a subgrammar of “eqtypes”, and in our implementation (which compiles to Haskell) we use typeclass constraints to pick out such types.

$$\begin{array}{l}
197 \quad \mathbf{bool} \xrightarrow{\text{rewrite}} \{1\} \\
198 \quad \mathbf{false} \xrightarrow{\text{rewrite}} \{\} \\
199 \quad \mathbf{true} \xrightarrow{\text{rewrite}} \{\langle \rangle\} \\
200 \quad \mathbf{when} (e) f \xrightarrow{\text{rewrite}} \mathbf{for} (\langle \rangle \in e) f \\
201 \\
202 \quad \{e \mid \varepsilon\} \xrightarrow{\text{rewrite}} \{e\} \\
203 \quad \{e \mid p \in f, \dots\} \xrightarrow{\text{rewrite}} \mathbf{for} (p \in f) \{e \mid \dots\} \\
204 \quad \{e \mid f, \dots\} \xrightarrow{\text{rewrite}} \mathbf{when} (f) \{e \mid \dots\}
\end{array}$$

FIGURE 2. Syntactic sugar

that this does *not* prevent mutual recursion, which can be expressed by taking a fixed point at product type, nor stratified fixed points *à la* Datalog.

2.3 Examples of Datafun Programs

In the examples, we make free use of strings and integers, ordered discretely ($x \leq y$ iff $x = y$). We also make use of some syntax sugar in the definitions:

- (1) First, we make free use of curried functions and pattern matching. Desugaring these is relatively standard, and so we will say little about it, with one exception. Namely, the box-elimination form **let** $[x] = e$ in e' is a pattern matching form, and so we allow it to occur inside of patterns. The effect of a box pattern $[p]$ is to ensure that all of the variables bound in the pattern p are treated as discrete variables.
- (2) We mentioned earlier that Datafun's boolean type **bool** is ordered $\mathbf{false} < \mathbf{true}$. This is because we encode booleans as sets of empty tuples, $\{1\}$, with \mathbf{false} being the empty set $\{\}$ and \mathbf{true} being the singleton $\{\langle \rangle\}$. At semilattice type we also permit a "one-sided" conditional test, **when** $(b) e$, which yields e if b is \mathbf{true} and \perp otherwise. Encoding booleans as sets has the advantage that **when** $(b) e$ is monotone in the condition b .
- (3) Finally, we make use of set comprehension notation. The desugaring we use is based on the translation of comprehensions to the monadic operators [Wadler 1992].

We summarize (except for pattern matching) the sugaring rules we use in figure 2.

2.3.1 Set Operations. Even before higher-order functions, one of the main benefits of Datafun relative to Datalog is that it offers the ability to manipulate relations as first class values. In this subsection we will show how a variety of standard operations on sets can be represented in Datafun.

Membership Tests. The first operation we consider is the membership test operation.

$$\begin{array}{l}
232 \quad \mathit{mem} : \square A \rightarrow \{A\} \rightarrow \mathbf{bool} \\
233 \quad \mathit{mem} [x] ys = \mathbf{for} (y \in ys) x = y \\
234
\end{array}$$

This checks if the input x is equal to any $y \in ys$. The argument x to mem is discrete, because an element is in a set or not – the test is not monotone in x .

Set Intersection. Using mem , it is possible to define set intersection, by taking the union of all the singleton sets $\{x\}$ sets where x is an element of xs also in the set ys .

$$\begin{array}{l}
240 \quad _ \wedge _ : \{A\} \rightarrow \{A\} \rightarrow \{A\} \\
241 \quad xs \wedge ys = \mathbf{for} (x \in xs) \mathbf{when} (\mathit{mem} [x] ys) \{x\} \\
242
\end{array}$$

Using comprehensions, this could alternately be written as:

$$243 \quad xs \wedge ys = \{x \mid x \in xs, \mathit{mem} [x] ys\}$$

246 *Relational Composition.* We can also define the composition of two relations in Datafun.

247 $_ \bullet _ : \{A \times B\} \rightarrow \{B \times C\} \rightarrow \{A \times C\}$
 248 $xs \bullet ys = \mathbf{for} (\langle a, b \rangle \in xs) \mathbf{for} (\langle b', c \rangle \in ys) \mathbf{when} (b = b') \{ \langle a, c \rangle \}$
 249

250 This is basically a transcription of the mathematical definition, where we build those pairs which
 251 agree on their B-typed components. It can also be written using set comprehension as:

252 $_ \bullet _ : \{A \times B\} \rightarrow \{B \times C\} \rightarrow \{A \times C\}$
 253 $xs \bullet ys = \{ \langle a, c \rangle \mid \langle a, b \rangle \in xs, \langle b', c \rangle \in ys, b = b' \}$
 254

255 *Transitive Closure.* Now, we define the transitive closure of a relation.

256 $tc : \square\{A \times A\} \rightarrow \{A \times A\}$
 257 $tc [e] = \mathbf{fix} \ s \ \mathbf{is} \ e \vee (e \bullet s)$
 258

259 This definition uses recursion, just like the mathematical definition – a transitive closure is
 260 the least relation containing the original relation xs and the composition of xs with the transitive
 261 closure. However, one feature of this definition peculiar to core Datafun is that the argument type is
 262 $\square\{A \times A\}$ – the transitive closure takes a *discrete* relation. This is because we must use the relation
 263 within the fixed point, and so its parameter needs to be discrete to occur within.

264 2.3.2 *Combinators for Regular Expression Matching.* Datafun permits tightly integrating the higher-order
 265 functional and bottom-up logic programming styles. In this section, we illustrate the benefits of
 266 doing so by showing how to implement a regular expression matching library in combinator parsing
 267 style. Like combinator parsers in functional languages, the code is very concise. However, support
 268 for the relational style ensures we can write naïve code *without* the exponential backtracking cliffs
 269 typical of parser combinators in functional languages.

270 We assume the existence of a function $pos : \mathbf{string} \rightarrow \{\mathbf{int}\}$ which takes a string and returns the
 271 set of valid indices in that string, and assume that string indexing is written $s[n]$, as in Java or C.
 272 Having done this, we define the type of regular expression matchers.

273 $\mathbf{type\ re} = \square\mathbf{string} \rightarrow \{\mathbf{int} \times \mathbf{int}\}$
 274

275 A regular expression takes a string (boxed so that it can be used inside fixed point expressions),
 276 and returns a set of pairs of integers. The idea is that if the regular expression matcher is passed
 277 the string argument s , then if (i, n) is one of the returned values, the substring $s_i, s_{i+1}, \dots, s_{n-1}$
 278 is matched by the regular expression. That is, it is inclusive on the left and exclusive on the right.

279 $sym : \mathbf{char} \rightarrow \mathbf{re}$
 280 $sym \ c \ [s] = \mathbf{for} \ (n \in pos \ s) \ \mathbf{when} \ (s[n] = c) \{ (n, n + 1) \}$
 281

282 The sym combinator takes a character and returns a set of substrings by returning the set $(n, n + 1)$
 283 where the n -th element of the string s is the character c .

284 $nil : \mathbf{re}$
 285 $nil \ [s] = \mathbf{for} \ (n \in pos \ s) \{ (n, n) \}$
 286

287 The call $nil \ [s]$ yields (n, n) for each position n in s , since an empty substring can start anywhere.

288 $seq : \mathbf{re} \rightarrow \mathbf{re} \rightarrow \mathbf{re}$
 289 $seq \ r_1 \ r_2 \ s = r_1 \ s \bullet r_2 \ s$
 290

291 The seq combinator takes two regular expressions r_1 and r_2 as arguments, applies its argument to
 292 both, and takes the relational composition of the results. Therefore, if the range (i, j) was in the
 293 result of r_1 , and (j, k) was in the result of r_2 , then we will return (i, k) , just as desired.

```

295  bot : re
296  bot _ = ⊥
297
298  alt : re → re → re
299  alt r1 r2 s = r1 s ∨ r2 s

```

We get the empty set of matches from *bot*, and *alt* unions the matches of its two arguments.

```

301  star : □re → re
302  star [r] [s] = nil [s] ∨ tc [r [s]]
303

```

The most interesting regular expression combinator is the Kleene star operation *star*. It uses *nil* to get the reflexive relation on positions, and then takes the transitive closure of the regular expression it received as an argument using the *tc* operation. This forces its argument to be boxed, since *tc* calculates a fixed point, and only discrete variables can occur inside of fixed point expressions.²

2.3.3 *Combinators for Regular Expression Matching, Take 2.* The combinators in the previous section found *all* matches within a given substring, but often we are not interested in all matches: we only want to know if a string can match starting at a particular location. We can easily refactor the combinators above to work in this style, which illustrates the benefits of tightly integrating functional and relational styles of programming – we can use functions to manage strict input/output divisions, and relations to manage nondeterminism and search.

```

315  type re = □(string × int) → {int}

```

Our new type of combinators takes a string and a starting position, and returns a set of ending positions. In contrast, the earlier type took a string and returned a set of start/end pairs.

```

319  sym : char → re
320  sym c [s, n] = when (s[n] = c) {n + 1}

```

sym c checks if *c* occurs at position *n*, returning *n + 1* if it does, and the empty set otherwise.

```

322  nil : re
323  nil [s, n] = {n}

```

```

325  seq : re → re → re
326  seq r1 r2 [s, n] = for (m ∈ r1 [s, n]) r2 [s, m]
327

```

The *nil* function simply returns the same index *n* it received as an argument, since an empty string matches starting from any position. Sequencing via *seq r₁ r₂* checks first to see the possible ending positions from matching *r₁*, and carries on with *r₂* from there.

```

331  bot : re
332  bot _ = ⊥
333
334  alt : re → re → re
335  alt r1 r2 [s, n] = x = r1 x ∨ r2 x

```

We still get the empty set from *bot*, and *alt* still unions the two sets of end positions.

```

338  star : □re → re
339  star [r] [s, n] = fix self is {n} ∨ for (m ∈ self) r [s, m]

```

²As a technical note, sets of pairs of integers do not form a finite-height lattice, so by typing this program is not an acceptable fixed point expression. However, since the positions in a string form a finite set, semantically it is fine. The original Datafun paper shows how one can define bounded fixed points to handle cases like this, so we will not be scrupulous.

344 As before, the *star* combinator takes a boxed regular expression as an argument, and for the same
 345 reason – we are implementing sequencing with a fixed point. One thing worth noting about this
 346 definition is that it is *left-recursive* – the definition takes the endpoints from the fixed point *self*, and
 347 then continues matching using the argument *r*. This should make clear that this is not just plain
 348 old functional programming – we are genuinely relying upon the fixed point semantics of Datafun.

350 3 From Seminaïve Evaluation to the Incremental λ -Calculus

351 Let's return to our example Datalog program, modified to consider graphs rather than ancestry:

```
352 path(X, Z) ← edge(X, Z).
353 path(X, Z) ← edge(X, Y), path(Y, Z).
```

355 Let's suppose that *edge* denotes a linear graph, $\{(a_1, a_2), (a_2, a_3), \dots, (a_{n-1}, a_n)\}$. Then *path*
 356 should denote its reachability relation, $\{(a_i, a_j) \mid 1 \leq i < j \leq n\}$. How can we compute this
 357 relation? The naïve approach is to begin with nothing in the *path* relation and repeatedly apply its
 358 rules until nothing more is deducible. We can make this precise by explicitly time-indexing our
 359 rules:

$$360 \text{path}_{i+1}(X, Z) \leftarrow \text{edge}(X, Z) \qquad \text{path}_{i+1}(X, Z) \leftarrow \text{path}_i(X, Y) \wedge \text{edge}(Y, Z)$$

362 By omission $\text{path}_0 = \emptyset$. From this it's easy to see that inductively $\text{path}_i \subseteq \text{path}_{i+1}$. Consequently,
 363 at step $i + 1$ we re-deduce every fact known at step i . For example, suppose $\text{path}_i(a_j, a_k)$. Then at
 364 step $i + 1$ we apply the second rule to $\text{edge}(a_{j-1}, a_j)$ and deduce $\text{path}_{i+1}(a_{j-1}, a_k)$. But since we
 365 also have $\text{path}_{i+1}(a_j, a_k)$, at time $i + 2$ we deduce the very same thing again, and again at $i + 3$,
 366 $i + 4$, and so on.

367 Because we append edges one at a time, path_i contains exactly paths of i or fewer edges.
 368 Therefore it takes n steps until we reach our fixed point $\text{path}_{n-1} = \text{path}_n$. Since step i involves
 369 $|\text{path}_i| \in \Theta(i^2)$ deductions, we make $\Theta(n^3)$ deductions in total. This seems wasteful, since there
 370 are only $\Theta(n^2)$ paths in the final result.

371 Seminaïve evaluation avoids this waste by transforming the rules for *path* to find the newly
 372 deducible paths, $d\text{path}_i$, at iteration i , and accumulating these changes to produce a final result:

$$373 \text{dpath}_0(X, Y) \leftarrow \text{edge}(X, Y)$$

$$374 \text{dpath}_{i+1}(X, Z) \leftarrow \text{edge}(X, Y) \wedge \text{dpath}_i(Y, Z)$$

$$375 \text{path}_{i+1}(X, Y) \leftarrow \text{path}_i(X, Y) \vee \text{dpath}_i(X, Y)$$

376 It's easy to show inductively that $d\text{path}_i$ contains only paths *exactly* $i + 1$ edges long. Consequently
 377 $|d\text{path}_i| \in \Theta(n - i)$ and we make $\Theta(n^2)$ deductions overall.³

380 3.1 Seminaïve evaluation as incremental computation

381 Now let's move from Datalog to Datafun. The transitive closure of *edge* is the fixed point of the
 382 monotone function *step* defined by:

$$383 \text{step path} = \text{edge} \cup \{(x, z) \mid (x, y) \in \text{edge}, (y, z) \in \text{path}\}$$

384 The naïve way to compute *step*'s fixed point is to iterate it: start with $\text{path}_0 = \emptyset$ and compute
 385 $\text{path}_{i+1} = \text{step path}_i$ for increasing i until $\text{path}_i = \text{path}_{i+1}$. But since $\text{path}_i \subseteq \text{step path}_i$, each
 386 iteration re-computes every path found by the previous iteration. Following Datalog, we'd prefer to
 387 compute only the *change* between iterations. So consider *step'* defined by:

390 ³Here we must assume the accumulation rule $\text{path}_{i+1}(X, Y) \leftarrow \text{path}_i(X, Y) \vee \text{dpath}_i(X, Y)$ is implemented using an
 391 union operator that is efficient when the sets being unioned are of greatly differing sizes.

$$step' \ dpath = \{(x, z) \mid (x, y) \in edge, (y, z) \in dpath\}$$

Observe that

$$\begin{aligned} & step \ path \cup \ step' \ dpath \\ &= edge \cup \{(x, z) \mid (x, y) \in edge, (y, z) \in path\} \cup \{(x, z) \mid (x, y) \in edge, (y, z) \in dpath\} \\ &= edge \cup \{(x, z) \mid (x, y) \in edge, (y, z) \in path \cup dpath\} \\ &= step \ (path \cup dpath) \end{aligned}$$

In other words, $step'$ tells us how $step$ changes as its input grows. Using this property, we can directly compute the changes $dpath_i$ between our iterations $path_i$:

$$\begin{aligned} dpath_0 &= step \ \emptyset = edge \\ dpath_{i+1} &= step' \ dpath_i = \{(x, z) \mid (x, y) \in edge, (y, z) \in dpath_i\} \\ path_{i+1} &= path_i \cup dpath_i \end{aligned}$$

These exactly mirror the derivative and accumulator rules for $path_i$ and $dpath_i$ we gave earlier.

The problem of semiaïve evaluation for Datafun, then, reduces to the problem of finding functions, like $step'$, which compute the change in a function's output given a change to its input. This is a problem of *incremental computation*, and since Datafun is a functional language, we turn to the *incremental λ -calculus* [Cai et al. 2014; Giarrusso et al. 2019].

3.2 Change structures

To make precise the notion of change, an incremental λ -calculus associates every type A with a *change structure*, consisting of:⁴

- (1) A type ΔA of possible changes to values of type A .
- (2) A relation $dx ::_A x \rightsquigarrow y$ for $dx : \Delta A$ and $x, y : A$, glossed as “ dx changes x into y ”.

Since the iterations of a fixed point grow monotonically, in Datafun we only need *increasing* changes. For example, sets change by gaining new elements:

$$\Delta\{A\}_{eq} = \{A\}_{eq} \qquad dx ::_{\{A\}_{eq}} x \rightsquigarrow x \cup dx$$

Set changes may be the most significant for fixed point purposes, but to handle all of Datafun we need a change structure for every type. For products and sums, for example, the change structure is pointwise:

$$\begin{array}{ccc} \Delta 1 = 1 & \Delta(A \times B) = \Delta A \times \Delta B & \Delta(A + B) = \Delta A + \Delta B \\ \langle \rangle ::_1 \langle \rangle \rightsquigarrow \langle \rangle & \frac{da ::_A a \rightsquigarrow a' \quad db ::_B b \rightsquigarrow b'}{\langle da, db \rangle ::_{A \times B} \langle a, b \rangle \rightsquigarrow \langle a', b' \rangle} & \frac{dx ::_{A_i} x \rightsquigarrow y}{in_i dx ::_{A_1 + A_2} in_i x \rightsquigarrow in_i y} \end{array}$$

Since we only consider increasing changes, and $\Box A$ is ordered discretely, the only “change” permitted is to stay the same. Consequently, no information is necessary to indicate what changed:

$$\Delta(\Box A) = 1 \qquad \langle \rangle ::_{\Box A} x \rightsquigarrow x$$

Finally we come to the most interesting case: functions.

⁴Our notion of change structure differs significantly from that of Cai et al. [2014], although it is similar to the logical relation given in Giarrusso et al. [2019]; we discuss this in §9. Although we do not use change structures *per se* in our proofs, they are an important source of intuition.

$$\Delta(A \rightarrow B) = \Box A \rightarrow \Delta A \rightarrow \Delta B \quad \frac{\text{FNCHANGE} \quad (\forall dx ::_A x \rightsquigarrow y) \text{ df } x \text{ dx} ::_B f x \rightsquigarrow g y}{\text{df} ::_{A \rightarrow B} f \rightsquigarrow g}$$

Observe that a function change df takes two arguments: a base point $x : \Box A$ and a change $dx : \Delta A$. To understand why we need both, consider incrementalizing function application: we wish to know how $f x$ changes as both f and x change. So fix $\text{df} :: f \rightsquigarrow g$ and $dx :: x \rightsquigarrow y$. How do we find the change $f x \rightsquigarrow g y$ that updates both function and argument?

If changes were given pointwise, taking only a base point, we'd stipulate that $\text{df} :: f \rightsquigarrow g$ iff $(\forall x) \text{ df } x :: f x \rightsquigarrow g x$. But this only gets us to $g x$, not $g y$: we've accounted for the change in the function, but not the argument. We can account for both by giving df an additional parameter: not just the base point x but also the change dx to it. Then by inverting **FNCHANGE** we have $\text{df } x \text{ dx} :: f x \rightsquigarrow g y$ as desired.

Note also the mixture of monotonicity and non-monotonicity in the type $\Box A \rightarrow \Delta A \rightarrow \Delta B$. Since our functions are monotone — increasing inputs yield increasing outputs — function *changes* are also monotone on input changes ΔA — a larger increase in the input yields a larger increase in the output. However, there's no reason to expect the change in the output to grow as the *base point* increases — hence the use of \Box .

3.3 Zero-changes, derivatives, and faster fixed points

If $dx ::_A x \rightsquigarrow x$, we call dx a *zero-change* to x . Usually zero-changes are rather boring — for example, a zero change to a set $x : \{A\}_{\text{eq}}$ is any $dx \subseteq x$, and so \emptyset is always a zero change. However, there is one very interesting exception: function zero changes. Suppose $\text{df} ::_{A \rightarrow B} f \rightsquigarrow f$. This implies that

$$dx ::_A x \rightsquigarrow y \implies \text{df } x \text{ dx} ::_B f x \rightsquigarrow f y \quad (1)$$

In other words, df yields the change in the output of f given a change to its input. This is exactly the property of step' that made it useful for seminaïve evaluation — indeed, step' is a zero-change to step , modulo not taking the base point x as an argument:

$$dx ::_{\{A\}_{\text{eq}}} x \rightsquigarrow y \implies \text{step}' dx ::_{\{A\}_{\text{eq}}} \text{step } x \rightsquigarrow \text{step } y$$

Function zero changes are so important we give them a special name: *derivatives*. We now have enough machinery to prove correct a general *seminaïve fixed point strategy*. First, observe that:

LEMMA 3.1. *At every semilattice type L , we have $\Delta L = L$ and $dx ::_L x \rightsquigarrow y \iff (x \vee dx) = y$.*

This holds by a simple induction on semilattice types L . Now, given a monotone map $f : L \rightarrow L$ and its derivative $f' : \Box L \rightarrow L \rightarrow L$, we can find f 's fixed-point as the limit of the sequence x_i defined:

$$\begin{aligned} x_0 &= \perp & x_{i+1} &= x_i \vee dx_i \\ dx_0 &= f \perp & dx_{i+1} &= f' x_i dx_i \end{aligned}$$

Let $\text{semifix}(f, f') = \bigvee_i x_i$. By induction and the derivative property, we have $dx_{i+1} :: x_i \rightsquigarrow f x_i$ and so $x_i = f^i \perp$, and therefore $\text{semifix}(f, f') = \text{fix } f$. Moreover, if L has no infinite ascending chains, we will reach our fixed point $x_i = x_{i+1}$ in a finite number of iterations.

This leads directly to our strategy for seminaïve Datafun. [Cai et al. \[2014\]](#) defines a static transformation *Derive e* which computes the change in e given the change in its free variables; it *incrementalizes e*. Our goal is not to incrementalize Datafun *per se*, but to find fixed points faster. Consequently, we define two mutually recursive transformations: ϕe , which computes e faster by replacing fixed points with calls to *semifix*; and δe , which incrementalizes ϕe so that we can

491		$[\varepsilon] = \varepsilon$
492	contexts $\Gamma ::= \varepsilon \mid \Gamma, H$	$[\Gamma, x : A] = [\Gamma]$
493	hypotheses $H ::= x : A \mid x :: A$	$[\Gamma, x :: A] = [\Gamma], x :: A$
494		
495	$\text{VAR} \quad \frac{x : A \in \Gamma}{\Gamma \vdash x : A}$	$\text{DVAR} \quad \frac{x :: A \in \Gamma}{\Gamma \vdash x : A}$
496	$\text{LAM} \quad \frac{\Gamma, x : A \vdash e : B}{\Gamma \vdash \lambda x. e : A \rightarrow B}$	$\text{APP} \quad \frac{\Gamma \vdash e : A \rightarrow B \quad \Gamma \vdash f : A}{\Gamma \vdash e f : B}$
497		$\text{UNIT} \quad \frac{}{\Gamma \vdash \langle \rangle : 1}$
498		
499	$\text{PAIR} \quad \frac{(\Gamma \vdash e_i : A_i)_i}{\Gamma \vdash \langle e_1, e_2 \rangle : A_1 \times A_2}$	$\text{PRJ} \quad \frac{\Gamma \vdash e : A_1 \times A_2}{\Gamma \vdash \pi_i e : A_i}$
500		$\text{INJ} \quad \frac{\Gamma \vdash e : A_i}{\Gamma \vdash \text{in}_i e : A_1 + A_2}$
501		
502		
503	$\text{CASE} \quad \frac{\Gamma \vdash e : A_1 + A_2 \quad (\Gamma, x_i : A_i \vdash f_i : B)_i}{\Gamma \vdash \text{case } e \text{ of } (\text{in}_i x_i \rightarrow f_i)_i : B}$	$\text{BOXI} \quad \frac{[\Gamma] \vdash e : A}{\Gamma \vdash [e] : \square A}$
504		$\text{BOXE} \quad \frac{\Gamma \vdash e : \square A \quad \Gamma, x :: A \vdash f : B}{\Gamma \vdash \text{let } [x] = e \text{ in } f : B}$
505		
506		
507	$\text{BOT} \quad \frac{}{\Gamma \vdash \perp : \underline{L}_{\text{eq}}}$	$\text{JOIN} \quad \frac{(\Gamma \vdash e_i : \underline{L}_{\text{eq}})_i}{\Gamma \vdash e_1 \vee e_2 : \underline{L}_{\text{eq}}}$
508		$\text{SET} \quad \frac{([\Gamma] \vdash e_i : \underline{A}_{\text{eq}})_i}{\Gamma \vdash \{e_i\}_i : \{\underline{A}_{\text{eq}}\}}$
509		$\text{SETFOR} \quad \frac{\Gamma \vdash e : \{A\} \quad \Gamma, x :: A \vdash f : \underline{L}_{\text{eq}}}{\Gamma \vdash \text{for } (x \in e) f : \underline{L}_{\text{eq}}}$
510		
511	$\text{EQ} \quad \frac{([\Gamma] \vdash e_i : \underline{A}_{\text{eq}})_i}{\Gamma \vdash e_1 = e_2 : \text{bool}}$	$\text{ISEMPTY} \quad \frac{[\Gamma] \vdash e : \{1\}}{\Gamma \vdash \text{empty? } e : 1 + 1}$
512		$\text{SPLIT} \quad \frac{\Gamma \vdash e : \square(A + B)}{\Gamma \vdash \text{split } e : \square A + \square B}$
513		$\text{FIX} \quad \frac{\Gamma \vdash e : \square(\underline{L}_{\text{fix}} \rightarrow \underline{L}_{\text{fix}})}{\Gamma \vdash \text{fix } e : \underline{L}_{\text{fix}}}$
514		

FIGURE 3. Datafun core syntax and typing rules

compute the derivative of fixed point functions. In order to define ϕ and δ and show them correct, however, we first need a fuller account of Datafun’s type system and semantics.

4 Typing and Semantics of Core Datafun

The syntax of core Datafun is given in figure 1 and its typing rules in figure 3. Contexts are lists of hypotheses H ; a hypothesis gives the type of either a monotone variable $x : A$ or a discrete variable $x :: A$. The stripping operation $[\Gamma]$ drops all monotone hypotheses from the context Γ , leaving only the discrete ones. The typing judgement $\Gamma \vdash e : A$ may be glossed as “under hypotheses Γ , the term e has the type A ”.

The VAR and DVAR rules say that both monotone hypotheses $x : A$ and discrete hypotheses $x :: A$ justify ascribing the variable x the type A . The LAM rule is the familiar rule for λ -abstraction. However, note that we introduce the argument variable $x : A$ as a *monotone* hypothesis, not a discrete one. (This is the “right” choice because in **Poset** the exponential object is the poset of monotone functions.) The application rule APP is standard, as are the rules UNIT, PAIR, PRJ, INJ, and CASE. As with LAM, the variables $x_i : A_i$ bound in the case branches f_i are monotone.

BoxI says that $[e]$ has type $\square A$ when e has type A in the stripped context $[\Gamma]$. This restricts e to refer only to discrete variables, ensuring we don’t smuggle any information we must treat monotonically into a discretely-ordered \square expression. The elimination rule BoxE for $(\text{let } [x] = e \text{ in } f)$ allows us to “cash in” a boxed expression $e : \square A$ by binding its result to a discrete variable $x :: A$ in the body f .

540 At this point, our typing rules correspond to standard constructive S4 modal logic [Pfenning and
 541 Davies 2001]. We get to Datafun by adding a handful of domain-specific types and operations.
 542 First, SPLIT rule says that $split\ e : \Box A + \Box B$ whenever $e : \Box(A + B)$, distributing box across sum
 543 types.⁵ The other direction, $\Box A + \Box B \rightarrow \Box(A + B)$, is already derivable, as is the isomorphism
 544 $\Box A \times \Box B \cong \Box(A \times B)$. In the examples we give, we use all these operations implicitly, to propagate
 545 discreteness onto the variables bound by a box pattern – in the pattern $[\langle in_1\ x, in_2\ y \rangle]$, we expect
 546 both the variables x and y to be discrete, which is information we propagate using these coercions.
 547 Semantically, all of these operations are the identity, as we will see in the following subsection.

548 This leaves only the rules for manipulating sets and other semilattices. Recall that sets form a
 549 semilattice in the inclusion order, with the empty set as least element and set union as join, and
 550 also that products of semilattices are semilattices (with order and operations given pointwise). So
 551 we take the *semilattice types* L to be the set type $\{A\}$, as well as units 1 and products of lattice types
 552 $L_1 \times L_2$. Then, the BOT rule says that the term \perp is the least element of any semilattice type L , and
 553 the JOIN rule gives the type of joins, saying that $e_1 \vee e_2$ is of semilattice type L when each e_i has
 554 type L . These constructors work for any semilattice L , but there are rules specialized to sets as well.
 555 The typing rule for eliminating sets, SETFOR, is *almost* the monadic bind. However, we generalize it
 556 by not requiring the term **for** $(x \in e)\ e'$ to eliminate into a set type. Instead, we permit e' to be *any*
 557 semilattice type, not just sets. Similarly, the introduction form SET is specific to set types, saying
 558 that $\{e_i\}_{i \in I}$ has type $\{A\}$ when each of the e_i has type A . Furthermore, each e_i has to typecheck
 559 in the discrete context, since membership relies on the non-monotonic property of equality. This
 560 can be seen in the EQ rule, which checks equality of two terms $e = f$ only when they are discrete,
 561 checking in stripped contexts. Finally, the rule FIX for fixed points *fix* e , takes an endofunction e of
 562 type $\Box(L \rightarrow L)$ and yields an expression of type L . In addition to being of semilattice type, we also
 563 demand here that L is of finite height, ensuring that the recursion will terminate.

564 4.1 Semantics

566 The syntax of core Datafun can be interpreted in **Poset**, the category of partially ordered sets
 567 and monotone functions between them. That is, an object A of **Poset** is a pair (A, \leq_A) consisting
 568 of a set A of elements, and a partial order relation $\leq_A \subseteq A \times A$ which is reflexive, transitive,
 569 and antisymmetric. A morphism $f : A \rightarrow B$ in **Poset** is a function on sets $f : A \rightarrow B$ which is
 570 monotone. That is, for all a and a' such that $a \leq_A a'$ we have $f(a) \leq_B f(a')$. The identity
 571 morphism is the identity function on the underlying sets, and composition of morphisms is just
 572 ordinary function composition. We will typically write composition in diagrammatic or “pipeline”
 573 order: given $f : A \rightarrow B$ and $g : B \rightarrow C$ we write $(f ; g) : A \rightarrow C$.

574 4.1.1 *Cartesian Closed Structure*. The cartesian closed structure of **Poset** is largely the same as in **Set**.

576 *Products*. Given partial orders (A, \leq_A) and (B, \leq_B) , we define the poset $(A, \leq_A) \times (B, \leq_B)$
 577 to be $(A \times B, \leq_{A \times B})$ with the order defined as:

$$578 \quad (a, b) \leq_{A \times B} (a', b') \iff a \leq_A a' \text{ and } b \leq_B b'$$

580 The verification that this forms a partial order is routine. To form the product object, we equip
 581 $A \times B$ with the projection maps $\pi_1 : A \times B \rightarrow A$ and $\pi_2 : A \times B \rightarrow B$, which are defined by the
 582 first and second projection on the underlying sets respectively. Given two maps $f : A \rightarrow B$ and
 583 $g : A \rightarrow C$, the universal property of products is witnessed by $\langle f, g \rangle : A \rightarrow B \times C$, where

$$584 \quad \langle f, g \rangle \triangleq a \mapsto \langle f(a), g(a) \rangle$$

586 ⁵An alternative syntax, pursued in Arntzenius and Krishnaswami [2016], would be to give two rules for **case**, depending on
 587 whether or not the scrutinee could be typechecked in a stripped context.

TYPE AND CONTEXT DENOTATIONS

589				
590	$\llbracket 1 \rrbracket = 1$	$\llbracket A \rightarrow B \rrbracket = \llbracket A \rrbracket \Rightarrow \llbracket B \rrbracket$		
591	$\llbracket \{A_{eq}\} \rrbracket = P\llbracket A_{eq} \rrbracket$	$\llbracket A \times B \rrbracket = \llbracket A \rrbracket \times \llbracket B \rrbracket$		
592				
593	$\llbracket \square A \rrbracket = \square \llbracket A \rrbracket$	$\llbracket A + B \rrbracket = \llbracket A \rrbracket + \llbracket B \rrbracket$		
594	$\llbracket \Gamma \rrbracket = \prod_{H \in \Gamma} \llbracket H \rrbracket$	$\llbracket x : A \rrbracket = \llbracket A \rrbracket$	$\llbracket x :: A \rrbracket = \square \llbracket A \rrbracket$	$\llbracket \Gamma \vdash A \rrbracket = \mathbf{Poset}(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket)$
595				
596				

TERM DENOTATIONS

597				
598	$\llbracket \Gamma \vdash x : A \rrbracket$	$= \pi_x ; \varepsilon$	(for $x :: A \in \Gamma$)	
599	$\llbracket \Gamma \vdash x : A \rrbracket$	$= \pi_x$	(for $x : A \in \Gamma$)	
600	$\llbracket \Gamma \vdash \lambda x. e : A \rightarrow B \rrbracket$	$= \lambda \llbracket \Gamma, x : A \vdash e : B \rrbracket$		
601	$\llbracket \Gamma \vdash f e : B \rrbracket$	$= \langle \llbracket \Gamma \vdash f : A \rightarrow B \rrbracket, \llbracket \Gamma \vdash e : A \rrbracket \rangle ; eval$		
602	$\llbracket \Gamma \vdash \langle e_1, e_2 \rangle : A_1 \times A_2 \rrbracket$	$= \langle \llbracket \Gamma \vdash e_1 : A_1 \rrbracket, \llbracket \Gamma \vdash e_2 : A_2 \rrbracket \rangle$		
603	$\llbracket \Gamma \vdash \pi_i e : A_i \rrbracket$	$= \llbracket \Gamma \vdash e : A_1 \times A_2 \rrbracket ; \pi_i$		
604	$\llbracket \Gamma \vdash [e] : \square A \rrbracket$	$= box_\Gamma(\llbracket \Gamma \rrbracket \vdash e : A \rrbracket)$		
605	$\llbracket \Gamma \vdash \mathbf{let} [x] = e \mathbf{in} f : B \rrbracket$	$= \langle id_\Gamma, \llbracket \Gamma \vdash e : \square A \rrbracket \rangle ; \llbracket \Gamma, x :: A \vdash f : B \rrbracket$		
606	$\llbracket \Gamma \vdash \perp : L \rrbracket$	$= !_\Gamma ; join_0^L$		
607	$\llbracket \Gamma \vdash e \vee f : L \rrbracket$	$= \langle \llbracket \Gamma \vdash e : L \rrbracket, \llbracket \Gamma \vdash f : L \rrbracket \rangle ; join_1^L$		
608	$\llbracket \Gamma \vdash \mathbf{empty?} e : 1 + 1 \rrbracket$	$= box_\Gamma(\llbracket \Gamma \rrbracket \vdash e : \{A\} \rrbracket) ; isEmpty$		
609	$\llbracket \Gamma \vdash \mathbf{split} e : \square A + \square B \rrbracket$	$= \llbracket \Gamma \vdash e : \square(A + B) \rrbracket ; dist_+^\square$		
610	$\llbracket e_1 = e_2 \rrbracket$	$= \langle box_\Gamma(\llbracket \Gamma \rrbracket \vdash e_1 : A \rrbracket), box_\Gamma(\llbracket \Gamma \rrbracket \vdash e_2 : A \rrbracket) \rangle ; eq$		
611	$\llbracket \Gamma \vdash \mathbf{fix} e : L \rrbracket$	$= \llbracket \Gamma \vdash e : \square(L \rightarrow L) \rrbracket ; fix$		
612	$\llbracket \{e_i\}_i \rrbracket$	$= \langle box_\Gamma(\llbracket \Gamma \rrbracket \vdash e_1 : A \rrbracket) ; singleton \rangle_i ; join^L$		
613	$\llbracket \Gamma \vdash \mathbf{for} (x \in e) f : L \rrbracket$	$= \langle id, \llbracket \Gamma \vdash e : \{A\} \rrbracket \rangle ; collect(\llbracket \Gamma, x :: A \vdash f : L \rrbracket)$		
614	$\llbracket \Gamma \vdash \mathbf{in}_i e : A_1 + A_2 \rrbracket$	$= \llbracket \Gamma \vdash e : A_i \rrbracket ; in_i$		
615	$\llbracket \Gamma \vdash \mathbf{case} e \mathbf{of} (\mathbf{in}_i x_i \rightarrow f_i)_i : B \rrbracket$	$= \langle id, \llbracket \Gamma \vdash e : A_1 + A_2 \rrbracket \rangle ; dist_+^\times ; [\llbracket \Gamma, x_i : A_i \vdash f_i : B \rrbracket]_i$		
616				
617				
618				
619				

AUXILLIARY OPERATIONS

620	$dist_+^\times : A \times (B_1 + B_2) \rightarrow (A \times B_1) + (A \times B_2)$	$box_\Gamma : \mathbf{Poset}(\llbracket \Gamma \rrbracket, A) \rightarrow \mathbf{Poset}(\llbracket \Gamma \rrbracket, \square A)$
621		
622	$dist_+^\times = \langle \pi_2 ; [\lambda(\langle \pi_2, \pi_1 \rangle ; in_i)]_i, \pi_1 \rangle ; eval$	$box_\Gamma(f) = \langle \pi_x ; \delta \rangle_{x :: A \in \Gamma} ; dist_\square^\times ; \square(f)$
623	$join_0^L : 1 \rightarrow L$	$join_2^L : L \times L \rightarrow L$
624	$join_0^1 = !_1$	$join_2^1 = !_{1 \times 1}$
625	$join_2^{L_1 \times L_2} = \langle join_0^{L_1}, join_0^{L_2} \rangle$	$join_2^{L_1 \times L_2} = \alpha_{L_1, L_2} ; join_2^{L_1} \times join_2^{L_2}$
626	$join_0^{PA} = join_0^{PA}$	$join_2^{PA} = join_2^{PA}$
627		
628		
629	$\alpha_{L, M} : (L \times M) \times (L \times M) \rightarrow (L \times L) \times (M \times M)$	
630	$\alpha_{L, M} = ((l_1, m_1), (l_2, m_2)) \mapsto ((l_1, l_2), (m_1, m_2))$	
631		
632		
633		
634		
635		
636		
637		

FIGURE 4. Semantics of Datafun

The verification that these maps are monotone and that the relevant diagrams commute is routine.

Exponentials. Given partial orders (A, \leq_A) and (B, \leq_B) , we define the poset $(A, \leq_A) \rightarrow (B, \leq_B)$ to be $(A \rightarrow B, \leq_{A \rightarrow B})$ with the order defined as:

$$f \leq_{A \rightarrow B} g \iff \forall a \in A. f(a) \leq_B g(a)$$

That this gives a partial order is immediate. We form the exponential object by equipping it with the evaluation map $eval_{A,B} \in \mathbf{Poset}((A \rightarrow B) \times A, B)$, defined via evaluation on sets:

$$eval \triangleq (f, x) \mapsto f(x)$$

Given $f : C \times A \rightarrow B$, the ‘‘exponential transpose’’ (i.e., currying operation) is defined as:

$$\lambda(f) : C \rightarrow (A \rightarrow B)$$

$$\lambda(f) \triangleq c \mapsto (\lambda a. f(c, a))$$

Verifying that these maps are monotone and that the relevant diagrams commute is straightforward.

Coproducts. Given partial orders (A, \leq_A) and (B, \leq_B) , we define the poset $(A, \leq_A) + (B, \leq_B)$ to be $(A + B, \leq_{A+B})$ with the order defined by (with the understanding that if no clause matches there is no order relationship):

$$in_1 a \leq_{A+B} in_1 a' \iff a \leq_A a'$$

$$in_2 b \leq_{A+B} in_2 b' \iff b \leq_B b'$$

It is easy to check this yields a partial order. We can construct the coproduct object by equipping this poset with the injections $in_1 : A \rightarrow A + B$ and $in_2 : B \rightarrow A + B$, which are given by the injections on the underlying sets. Given two maps $f : A \rightarrow C$ and $g : B \rightarrow C$, the universal property of coproducts is witnessed by $[f, g] : A + B \rightarrow C$, where

$$[f, g] \triangleq x \mapsto \begin{cases} f(a) & \text{when } x = in_1 a \\ g(b) & \text{when } x = in_2 b \end{cases}$$

It is easy to check the defined maps are monotone and that the relevant diagrams commute.

4.1.2 The Discreteness Comonad. Given a poset (A, \leq_A) we define the discreteness operator $\square(A, \leq_A)$ as $(A, \leq_{\square A})$, where

$$a \leq_{\square A} a' \iff a = a'$$

That is, the discrete order preserves the underlying elements, but reduces the partial order to mere equality. Given a morphism $f : A \rightarrow B$, the functorial action is defined as:

$$\square(f) \triangleq f$$

That is, the action on morphisms just gives back the same underlying function on sets. Checking that this preserves monotonicity is trivial, as is the preservation of identities and composition. This functor forms a comonad, with the extraction $\varepsilon_A : \square A \rightarrow A$ and duplication $\delta_A : \square A \rightarrow \square \square A$ given by identities on the underlying sets:

$$\varepsilon_A : \square A \rightarrow A$$

$$\delta_A : \square A \rightarrow \square \square A$$

$$\varepsilon_A = a \mapsto a$$

$$\delta_A = a \mapsto a$$

Once we check that the monotonicity requirements are satisfied, this makes checking the comonad diagrams easy. It is also immediate that \square is a comonad monoidal with respect to *both* products and coproducts. That is, $\square(A \times B) \simeq \square A \times \square B$, and also $\square(A + B) \simeq \square A + \square B$. In both cases the isomorphism is witnessed in both directions by identity on the underlying elements. We will

687 write $dist_{\times}^{\square}$ to name the map witnessing distributivity of \square over products, and $dist_{+}^{\square}$ to name the
 688 map witnessing distributivity of \square over coproducts.

689
 690 4.1.3 *Sets and Lattices.* Given a poset (A, \leq_A) we define the finite powerset poset $P(A, \leq_A)$ as
 691 (PA, \subseteq) , with subsets of A as elements ordered by subset inclusion. This is obviously a semilattice,
 692 with the least element $join_0^{PA}$ given by the empty set, and binary join $join_2^{PA}$ given by union. If
 693 A itself is finite, then this is also a *complete* semilattice. We can define the following collection of
 694 useful morphisms on sets:

$$\begin{array}{ll}
 695 \quad join_0^{PA} & : 1 \rightarrow P(A) & join_2^{PA} & : P(A) \times P(A) \rightarrow P(A) \\
 696 \quad join_0^{PA} & = \langle \rangle \mapsto \emptyset & join_2^{PA} & : \langle X, Y \rangle \mapsto X \cup Y \\
 697 \\
 698 \quad singleton & : \square A \rightarrow PA & isEmpty & : \square PA \rightarrow 1 + 1 \\
 699 \quad singleton & = a \mapsto \{a\} & isEmpty & = X \mapsto \begin{cases} in_1 \langle \rangle & \text{when } X = \emptyset \\ in_2 \langle \rangle & \text{otherwise} \end{cases}
 \end{array}$$

701 The *singleton* function takes a value and makes a singleton set out of it. The domain must be discrete,
 702 as otherwise the map will not be monotone (sets are ordered by inclusion, and set membership
 703 relies on equality, not the partial order). Similarly, the emptiness test *isEmpty* also takes a discrete
 704 set-valued argument, because otherwise the boolean test would not be monotone.

705 Finally, if L is a complete lattice, and $f : A \times \square B \rightarrow L$, then we can define we can also define
 706 the morphism $collect(f) : A \times PA \rightarrow L$ as follows:

$$707 \quad collect(f) = (a, X) \mapsto \bigvee_{b \in X} f(a, b)$$

708 We will use this to interpret for-comprehensions, since all the definable lattice types (sets, and
 709 products of lattices) in core Datafun are complete lattices. However, it is worth noting that the
 710 discreteness restrictions on *singleton* mean that powersets do not quite form a monad in **Poset**.

711
 712 4.1.4 *Fixed Points.* Given a finite-height semilattice L , we can define a fixed point operation fix :
 713 $\square(L \rightarrow L) \rightarrow L$ as follows:

$$714 \quad fix = f \mapsto \bigvee_{n \in \mathbb{N}} f^n(\perp)$$

715 A routine inductive argument shows this must yield a least fixed point.

716
 717 4.1.5 *Interpretation.* The semantic interpretation (defined over typing derivations) is given in [figure 4](#).
 718 We give the interpretation in combinatory style, and to increase readability, we freely use n-ary
 719 products to elide the book-keeping associated with reassociating binary products. The interpretation
 720 itself mostly follows the usual interpretation for constructive S4 [[Alechina et al. 2001](#)], with what
 721 novelty there is occurring in the interpretation of sets and fixed points. Even there, the semantics is
 722 straightforward, making fairly direct use of the combinators defined above.

723 4.2 Metatheory

724 If we were presenting core Datafun in isolation, the usual thing to do would be to prove the
 725 soundness of syntactic substitution, show that syntactic and semantic substitution agree, and then
 726 establish the equational theory. However, that is not our goal in this paper. We want to prove the
 727 correctness of the seminaïve translation, which we will do with a logical relations argument. Since
 728 we can harvest almost all the properties we need from the logical relation, only a small residue
 729 of metatheory needs to be established manually – indeed, the only thing we need to prove at this
 730 stage is the type-correctness of weakening.

731
 732
 733
 734
 735

$$\begin{array}{c}
\text{EMPTY} \\
\frac{}{\varepsilon \sqsubseteq \varepsilon} \\
\text{CONS} \\
\frac{\Gamma \sqsubseteq \Delta}{\Gamma, H \sqsubseteq \Delta, H} \\
\text{DROP} \\
\frac{\Gamma \sqsubseteq \Delta}{\Gamma \sqsubseteq \Delta, H} \\
\text{Disc} \\
\frac{\Gamma \sqsubseteq \Delta}{\Gamma, x : A \sqsubseteq \Delta, x :: A}
\end{array}$$

FIGURE 5. Weakening relation

$$\begin{array}{ll}
\Phi 1 = 1 & \Delta 1 = 1 \\
\Phi\{A\}_{\text{eq}} = \{\Phi A\}_{\text{eq}} \quad (\text{see lemma 5.1}) & \Delta\{A\}_{\text{eq}} = \{A\}_{\text{eq}} \\
\Phi(\Box A) = \Box(\Phi A \times \Delta \Phi A) & \Delta(\Box A) = 1 \\
\Phi(A \times B) = \Phi A \times \Phi B & \Delta(A \times B) = \Delta A \times \Delta B \\
\Phi(A + B) = \Phi A + \Phi B & \Delta(A + B) = \Delta A + \Delta B \\
\Phi(A \rightarrow B) = \Phi A \rightarrow \Phi B & \Delta(A \rightarrow B) = \Box A \rightarrow \Delta A \rightarrow \Delta B
\end{array}$$

FIGURE 6. Δ and Φ type transformations

We define the weakening relation $\Gamma \sqsubseteq \Delta$ in figure 5. This says that Δ is a weakening of Γ , either because (as usual) it has extra hypotheses (in rule DROP), or (additionally!) because a hypothesis in Γ becomes discrete in Δ (rule DISC). The idea is that making a hypothesis discrete only increases the number of places it can be used.

LEMMA 4.1. *If $\Gamma \vdash e : A$ and $\Gamma \sqsubseteq \Delta$ then $\Delta \vdash e : A$.*

This follows by the usual induction on typing derivations.

5 The ϕ and δ Transformations

We use two static transformations, ϕ and δ , defined in figures 7 and 8 respectively. Rather than dive into the gory details immediately, we first build some intuition.

The speed-up transform ϕe computes fixed points seminaïvely by replacing *fix* f by *semifix* (f, f'). But to find the derivative f' of f we'll need a second transform, called δe . Since a derivative is a zero change, can δe simply find a zero change to e ? Unfortunately, this is not strong enough. For example, the derivative of $\lambda x. e$ depends on how e changes as its free variable x changes — which is not necessarily a zero change. To compute derivatives, we need to solve the general problem of computing *changes*. So, modelled on the incremental λ -calculus' *Derive* [Cai et al. 2014], δe will compute how ϕe changes as its free variables change.

However, to speed up *fix* e we don't want the change to e ; we want its derivative. Since derivatives are zero-changes, function changes and derivatives coincide if *the function cannot change*. This is why the typing rule for *fix* e requires that $e : \Box(L_{\text{fix}} \rightarrow L_{\text{fix}})$: the use of \Box prevents e from changing! So the key strategy of our speed-up transformation is to **decorate expressions of type $\Box A$ with their zero-changes**. This makes derivatives available exactly where we need them: at *fix* expressions.

5.1 Typing ϕ and δ

In order to decorate expressions with extra information, ϕ also needs to decorate their types. In figure 6 we give a type translation ΦA capturing this. In particular, if $e : \Box A$ then ϕe will have type $\Phi(\Box A) = \Box(\Phi A \times \Delta \Phi A)$. The idea is that evaluating ϕe will produce a pair $[(x, dx)]$ where $x : \Phi A$ is the sped-up result and $dx : \Delta \Phi A$ is a zero-change to x . Thus, if $e : \Box(L_{\text{fix}} \rightarrow L_{\text{fix}})$, then ϕe will compute $[(f, f')]$, where f' is the derivative of f .

On types other than $\Box A$, there is no information we need to add, so Φ simply distributes. In particular, source programs and sped-up programs agree on the shape of first-order data:

LEMMA 5.1. $\Phi A_{\text{eq}} = A_{\text{eq}}$.

PROOF. Induct on A_{eq} . □

As we'll see in §5.3 and 5.4, ϕ and δ are mutually recursive. To make this work, δe must find the change to ϕe rather than e . So if $e : A$ then $\phi e : \Phi A$ and $\delta e : \Delta \Phi A$. However, so far we have neglected to say what ϕ and δ do to typing contexts. To understand this, it's helpful to look at what Φ and $\Delta \Phi$ do to functions and to \Box . This is because expressions denote functions of their free variables. Moreover, in Datafun free variables come in two flavors, monotone and discrete, and discrete variables are semantically \Box -ed.

If we view expressions as functions of their free variables, δe will denote the *derivative* of the function ϕe denotes. And just as the derivative of a unary function $f x$ has *two* arguments, $df x dx$, the derivative of an expression e with n variables x_1, \dots, x_n will have $2n$ variables: the original x_1, \dots, x_n and their changes dx_1, \dots, dx_n .⁶ However, this says nothing yet about monotonicity or discreteness. To make this precise, we'll use three context transformations, named according to the analogous type operators \Box , Φ , and Δ :

$$\begin{array}{ll} \Box(x : A) = x :: A & \Box(x :: A) = x :: A \\ \Phi(x : A) = x : \Phi A & \Phi(x :: A) = x :: \Phi A, dx :: \Delta \Phi A \\ \Delta(x : A) = dx : \Delta A & \Delta(x :: A) = \varepsilon \quad (\text{the empty context}) \end{array}$$

(Otherwise all three operators distribute; e.g. $\Box \varepsilon = \varepsilon$ and $\Box(\Gamma_1, \Gamma_2) = \Box \Gamma_1, \Box \Gamma_2$.)

Intuitively, $\Box \Gamma$, $\Phi \Gamma$, and $\Delta \Gamma$ mirror the effect of \Box , Φ , and Δ on the semantics of Γ :

$$\begin{array}{lll} \llbracket \Box \Gamma \rrbracket \cong \Box \llbracket \Gamma \rrbracket & \llbracket \Phi(x : A) \rrbracket \cong \llbracket \Phi A \rrbracket & \llbracket \Delta(x : A) \rrbracket \cong \llbracket \Delta A \rrbracket \\ & \llbracket \Phi(x :: A) \rrbracket \cong \llbracket \Phi \Box A \rrbracket & \llbracket \Delta(x :: A) \rrbracket \cong \llbracket \Delta \Box A \rrbracket \end{array}$$

These defined, we can state the types of ϕe and δe :

THEOREM 5.2 (WELL-TYPEDNESS). *If $\Gamma \vdash e : A$, then*

$$\begin{array}{l} \Phi \Gamma \vdash \phi e : \Phi A \\ \Box \Phi \Gamma, \Delta \Phi \Gamma \vdash \delta e : \Delta \Phi A \end{array}$$

As expected if we view expressions as functions of their free variables, if we pretend Γ is a type, these correspond to $\Phi(\Gamma \rightarrow A)$ and $\Delta \Phi(\Gamma \rightarrow A)$ respectively:

$$\Phi(\Gamma \rightarrow A) = \Phi \Gamma \rightarrow \Phi A \quad \Delta \Phi(\Gamma \rightarrow A) = \Box \Phi \Gamma \rightarrow \Delta \Phi \Gamma \rightarrow \Delta \Phi A$$

To get the hang of these context and type transformations, suppose $x :: A, y : B \vdash e : C$. Then theorem 5.2 tells us:

$$\begin{array}{l} x :: \Phi A, dx :: \Delta \Phi A, y : \Phi B \vdash \phi e : \Phi C \\ x :: \Phi A, dx :: \Delta \Phi A, y :: \Phi B, dy : \Delta \Phi B \vdash \delta e : \Delta \Phi C \end{array}$$

Along with the original program's variables, ϕe requires zero change variables dx for every discrete source variable x . Meanwhile, δe requires changes for *every* source program variable (for discrete variables these will be zero changes), and moreover is *discrete* with respect to the source program variables (the "base points").

⁶We assume throughout the paper as a matter of notational convenience that source programs contain no variables starting with the letter d .

834	$\phi x = x$	$\phi x = x$
835	$\phi(\lambda x. e) = \lambda x. \phi e$	$\phi(e f) = \phi e \phi f$
836	$\phi\langle e_i \rangle_i = \langle \phi e_i \rangle_i$	$\phi(\pi_i e) = \pi_i \phi e$
837	$\phi(\text{in}_i e) = \text{in}_i \phi e$	$\phi(\text{case } e \text{ of } (\text{in}_i x \rightarrow f_i)_i) = \text{case } \phi e \text{ of } (\text{in}_i x \rightarrow \phi f_i)_i$
838	$\phi \perp = \perp$	$\phi(e \vee f) = \phi e \vee \phi f$
839	$\phi(\{e_i\}_i) = \{\phi e_i\}_i$	$\phi(\text{for } (x \in e) f) = \text{for } (x \in \phi e) \text{ let } [dx] = [0 x] \text{ in } \phi f$
840	$\phi[e] = [\langle \phi e, \delta e \rangle]$	$\phi(\text{let } [x] = e \text{ in } f) = \text{let } [\langle x, dx \rangle] = \phi e \text{ in } \phi f$
841	$\phi(e = f) = (\phi e = \phi f)$	$\phi(\text{empty? } e) = \text{empty? } \phi e$
842	$\phi(\text{fix } e) = \text{semifix } \phi e$	$\phi(\text{split } e) = \text{case } \phi e \text{ of}$
843		$[\langle \text{in}_i x, \text{in}_i dx \rangle] \rightarrow \text{in}_i [\langle x, dx \rangle]_i$
844		$[\langle \text{in}_i x, \text{in}_j _ \rangle] \rightarrow \text{in}_i [\langle x, \text{dummy } x \rangle]_{i \neq j}$
845		
846		
847		
848		

FIGURE 7. Seminaïve speed-up translation, ϕ

850		
851		
852		
853	$\delta \perp = \delta\{e_i\}_i = \delta(e = f) = \delta(\text{fix } e) = \perp$	
854	$\delta x = dx$	$\delta x = dx$
855	$\delta(\lambda x. e) = \lambda[x]. \lambda dx. \delta e$	$\delta(e f) = \delta e [\phi e] \delta f$
856	$\delta\langle e_i \rangle_i = \langle \delta e_i \rangle_i$	$\delta(\pi_i e) = \pi_i \delta e$
857	$\delta(\text{in}_i e) = \text{in}_i \delta e$	$\delta(e \vee f) = \delta e \vee \delta f$
858	$\delta[e] = \langle \rangle$	$\delta(\text{let } [x] = e \text{ in } f) = \text{let } [\langle x, dx \rangle] = \phi e \text{ in } \delta f$
859	$\delta(\text{empty? } e) = \text{empty? } \phi e$	$\delta(\text{split } e) = \text{case } \phi e \text{ of } [\langle \text{in}_i _, _ \rangle] \rightarrow \text{in}_i \langle \rangle_i$
860		
861	$\delta(\text{case } e \text{ of } (\text{in}_i x \rightarrow f_i)_i) = \text{case split } [\phi e], \delta e \text{ of}$	
862		$(\text{in}_i [x], \text{in}_i dx \rightarrow \delta f_i)_i$
863		$(\text{in}_i [x], \text{in}_j _ \rightarrow \text{let } dx = \text{dummy } x \text{ in } \delta f_i)_{i \neq j}$
864	$\delta(\text{for } (x \in e) f) = (\text{for } (x \in \delta e) \text{ let } [dx] = 0 x \text{ in } \phi f)$	
865		$\vee (\text{for } (x \in \phi e \vee \delta e) \text{ let } [dx] = 0 x \text{ in } \delta f)$
866		
867		
868		
869		

FIGURE 8. Seminaïve derivative translation, δ

We now have enough information to tackle the definitions of ϕ and δ given in figures 7 and 8. In the remainder of this section, we'll examine the most interesting and important parts of these definitions in detail.

5.2 Fixed points

The whole purpose of ϕ and δ is to speed up fixed points, so let's start there. In a fixed point expression $\text{fix } e$, we know $e : \square(\mathbb{L}_{\text{fix}} \rightarrow \mathbb{L}_{\text{fix}})$. Consequently the type of ϕe is

$$\begin{aligned}
883 \quad \Phi(\Box_{\text{fix}}(\mathbb{L} \rightarrow \mathbb{L})) &= \Box(\Phi(\mathbb{L}_{\text{fix}} \rightarrow \mathbb{L}_{\text{fix}}) \times \Delta\Phi(\mathbb{L}_{\text{fix}} \rightarrow \mathbb{L}_{\text{fix}})) \\
884 \quad &= \Box((\Phi\mathbb{L}_{\text{fix}} \rightarrow \Phi\mathbb{L}_{\text{fix}}) \times (\Box\Phi\mathbb{L}_{\text{fix}} \rightarrow \Delta\Phi\mathbb{L}_{\text{fix}} \rightarrow \Delta\Phi\mathbb{L}_{\text{fix}})) \\
885 \quad &= \Box((\mathbb{L}_{\text{fix}} \rightarrow \mathbb{L}_{\text{fix}}) \times (\Box\mathbb{L}_{\text{fix}} \rightarrow \Delta\mathbb{L}_{\text{fix}} \rightarrow \Delta\mathbb{L}_{\text{fix}})) && \text{by lemma 5.1, } \Phi\mathbb{L}_{\text{fix}} = \mathbb{L}_{\text{fix}} \\
886 \quad &= \Box((\mathbb{L}_{\text{fix}} \rightarrow \mathbb{L}_{\text{fix}}) \times (\Box\mathbb{L}_{\text{fix}} \rightarrow \mathbb{L}_{\text{fix}} \rightarrow \mathbb{L}_{\text{fix}})) && \text{by lemma 3.1, } \Delta\mathbb{L}_{\text{fix}} = \mathbb{L}_{\text{fix}}
\end{aligned}$$

888 The behavior of ϕe is to compute a boxed pair $[\langle f, f' \rangle]$, where $f : \mathbb{L}_{\text{fix}} \rightarrow \mathbb{L}_{\text{fix}}$ is a sped-up function
889 and $f' : \Box\mathbb{L}_{\text{fix}} \rightarrow \mathbb{L}_{\text{fix}} \rightarrow \mathbb{L}_{\text{fix}}$ is its derivative. This is exactly what we need to call *semifix*. Therefore
890 $\phi(\text{fix } e) = \text{semifix } \phi e$. However, if we're going to use *semifix* in the output of ϕ , we ought to give it
891 a typing rule and semantics:
892

$$\begin{array}{c}
893 \quad \Gamma \vdash e : \Box((\mathbb{L}_{\text{fix}} \rightarrow \mathbb{L}_{\text{fix}}) \times (\Box\mathbb{L}_{\text{fix}} \rightarrow \mathbb{L}_{\text{fix}} \rightarrow \mathbb{L}_{\text{fix}})) \\
894 \quad \hline \\
895 \quad \Gamma \vdash \text{semifix } e : \mathbb{L}_{\text{fix}}
\end{array}
\quad
\begin{array}{c}
\llbracket \text{semifix } e \rrbracket \gamma = \text{semifix } (f, f') \\
\text{where } (f, f') = \llbracket e \rrbracket \gamma
\end{array}$$

896 As for $\delta(\text{fix } e)$, since e can't change (having \Box type), neither can $\text{fix } e$ (or *semifix* ϕe). All we
897 need is a zero change at type \mathbb{L}_{fix} ; by lemma 3.1, \perp suffices.

898 5.3 Variables, λ , and application

899 At the core of a functional language are variables, λ , and application. The ϕ translation leaves these
900 alone, simply distributing over subexpressions. On variables, δ yields the corresponding change
901 variables. On functions and application, δ is more interesting:
902

$$903 \quad \Delta\Phi(A \rightarrow B) = \Box\Phi A \rightarrow \Delta\Phi A \rightarrow \Delta\Phi B \quad \delta(\lambda x. e) = \lambda[x]. \lambda dx. \delta e \quad \delta(e f) = \delta e [\phi f] \delta f$$

904 The intuition behind $\delta(\lambda x. e) = \lambda[x]. \lambda dx. \delta e$ is that a function change takes two arguments,
905 a base point x and a change dx , and yields the change in the result of the function, δe . However,
906 we are given an argument of type $\Box\Phi A$, but consulting theorem 5.2 for the type of δe , we need a
907 discrete variable $x :: \Phi A$, so we use pattern-matching to unbox our argument.
908

909 The intuition behind $\delta(e f) = \delta e [\phi f] \delta f$ is much the same: δe needs two arguments, the original
910 input ϕf and its change δf , to return the change in the function's output. Moreover, it's discrete in
911 its first argument, so we need to box it, $[\phi f]$.

912 One might wonder why this type-checks, since ϕe and δe don't use the same typing context.
913 We're even boxing ϕf , hiding all monotone variables; consequently, it gets the context $[\Box\Phi\Gamma, \Delta\Phi\Gamma]$.
914 However, \Box makes every variable discrete, and $[-]$ leaves discrete variables alone, so this includes
915 at least $\Box\Phi\Gamma$. The context ϕf needs is $\Phi\Gamma$. Since \Box only makes a context stronger, we're safe. To
916 emphasize this, we've marked all *discrete* uses of ϕe inside δe in pink. The same argument applies
917 (all the more easily) when ϕe is used in a monotone rather than a discrete position.
918

919 5.4 The discreteness comonad, \Box

920 Our strategy hinges on decorating expressions of type $\Box A$ with their zero-changes, so the transla-
921 tions of $[e]$ and $(\text{let } [x] = e \text{ in } f)$ are of particular interest. The most trivial of these is $\delta[e] = \langle \rangle$;
922 this follows from $\Delta\Phi\Box A = 1$, since boxed values cannot change.

923 Next, consider $\phi[e] = [\langle \phi e, \delta e \rangle]$. The intuition here is straightforward: ϕ needs to decorate
924 e with its zero change; since e is discrete and cannot change, we use δe . However! In general,
925 one cannot use δ inside the ϕ translation and expect the result to be well-typed; ϕ and δ require
926 different typing contexts. To see this, let's apply theorem 5.2 to singleton contexts:
927

$$\begin{array}{c}
928 \quad \Gamma \qquad \Phi\Gamma \qquad \Box\Phi\Gamma, \Delta\Phi\Gamma \\
929 \quad \hline \\
930 \quad x : A \quad x : \Phi A \qquad x :: \Phi A, dx : \Delta\Phi A \\
931 \quad x :: A \quad x :: \Phi A, dx :: \Delta\Phi A \qquad x :: \Phi A, dx :: \Delta\Phi A
\end{array}$$

$$\begin{array}{ll}
\text{dummy}_{\{\mathbb{A}\}} _ = \{\} & \text{dummy}_{\mathbb{A} \times \mathbb{B}} \langle x, y \rangle = \langle \text{dummy } x, \text{dummy } y \rangle \\
\text{dummy}_1 \langle \rangle = \langle \rangle & \text{dummy}_{\mathbb{A} + \mathbb{B}} (\text{in}_i x) = \text{in}_i (\text{dummy } x) \\
\text{dummy}_{\square \mathbb{A}} [x] = [\text{dummy } x] & \text{dummy}_{\mathbb{A} \rightarrow \mathbb{B}} f = \lambda x. \text{dummy } (f x)
\end{array}$$

FIGURE 9. The function $\text{dummy}_{\mathbb{A}} : \mathbb{A} \rightarrow \Delta \mathbb{A}$

Luckily, although $\Phi\Gamma$ and $\square\Phi\Gamma$, $\Delta\Phi\Gamma$ differ on monotone variables, they agree on discrete variables. And since e is discrete, there are no monotone variables in e , justifying the use of δe in $\phi[e] = \langle \phi e, \delta e \rangle$.

Next we come to $(\text{let } [x] = e \text{ in } f)$, whose ϕ and δ translations are very similar:

$$\begin{array}{l}
\phi(\text{let } [x] = e \text{ in } f) = \text{let } [\langle x, dx \rangle] = \phi e \text{ in } \phi f \\
\delta(\text{let } [x] = e \text{ in } f) = \text{let } [\langle x, dx \rangle] = \phi e \text{ in } \delta f
\end{array}$$

Since x is a discrete variable, both ϕf and δf need access to its zero change dx . Luckily, $\phi e : \square(\Phi \mathbb{A} \times \Delta \Phi \mathbb{A})$ provides it, so we simply unpack it. We don't use δe in δf , but this is unsurprising when you consider that its type is $\Delta \Phi \square \mathbb{A} = 1$.

5.5 Case analysis, split, and dummy

The derivative of case-analysis, $\delta(\text{case } e \text{ of } (\text{in}_i x_i \rightarrow f_i)_i)$, is complex. Suppose ϕe evaluates to $\text{in}_i x$ and its change δe evaluates to $\text{in}_j dx$. Since δe is a change to ϕe , the change structure on sums tells us that $i = j$! (This is because sums are ordered disjointly; the value x can increase, but the tag in_i must remain the same.) So the desired change $\delta(\text{case } e \text{ of } \dots)$ is given by δf_i in a context supplying a discrete base point x (the value x) and the change dx . To bind x discretely, we need to use $[\phi e] : \square(\Phi \mathbb{A} + \Phi \mathbb{B})$; to pattern-match on this, we need *split* to distribute the \square .

This handles the first two cases, $(\text{in}_i [x], \text{in}_i dx \rightarrow \delta f_i)_i$. Since we know the tags on ϕe and δe agree, these are the only possible cases. However, to appease our type-checker we must handle the *impossible* case that $i \neq j$. This case is dead code: it needs to typecheck, but is otherwise irrelevant. It suffices to generate a dummy change $dx : \Delta \Phi \mathbb{A}_i$ from our base point $x : \Phi \mathbb{A}_i$. We do this using a simple function $\text{dummy}_{\mathbb{A}} : \mathbb{A} \rightarrow \Delta \mathbb{A}$ (figure 9).

We also need *dummy* in the definition of $\phi(\text{split } e)$. In effect $\text{split} : \square(\mathbb{A} + \mathbb{B}) \rightarrow \square \mathbb{A} + \square \mathbb{B}$. Observe that

$$\begin{array}{l}
\Phi(\square(\mathbb{A} + \mathbb{B})) = \square((\Phi \mathbb{A} + \Phi \mathbb{B}) \times (\Delta \Phi \mathbb{A} + \Delta \Phi \mathbb{B})) \\
\Phi(\square \mathbb{A} + \square \mathbb{B}) = \square(\Phi \mathbb{A} \times \Delta \Phi \mathbb{A}) + \square(\Phi \mathbb{B} \times \Delta \Phi \mathbb{B})
\end{array}$$

So while ϕe yields a boxed pair of tagged values, $[\langle \text{in}_i x, \text{in}_j dx \rangle]$, we need $\phi(\text{split } e)$ to yield a tagged boxed pair, $\text{in}_i [\langle x, dx \rangle]$. Again we use *dummy* to handle the impossible case $i \neq j$.

Finally, observe that $\delta(\text{split } e)$ has type $\Delta \Phi(\square \mathbb{A} + \square \mathbb{B}) = \Delta \Phi \square \mathbb{A} + \Delta \Phi \square \mathbb{B} = 1 + 1$. All it must do is return $(\text{in}_i \langle \rangle)$ with a tag that matches $\phi(\text{split } e)$ and ϕe ; **case**-analysing ϕe suffices.

5.6 Semilattices and comprehensions

The translation $\phi(e \vee f) = \phi e \vee \phi f$ is as simple as it seems. However, $\delta(e \vee f) = \delta e \vee \delta f$ is slightly cleverer. In particular, let's restrict to sets, and suppose that dx changes x into x' and dy changes y to y' . In particular, let's suppose these changes are *precise*: that $dx = x' \setminus x$ and $dy = y' \setminus y$. Then the precise change from $x \cup y$ into $x' \cup y'$ is:

$$(x' \cup y') \setminus (x \cup y) = (x' \setminus x \setminus y) \cup (y' \setminus y \setminus x) = (dx \setminus y) \cup (dy \setminus x)$$

981 This suggests letting $\delta(e \cup f) = (\delta e \setminus \phi f) \cup (\delta f \setminus \phi e)$. This is a valid derivative, but it involves
 982 recomputing ϕe and ϕf , and our goal is to avoid recomputation. So instead, we *overapproximate*
 983 the derivative: $\delta e \cup \delta f$ might contain some unnecessary elements, but we expect it to be cheaper
 984 to include these than to recompute ϕe and ϕf . This overapproximation agrees with seminaïve
 985 evaluation in Datalog: Datalog implicitly unions the results of different rules for the same predicate
 986 (e.g. those for *path* in §3), and the seminaïve translations of these rules do not include negated
 987 premises to compute a more precise difference.

988 Now let's consider **for** $(x \in e) f$. Its ϕ -translation is straightforward, with one hitch: because
 989 $x :: A_{\text{eq}}$ is a discrete variable, the inner loop ϕf needs access to its zero change $dx :: \Delta A_{\text{eq}}$. And at
 990 eqtypes (although not in general), the *dummy* function computes zero changes:

991 LEMMA 5.3. *dummy* $x :: A_{\text{eq}} \ x \rightsquigarrow x$ for any $x : A_{\text{eq}}$.

992 For clarity, we write **0** rather than *dummy* when we use it to produce zero changes; we only call it
 993 *dummy* in dead code.

994 Finally, we come to $\delta(\mathbf{for} (x \in e) f)$, the computational heart of the seminaïve transformation,
 995 as **for** is what enables embedding relational algebra (the right-hand-sides of Datalog clauses) into
 996 Datafun. Here there are two things to consider, corresponding to the two **for**-clauses $\delta(\mathbf{for} (x \in e) f)$
 997 generates. First, if the set ϕe we're looping over gains new elements $x \in \delta e$, we need to compute
 998 ϕf over these new elements. Second, if the inner loop ϕf changes, we need to add in its changes
 999 δf for every element, new or old, in the looped-over set, $\phi e \vee \delta e$. Just as in the ϕ -translation, we
 1000 use **0**/*dummy* to calculate zero-changes to set elements.

1003 5.7 Leftovers

1004 The ϕ rules we haven't yet discussed simply distribute ϕ over subexpressions. The remaining δ
 1005 rules mostly do the same, with a few exceptions. In the case of $\delta(\{e_i\}_i) = \delta(e = f) = \perp$, the
 1006 sub-expressions are discrete and cannot change, so we produce a zero-change \perp . This is also the
 1007 case for $\delta(\text{empty? } e) = \text{empty? } \phi e$, but as with $\delta(\text{split } e)$, the zero-change here is at type $1 + 1$, so
 1008 to get the tag right we use ϕe .

1010 6 Proving the Seminaïve Transformation Correct

1011 We formalize the intended behavior of ϕe and δe using a logical relation. Inductively on types A ,
 1012 letting $a, b \in \llbracket A \rrbracket$, $x, y \in \llbracket \Phi A \rrbracket$, and $dx \in \llbracket \Delta \Phi A \rrbracket$, we define $dx :: A \ x \not\downarrow a \rightarrow y \not\downarrow b$, which may be
 1013 glossed as “ x, y speed up a, b respectively, and dx changes x into y ”, as follows:

$$\begin{aligned}
 & \langle \rangle ::_1 \langle \rangle \not\downarrow \langle \rangle \rightarrow \langle \rangle \not\downarrow \langle \rangle \iff \top \\
 & dx ::_{\{A\}} x \not\downarrow a \rightarrow y \not\downarrow b \iff (x, y, x \cup dx) = (a, b, y) \\
 & \langle \rangle ::_{\square A} (x, dx) \not\downarrow a \rightarrow (y, dy) \not\downarrow b \iff (a, x, dx) = (b, y, dy) \wedge dx ::_A x \not\downarrow a \rightarrow y \not\downarrow b \\
 & \vec{dx} ::_{A_1 \times A_2} \vec{x} \not\downarrow \vec{a} \rightarrow \vec{y} \not\downarrow \vec{b} \iff (\forall i) dx_i ::_{A_i} x_i \not\downarrow a_i \rightarrow y_i \not\downarrow b_i \\
 & in_i dx ::_{A_1 + A_2} in_i x \not\downarrow in_k a \rightarrow in_l y \not\downarrow in_m b \iff i = j = k = l = m \wedge dx ::_{A_i} x \not\downarrow a \rightarrow y \not\downarrow b \\
 & df ::_{A \rightarrow B} f \not\downarrow f_s \rightarrow g \not\downarrow g_s \iff (\forall dx ::_A x \not\downarrow a \rightarrow y \not\downarrow b) \\
 & \quad \quad \quad df \ x \ dx ::_B f \ x \not\downarrow f_s \ a \rightarrow g \ y \not\downarrow g_s \ b
 \end{aligned}$$

1024 This extends to contexts Γ , letting $\rho, \rho' \in \llbracket \Gamma \rrbracket$, $\gamma, \gamma' \in \llbracket \Phi \Gamma \rrbracket$, and $d\gamma \in \llbracket \Delta \Phi \Gamma \rrbracket$:

$$\begin{aligned}
 & d\gamma ::_{\Gamma} \gamma \not\downarrow \rho \rightarrow \gamma' \not\downarrow \rho' \iff (\forall x : A \in \Gamma) d\gamma_{dx} ::_A \gamma_x \not\downarrow \rho_x \rightarrow \gamma'_x \not\downarrow \rho'_x \\
 & \quad \quad \quad \wedge (\forall x :: A \in \Gamma) \langle \rangle ::_{\square A} (\gamma_{dx}, \gamma_x) \not\downarrow \rho_x \rightarrow (\gamma'_{dx}, \gamma'_x) \not\downarrow \rho'_x
 \end{aligned}$$

1028 Our fundamental result is that ϕ and δ generate expressions which preserve this logical relation:

1029

THEOREM 6.1 (FUNDAMENTAL). *If $\Gamma \vdash e : A$ and $d\gamma ::_{\Gamma} \gamma \not\leq \rho \rightarrow \gamma' \not\leq \rho'$ then*

$$\llbracket \delta e \rrbracket \langle \gamma, d\gamma \rangle ::_A \llbracket \phi e \rrbracket \gamma \not\leq \llbracket e \rrbracket \rho \rightarrow \llbracket \phi e \rrbracket \gamma' \not\leq \llbracket e \rrbracket \rho'$$

At eqtypes, it's easy to show inductively that $d\alpha ::_{\text{eq}} \alpha \not\leq a \rightarrow \beta \not\leq b$ implies $\alpha = a$. Consequently, first-order closed programs compute the same result when ϕ -translated:

COROLLARY 6.2 (FIRST-ORDER CORRECTNESS). *If $\varepsilon \vdash e : A_{\text{eq}}$ then $\llbracket e \rrbracket = \llbracket \phi e \rrbracket$.*

7 Applying the Seminaïve Transformation to Transitive Closure

Let's try applying the seminaïve transform to a simple Datafun program: the transitive closure function tc from §2.3.1:

$$\begin{aligned} tc \ [e] &= \mathbf{fix} \ p \ \mathbf{is} \ e \cup (e \bullet p) \\ s \bullet t &= \mathbf{for} \ (\langle x, y_1 \rangle \in s) \ \mathbf{for} \ (\langle y_2, z \rangle \in t) \ \mathbf{when} \ (y_1 = y_2) \ \{\langle x, z \rangle\} \end{aligned}$$

In the process we'll discover that besides ϕ itself we need a few simple optimisations to actually speed up our program: most importantly, we need to propagate \perp expressions. In our experience, performing ϕ and δ by hand is easiest when you work inside-out. At the core of transitive closure is a relation composition, $(e \bullet p)$, and at the core of relation composition is a **when**-expression. Let's take a look at its ϕ and δ translations:

$$\begin{aligned} \phi(\mathbf{when} \ (y_1 = y_2) \ \{\langle x, z \rangle\}) &= \phi(\mathbf{for} \ (\langle \rangle \in y_1 = y_2) \ \{\langle x, z \rangle\}) && \text{desugaring} \\ &= \mathbf{for} \ (\langle \rangle \in y_1 = y_2) \ \phi\{\langle x, z \rangle\} && \text{omitting a needless let-binding} \\ &= \mathbf{when} \ (y_1 = y_2) \ \{\langle x, z \rangle\} && \text{resugaring} \end{aligned}$$

Frequently, as in this case, ϕ does nothing interesting. For brevity we'll skip such no-op translations.

$$\begin{aligned} \delta(\mathbf{when} \ (y_1 = y_2) \ \{\langle x, z \rangle\}) &= \delta(\mathbf{for} \ (\langle \rangle \in y_1 = y_2) \ \{\langle x, z \rangle\}) && \text{desugaring when} \\ &= \mathbf{for} \ (\langle \rangle \in \delta(y_1 = y_2)) \ \phi\{\langle x, z \rangle\} && \text{omitting needless let-bindings} \\ &\cup \mathbf{for} \ (\langle \rangle \in \phi(y_1 = y_2) \cup \delta(y_1 = y_2)) \ \delta\{\langle x, z \rangle\} \\ &= \mathbf{for} \ (\langle \rangle \in \perp) \ \{\langle x, z \rangle\} \cup \mathbf{for} \ (\langle \rangle \in \phi(y_1 = y_2) \cup \perp) \ \perp && \text{rules for } \phi(e = f) \text{ and } \delta\{e_i\}_i \\ &= \perp && \text{propagating } \perp \end{aligned}$$

The core insight here is that $y_1 = y_2$ can't change, and neither can $\{\langle x, z \rangle\}$. By propagating this information — for example, rewriting $(\mathbf{for} \ (x \in \perp) \ e)$ to \perp — we can simplify our derivative down to nothing. Now let's pull out and examine $\mathbf{for} \ (\langle y_2, z \rangle \in t) \ \mathbf{when} \ (y_1 = y_2) \ \{\langle x, z \rangle\}$. The ϕ translation is again a no-op.

$$\begin{aligned} \delta(\mathbf{for} \ (\langle y_2, z \rangle \in t) \ \mathbf{when} \ (y_1 = y_2) \ \{\langle x, z \rangle\}) &= \mathbf{for} \ (\langle y_2, z \rangle \in dt) \ \phi(\mathbf{when} \ (y_1 = y_2) \ \{\langle x, z \rangle\}) && \text{omitting needless let-bindings} \\ &\cup \mathbf{for} \ (\langle y_2, z \rangle \in t \cup dt) \ \delta(\mathbf{when} \ (y_1 = y_2) \ \{\langle x, z \rangle\}) \\ &= \mathbf{for} \ (\langle y_2, z \rangle \in dt) \ \mathbf{when} \ (y_1 = y_2) \ \{\langle x, z \rangle\} && \text{propagating } \perp \end{aligned}$$

Tackling the outermost **for** loop:

1079 $\delta(\mathbf{for} (\langle x, y_1 \rangle \in s) \mathbf{for} (\langle y_2, z \rangle \in t) \mathbf{when} (y_1 = y_2) \{\langle x, z \rangle\})$
 1080 $= \mathbf{for} (\langle x, y_1 \rangle \in ds) \phi(\mathbf{for} (\langle y_2, z \rangle \in t) \mathbf{when} (y_1 = y_2) \{\langle x, z \rangle\})$ definition of $\delta(\mathbf{for} \dots)$
 1081 $\cup \mathbf{for} (\langle x, y_1 \rangle \in s \cup ds) \delta(\mathbf{for} (\langle y_2, z \rangle \in t) \mathbf{when} (y_1 = y_2) \{\langle x, z \rangle\})$
 1082 $= \mathbf{for} (\langle x, y_1 \rangle \in ds) \mathbf{for} (\langle y_2, z \rangle \in t) \mathbf{when} (y_1 = y_2) \{\langle x, z \rangle\}$ applying previous work
 1083 $\cup \mathbf{for} (\langle x, y_1 \rangle \in s \cup ds) \mathbf{for} (\langle y_2, z \rangle \in dt) \mathbf{when} (y_1 = y_2) \{\langle x, z \rangle\}$
 1084 $= (ds \bullet t) \cup ((s \cup ds) \bullet dt)$ rewriting in terms of \bullet

1087 This, then, is the derivative $\delta(s \bullet t)$ of relation composition. With a bit of rewriting, this is equivalent
 1088 to $(ds \bullet t) \cup (s \bullet dt) \cup (ds \bullet dt)$, which is perhaps the derivative a human would give.

1089 Let's use this to figure out $\phi(tc [e])$. Working inside out, we start with the derivative of the loop
 1090 body, $\delta(e \cup (e \bullet p))$:

1091 $\delta(e \cup (e \bullet p)) = \delta e \cup \delta(e \bullet p)$
 1092 $= \delta e \cup (\delta e \bullet p) \cup ((e \cup \delta e) \bullet dp)$
 1093 $= \perp \cup (\perp \bullet p) \cup ((e \cup \perp) \bullet dp)$ δe is a zero change
 1094 $= e \bullet dp$ propagate \perp

1097 This requires a new optimization: by definition, $\delta e = de$. However, since e is discrete we know it's
 1098 not changing, and since it's of set type, de may as well be the empty set. So we replace δe with \perp
 1099 instead. Finally, putting everything together:

1100 $\phi(\mathbf{fix} p \text{ is } e \cup (e \bullet p)) = \phi(\mathbf{fix} [\lambda p. e \cup (e \bullet p)])$ desugaring
 1101 $= \mathit{semifix} \phi[\lambda p. e \cup (e \bullet p)]$
 1102 $= \mathit{semifix} [\langle \phi(\lambda p. e \cup (e \bullet p)), \delta(\lambda p. e \cup (e \bullet p)) \rangle]$
 1103 $= \mathit{semifix} [\langle (\lambda p. e \cup (e \bullet p)), (\lambda [p]. \lambda dp. e \bullet dp) \rangle]$ previous work

1106 Examining the recurrence produced by this use of *semifix*, we recover exactly the seminaïve
 1107 transitive closure algorithm we gave in §3.1:

1108 $x_0 = \perp$ $x_{i+1} = x_i \cup dx_i$
 1109 $dx_0 = (\lambda p. e \cup (e \bullet p)) \perp = e$ $dx_{i+1} = (\lambda [p]. \lambda dx. e \bullet dp) [x_i] dx_i = e \bullet dx_i$

1112 8 Implementation and Optimization

1113 To test whether the ϕ translation can produce the asymptotic performance gains we claim, we
 1114 have implemented a compiler from a fragment of Datafun (omitting sum types) to Haskell. We
 1115 use Haskell's `Data.Set` to represent Datafun sets, and typeclasses to implement Datafun's notions
 1116 of equality and semilattice types. We do no query planning; relational joins, written in Datafun
 1117 as nested **for**-loops, are compiled into nested loops. Consequently our performance is worse than
 1118 any real Datalog engine. However, we do implement the ϕ translation, along with the following
 1119 optimizations:

- 1120 • Propagating \perp ; for example, rewriting $(e \vee \perp) \rightsquigarrow e$ and $(\mathbf{for} (x \in e) \perp) \rightsquigarrow \perp$.
- 1121 • Replacing lattice-valued expressions that produce zero changes (for example, changes to
 1122 discrete variables δx) by \perp . This makes \perp -propagation more effective.
- 1123 • Recognising complex zero change expressions; for example, $\delta e \phi f \delta f$ is a zero change if δe
 1124 and δf are. This allows more zero changes to be replaced by \perp , especially in higher-order
 1125 code such as our regular expression example.

1126
1127




FIGURE 10. Naïve and seminaïve transitive closure on a linear graph

We benchmarked the transitive closure function tc from §2.3.1, compiled both naïvely and seminaïvely (i.e. omitting or using ϕ), against the linear graph $\{(i, i + 1) \mid 0 \leq i \leq n\}$. The results (figure 10) are consistent with our expectation that ϕ -translated code, with appropriate optimisation, can perform asymptotically better than naïve evaluation.

9 Discussion and Related Work

Nested fixed points. The typing rule for $fix\ e$ requires $e : \square(\underset{fix}{L} \rightarrow \underset{fix}{L})$. The ϕ translation takes advantage of this \square , decorating expressions of type $\square A$ with their zero-changes. However, it also prevents an otherwise valid idiom: in a nested fixed-point expression $fix\ x\ is\ \dots\ (fix\ y\ is\ e)\ \dots$, the inner fixed point body e cannot use the monotone variable x ! This restriction is not present in Arntzenius and Krishnaswami [2016]; its addition brings Datafun closer to Datalog, whose syntax cannot express this sort of nested fixed point.

We suspect it is possible to lift this restriction without losing seminaïve evaluation, by decorating *all* expressions and variables (not just discrete ones) with zero-changes. However, this also invalidates $\delta(fix\ f) = \perp$: now that f can change, so can $fix\ f$. Luckily, there is a simple and correct solution: $\delta(fix\ f) = fix\ [\delta f\ [fix\ f]]$ [Arntzenius 2017]. However, to compute this new fixed point seminaïvely, we need a *second derivative*: the zero-change to $\delta f\ [fix\ f]$. Indeed, for a program with fixed points nested n deep, we need n^{th} derivatives. We leave this to future work.

Related Work. The incremental lambda calculus was introduced by Cai et al. [2014], as a static program transformation which associated a type of *changes* to each base type, along with operations to update a value based on a change. Then, a program transformation on the simply-typed lambda calculus with base types and functions was defined, which rewrote lambda terms into incremental functions which propagated changes as needed to reduce recomputation. The fundamental idea of the incremental function type taking two arguments (a base point and a change) is one we have built on, though we have extended the transformation to support many more types like sums, sets, modalities, and fixed points. Subsequently, Giarrusso et al. [2019] extended this work to support the *untyped* lambda calculus, additionally also extending the incremental transform to support additional *caching*. In this work, the overall correctness of change propagation was proven using a step-indexed logical relation, which defined which changes were valid in a fashion very similar to our own.

1177 The motivating examples of this line of work was to optimize bulk collection operations, and
 1178 benchmarks showing asymptotic performance improvements were demonstrated. However, all of
 1179 the intuitions were phrased in terms of calculus – a change structure can be thought of as a space
 1180 paired with its tangent space, a zero change on functions is a derivative, and so on. The idea of a
 1181 derivative as a linear approximation is taken most seriously in the work on the differential lambda
 1182 calculus [Ehrhard and Regnier 2003]. These calculi have the beautiful property that the *syntactic*
 1183 linearity in the lambda calculus corresponds to the *semantic* notion of linear transformation.

1184 Unfortunately, the intuition of a derivative has its limits. A function’s derivative is *unique*, a prop-
 1185 erty which models of differential lambda calculi have gone to considerable length to enforce [Blute
 1186 et al. 2006]. This is problematic from the point of view of seminaïve evaluation, since we must
 1187 have the freedom to overapproximate. In §5.6, we took the derivative $\delta(e \vee f)$ to be $\delta(e) \vee \delta(f)$,
 1188 which may overapproximate the change to $e \vee f$. This spares us from having to do expensive
 1189 recomputations to construct set differences, and without this freedom it is doubtful that seminaïve
 1190 evaluation would even be useful at all!

1191 Alvarez-Picallo et al. [2019] offer an alternative formulation of change structures, by requiring
 1192 changes to form a monoid, and representing the change itself with a monoid action. They use
 1193 change actions to prove the correctness of seminaïve evaluation for Datalog, and express the hope
 1194 that it could apply to Datafun. Unfortunately, it does not seem to – the natural notion of function
 1195 change in their setting is pointwise, which does not seem to lead to the derivatives we want in the
 1196 examples we considered.

1197 Overall, there seems to be a lot of freedom in the design space for incremental calculi, and the
 1198 tradeoffs different choices are making remain unclear. Much further investigation is warranted!

1199

1200 References

- 1201 Foto Afrati and Christos H. Papadimitriou. 1993. The Parallel Complexity of Simple Logic Programs. *J. ACM* 40, 4 (Sep
 1202 1993), 891–916. <https://doi.org/10.1145/153724.153752>
- 1203 Natasha Alechina, Michael Mendler, Valeria de Paiva, and Eike Ritter. 2001. Categorical and Kripke Semantics for Constructive
 1204 S4 Modal Logic. In *Computer Science Logic, 15th International Workshop, CSL 2001. 10th Annual Conference of the EACSL,*
 1205 *Paris, France, September 10-13, 2001, Proceedings (Lecture Notes in Computer Science)*, Laurent Fribourg (Ed.), Vol. 2142.
 1206 Springer, 292–307. https://doi.org/10.1007/3-540-44802-0_21
- 1207 Mario Alvarez-Picallo, Alex Eysers-Taylor, Michael Peyton Jones, and C.-H. Luke Ong. 2019. Fixing Incremental Computation -
 1208 Derivatives of Fixpoints, and the Recursive Semantics of Datalog. In *Programming Languages and Systems - 28th European*
 1209 *Symposium on Programming, ESOP 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software,*
 1210 *ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings (Lecture Notes in Computer Science)*, Luís Caires (Ed.),
 1211 Vol. 11423. Springer, 525–552. https://doi.org/10.1007/978-3-030-17184-1_19
- 1212 Peter Alvaro, Neil Conway, Joseph M. Hellerstein, and William R. Marczak. 2011. Consistency Analysis in Bloom: a CALM
 1213 and Collected Approach. In *CIDR 2011, Fifth Biennial Conference on Innovative Data Systems Research, Asilomar, CA, USA,*
 1214 *January 9-12, 2011, Online Proceedings*. 249–260.
- 1215 Molham Aref, Balder ten Cate, Todd J. Green, Benny Kimelfeld, Dan Olteanu, Emir Pasalic, Todd L. Veldhuizen, and
 1216 Geoffrey Washburn. 2015. Design and Implementation of the LogicBlox System. In *Proceedings of the 2015 ACM SIGMOD*
 1217 *International Conference on Management of Data, Melbourne, Victoria, Australia, May 31 - June 4, 2015*. 1371–1382.
- 1218 Michael Arntzenius. 2017. Static differentiation of monotone fixed points. <http://www.rntz.net/files/fixderiv.pdf>. Unpub-
 1219 lished note.
- 1220 Michael Arntzenius and Neelakantan R. Krishnaswami. 2016. Datafun: A Functional Datalog. In *Proceedings of the 21st*
 1221 *ACM SIGPLAN International Conference on Functional Programming (ICFP 2016)*. ACM, New York, NY, USA, 214–227.
 1222 <https://doi.org/10.1145/2951913.2951948>
- 1223 François Bancilhon. 1986. Naive Evaluation of Recursively Defined Relations. In *On Knowledge Base Management Systems:*
 1224 *Integrating Artificial Intelligence and Database Technologies*, Michael L Brodie and John Mylopoulos (Eds.). Springer-Verlag
 1225 New York, Inc., New York, NY, USA, 165–178. <http://dl.acm.org/citation.cfm?id=8789.8804>
- 1226 François Bancilhon, David Maier, Yehoshua Sagiv, and Jeffrey D Ullman. 1986. Magic Sets and Other Strange Ways to
 1227 Implement Logic Programs (Extended Abstract). In *Proceedings of the Fifth ACM SIGACT-SIGMOD Symposium on Principles*
 1228 *of Database Systems (PODS '86)*. ACM, New York, NY, USA, 1–15. <https://doi.org/10.1145/6012.15399>

1229

- 1226 Moritz Y. Becker, Cédric Fournet, and Andrew D. Gordon. 2010. SecPAL: Design and semantics of a decentralized
 1227 authorization language. *Journal of Computer Security* 18, 4 (2010), 619–665.
- 1228 R. F. Blute, J. R. B. Cockett, and R. A. G. Seely. 2006. Differential categories. *Mathematical Structures in Computer Science*
 1229 16, 6 (2006), 1049–1083. <https://doi.org/10.1017/S0960129506005676>
- 1230 Yufei Cai, Paolo G. Giarrusso, Tillmann Rendel, and Klaus Ostermann. 2014. A Theory of Changes for Higher-order
 1231 Languages: Incrementalizing λ -calculi by Static Differentiation. In *Proceedings of the 35th ACM SIGPLAN Conference on*
 1232 *Programming Language Design and Implementation (PLDI '14)*. ACM, New York, NY, USA, 145–155. <https://doi.org/10.1145/2594291.2594304>
- 1233 Stefano Ceri, Georg Gottlob, and Letizia Tanca. 1989. What you Always Wanted to Know About Datalog (And Never Dared
 1234 to Ask). *IEEE Trans. Knowl. Data Eng.* 1, 1 (1989), 146–166.
- 1235 Evgeny Dantsin, Thomas Eiter, Georg Gottlob, and Andrei Voronkov. 2001. Complexity and Expressive Power of Logic
 1236 Programming. *Comput. Surveys* 33, 3 (Sep 2001), 374–425. <https://doi.org/10.1145/502807.502810>
- 1237 Oege de Moor, Damien Sereni, Mathieu Verbaere, Elnar Hajiye, Pavel Avgustinov, Torbjörn Ekman, Neil Ongkingco, and
 1238 Julian Tibble. 2007. λ QL: Object-Oriented Queries Made Easy. In *Generative and Transformational Techniques in Software*
 1239 *Engineering II, International Summer School, GTTSE 2007, Braga, Portugal, July 2-7, 2007. Revised Papers*. 78–133.
- 1240 Thomas Ehrhard and Laurent Regnier. 2003. The differential lambda-calculus. *Theoretical Computer Science* 309, 1 (2003),
 1241 1 – 41. [https://doi.org/10.1016/S0304-3975\(03\)00392-X](https://doi.org/10.1016/S0304-3975(03)00392-X)
- 1242 George Fourtounis and Yannis Smaragdakis. 2019. Deep Static Modeling of invokedynamic. In *ECOOP (LIPIcs)*, Vol. 109.
 1243 Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik.
- 1244 Paolo G. Giarrusso, Yann Régis-Gianas, and Philipp Schuster. 2019. Incremental λ -Calculus in Cache-Transfer Style
 1245 - Static Memoization by Program Transformation. In *ESOP (Lecture Notes in Computer Science)*, Vol. 11423. Springer,
 1246 553–580.
- 1247 Rich Hickey, Stuart Halloway, and Justin Gehrtland. 2012. Datomic: The fully transactional, cloud-ready, distributed database.
 1248 <http://www.datomic.com>. Accessed: 5 July 2019.
- 1249 Martin Hofmann. 1997. A Mixed Modal/Linear Lambda Calculus with Applications to Bellantoni-Cook Safe Recursion. In
 1250 *CSL (Lecture Notes in Computer Science)*, Vol. 1414. Springer, 275–294.
- 1251 Herbert Jordan, Bernhard Scholz, and Pavle Subotic. 2016. Soufflé: On Synthesis of Program Analyzers. In *CAV (2) (Lecture*
 1252 *Notes in Computer Science)*, Vol. 9780. Springer, 422–430.
- 1253 Daan Leijen and Erik Meijer. 1999. Domain Specific Embedded Compilers. In *Proceedings of the 2Nd Conference on Domain-*
 1254 *specific Languages (DSL '99)*. ACM, 109–122. <https://doi.org/10.1145/331960.331977> event-place: Austin, Texas,
 1255 USA.
- 1256 Magnus Madsen, Ming-Ho Yee, and Ondrej Lhoták. 2016. From Datalog to FLIX: A Declarative Language for Fixed Points
 1257 on Lattices. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation,*
 1258 *PLDI 2016, Santa Barbara, CA, USA, June 13-17, 2016*, Chandra Krintz and Emery Berger (Eds.). ACM, 194–208.
 1259 <https://doi.org/10.1145/2908080.2908096>
- 1260 Frank Pfenning and Rowan Davies. 2001. A judgmental reconstruction of modal logic. *Mathematical Structures in Computer*
 1261 *Science* 11, 4 (2001), 511–540. <https://doi.org/10.1017/S0960129501003322>
- 1262 Max Schäfer and Oege de Moor. 2010. Type inference for datalog with complex type hierarchies. In *Proceedings of the 37th*
 1263 *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010, Madrid, Spain, January 17-23,*
 1264 *2010*. 145–156.
- 1265 Yannis Smaragdakis and George Balatsouras. 2015. *Pointer Analysis*. Now Foundations and Trends. [https://ieeexplore.1266
 1267
 1268
 1269
 1270
 1271
 1272
 1273
 1274](https://ieeexplore.ieee.org/document/8186778)
- 1261 Philip Wadler. 1992. Comprehending Monads. *Mathematical Structures in Computer Science* 2, 4 (1992), 461–493.
- 1262 John Whaley. 2007. *Context-Sensitive Pointer Analysis using Binary Decision Diagrams*. Ph.D. Dissertation. Stanford University.
- 1263 John Whaley and Monica S. Lam. 2004. Cloning-based context-sensitive pointer alias analysis using binary decision
 1264 diagrams. In *Proceedings of the ACM SIGPLAN 2004 Conference on Programming Language Design and Implementation*
 1265 *2004, Washington, DC, USA, June 9-11, 2004*, William Pugh and Craig Chambers (Eds.). ACM, 131–144. <https://doi.org/10.1145/996841.996859>