



HAL
open science

Incremental λ -Calculus in Cache-Transfer Style Static Memoization by Program Transformation

Paolo G Giarrusso, Yann Régis-Gianas, Philipp Schuster

► **To cite this version:**

Paolo G Giarrusso, Yann Régis-Gianas, Philipp Schuster. Incremental λ -Calculus in Cache-Transfer Style Static Memoization by Program Transformation. ESOP 2019 - European Symposium on Programming, Apr 2019, Prague, Czech Republic. hal-02405864

HAL Id: hal-02405864

<https://inria.hal.science/hal-02405864>

Submitted on 11 Dec 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Incremental λ -Calculus in Cache-Transfer Style

Static Memoization by Program Transformation

Paolo G. Giarrusso¹, Yann Régis-Gianas², and Philipp Schuster³

¹ LAMP — EPFL

² IRIF, University of Paris Diderot, INRIA

³ University of Tübingen

Abstract. Incremental computation requires propagating changes and reusing intermediate results of base computations. Derivatives, as produced by static differentiation [7], propagate changes but do not reuse intermediate results, leading to wasteful recomputation. As a solution, we introduce conversion to *Cache-Transfer-Style*, an additional program transformations producing purely incremental functional programs that create and maintain nested tuples of intermediate results. To prove CTS conversion correct, we extend the correctness proof of static differentiation from STLC to untyped λ -calculus via *step-indexed logical relations*, and prove sound the additional transformation via simulation theorems. To show ILC-based languages can improve performance relative to from-scratch recomputation, and that CTS conversion can extend its applicability, we perform an initial performance case study. We provide derivatives of primitives for operations on collections and incrementalize selected example programs using those primitives, confirming expected asymptotic speedups.

1 Introduction

After computing a base output from some base input, we often need to produce updated outputs corresponding to updated inputs. Instead of rerunning the same *base program* on the updated input, incremental computation transforms the input change to an output change, potentially reducing asymptotic time complexity and significantly improving efficiency, especially for computations running on large data sets.

Incremental λ -Calculus (ILC) is a recently introduced framework [7] for *higher-order* incremental computation. ILC represents changes from a base value v_1 to an updated value v_2 as a first-class *change value* dv . Since functions are first-class values, change values include *function changes*.

ILC also statically transforms *base programs* to *incremental programs* or *derivatives*, that are functions mapping input changes to output changes. Incremental language designers can then provide their language with (higher-order) primitives (with their derivatives) that efficiently encapsulate incrementalizable computation skeletons (such as tree-shaped folds), and ILC will incrementalize higher-order programs written in terms of these primitives.

Alas, ILC only incrementalizes efficiently *self-maintainable computations* [7, Sec. 4.3], that is, computations whose output changes can be computed using only input changes, but not the inputs themselves [11]. Few computations are self-maintainable: for instance, mapping self-maintainable functions on a sequence is self-maintainable, but dividing numbers is not! Cai et al. [7] already hint at this problem; we elaborate on it in Section 2.1. In this paper, we extend ILC to non-self-maintainable computations. To this end, we must enable derivatives to reuse intermediate results created by the base computation.

Many incrementalization approaches rely on some forms of dynamic memoization to remember intermediate results: they typically use hashables to memoize function results, or dynamic dependence graphs [1] to remember the computation trace. However, looking up intermediate results in such dynamic data structure has a cost in time, and typical general-purpose optimizers cannot predict results from memoization lookups. Besides, reasoning on dynamic dependence graphs and computation traces is often complex. Instead, ILC aims to produce purely functional programs that are suitable for further optimizations and equational reasoning.

To that end, we eschew standard dynamic memoization in favor of *static memoization*: we transform programs to *cache-transfer style (CTS)*, following ideas from Liu and Teitelbaum [20]. CTS functions output *caches* of intermediate results along with their primary results. Caches are just nested tuples whose structure is derived from code, and accessing them does not involve looking up keys depending on inputs. On the contrary, intermediate results can be fetched from these tuples using statically known locations. We also extend differentiation to produce *CTS derivatives*, which can extract from caches any intermediate results they need and which are responsible for updating the caches for the next computation step. This approach was inspired and pioneered by Liu and Teitelbaum [20] for untyped first-order functional languages; we integrate it with ILC and extend it to higher-order languages.

The correctness proof of static differentiation in CTS is quite challenging. First, it requires to show a forward simulation relation between two triples of reduction traces (the first triple being made of the source base evaluation, the source updated evaluation and the source derivative evaluation; the second triple being made of the corresponding CTS-translated evaluations). Dealing with six distinct evaluation environments at the same time was error prone on paper and for this reason, we conducted the proof using Coq [26]. Second, the simulation relation must not only track values but also caches, which are only partially updated while in the middle of the evaluation of derivatives. Finally, we study the translation for an untyped λ -calculus, while previous ILC correctness proofs were restricted to simply-typed λ -calculus. Hence, we define which changes are valid via a *logical relation* and show its *fundamental property*. Being in an untyped setting, our logical relations are *step-indexed*, not indexed by types. We study an untyped language with the intention to make it applicable to the erasure of typed languages. Formalizing a type-preserving translation is left for future work because giving a type to CTS programs is challenging, as we shall explain.

In addition to the correctness proof, we present preliminary experimental results from three case studies. We obtain efficient incremental programs even on non self-maintainable functions.

Our contributions are presented as follows. First, we summarize ILC and illustrate the need to extend it to remember intermediate results via CTS (Sec. 2). Second, in our mechanized formalization (Sec. 3), we give a novel proof of correctness for ILC differentiation for untyped λ -calculus, based on step-indexed logical relations (Sec. 3.4). Third, building on top of ILC differentiation, we show how to transform untyped higher-order programs to CTS (Sec. 3.5) and we show that CTS functions and derivatives *simulate* correctly their non-CTS counterparts (Sec. 3.7). Finally, in our case studies (Sec. 4), we compare the performance of the generated code to the base programs. Sec. 4.4 discusses limitations and future work. Sec. 5 discusses related work and Sec. 6 concludes.

2 ILC and CTS Primer

In this section we exemplify ILC by applying it on an average function, show why the resulting incremental program is asymptotically inefficient, and use CTS conversion and differentiation to incrementalize our example efficiently and speed it up asymptotically (as confirmed by benchmarks in Sec. 4.1). Further examples in Sec. 4 apply CTS to higher-order programs and suggest that CTS enables incrementalizing efficiently some core database primitives such as joins.

2.1 Incrementalizing *average* via ILC

Our example computes the average of a bag of numbers. After computing the *base output* y_1 of the average function on the *base input* bag xs_1 , we want to update the output in response to a stream of updates to the input bag. For simplicity, we assume we have two *updated inputs* xs_2 and xs_3 and want to compute two *updated outputs* y_2 and y_3 , so the overall program can be described in Haskell as:

```

average  :: Bag ℤ → ℤ
average xs = let s = sum xs; n = length xs; r = div s n in r
average3 = let y1 = average xs1; y2 = average xs2; y3 = average xs3 in (y1, y2, y3)

```

Throughout the paper, we contrast *base* vs *updated* inputs and outputs. We also have base and updated values, computations, inputs, outputs, and so on.

We want to compute the updated outputs y_2 and y_3 in *average₃* faster using ILC. For that, we assume that we receive not only updated inputs xs_2 and xs_3 but also *input change* dxs_1 from xs_1 to xs_2 and input change dxs_2 from xs_2 to xs_3 . An input change dx from x_1 to x_2 describes the changes from base input x_1 to updated input x_2 , so that x_2 can be computed via the *update operator* \oplus as $x_1 \oplus dx$.

We use ILC to automatically transform *average* to its derivative *daverage* :: $Bag\ \mathbb{Z} \rightarrow \Delta(Bag\ \mathbb{Z}) \rightarrow \Delta\mathbb{Z}$. A derivative is guaranteed to map input changes to

output changes. In this example, this means that $dy_1 = daverage\ xs_1\ dxs_1$ is a change from base output $y_1 = average\ xs_1$ to updated output $y_2 = average\ xs_2$, hence $y_2 = y_1 \oplus dy_1$.

Thanks to *daverage*'s correctness, we can avoid expensive calls to *average* on updated inputs, and use *daverage* instead. This lets us rewrite *average₃* to *incrementalAverage₃*:

```

incrementalAverage3 :: (Z, Z, Z)
incrementalAverage3 =
  let y1 = average xs1; dy1 = daverage xs1 dxs1
      y2 = y1 ⊕ dy1; dy2 = daverage xs2 dxs2
      y3 = y2 ⊕ dy2
  in (y1, y2, y3)

```

As shown in previous work [10], the derivative df of a function f is the nil change of f since $(f \oplus df)(a \oplus da) = f\ a \oplus df\ a\ da = f(a \oplus da)$

2.2 Self-maintainability and efficiency of derivatives

Alas, derivatives are efficient only if they are *self-maintainable*, and *daverage* is not! So, *incrementalAverage₃* is no faster than *average₃*. Let us inspect the code generated by differentiating *average*.

```

daverage :: Bag Z → Δ(Bag Z) → ΔZ
daverage xs dxs = let s = sum xs; ds = dsum xs dxs;
                    n = length xs; dn = dlength xs dxs;
                    r = div s n; dr = ddiv s ds n dn
                  in dr

```

Since *average* combines *sum*, *length*, and *div*, its derivative *daverage* combines those functions and their derivatives accordingly. It recomputes s , n and r just like *average*, but r is not used anywhere so its recomputation could be avoided. On the other hand, expensive intermediate results s and n are passed to *ddiv*. If *ddiv* does not actually use its inputs s and n , then their computation could be avoided too. Cai et al. [7] call a derivative *self-maintainable* if it does not inspect its base inputs, but only its change inputs. Derivatives produced by ILC are only efficient if they do not compute inputs of self-maintainable derivatives. Cai et al. [7] leave efficient support for non-self-maintainable derivatives for future work. But *ddiv* is *not* self-maintainable! Consider its implementation:

```

ddiv :: Z → ΔZ → Z → ΔZ → ΔZ
ddiv a da b db = div (a ⊕ da) (b ⊕ db) - div a b

```

Function *ddiv* computes the difference between the updated and the original result. It uses its base inputs a and b , hence is not self-maintainable, and a derivative calling it, such as *daverage*, will not be efficient.

But not all is lost: executing *daverage xs dxs* will compute exactly the same s and n as executing *average xs*, so to avoid recomputation we must simply save

and reuse them. Hence, we CTS-convert each function f to a *CTS function* fC and a *CTS derivative* dfC : CTS function fC produces, together with its final result, a *cache* containing intermediate results, that the caller must pass to CTS derivative dfC .

CTS-converting our example produces the following code, which requires no wasteful recomputation.

```

type AverageC = (Z, SumC, Z, LengthC, Z, DivC)
averageC :: Bag Z → (Z, AverageC)
averageC xs =
  let (s, cs1) = sumC xs; (n, cn1) = lengthC xs; (r, cr1) = divC s n
  in (r, (s, cs1, n, cn1, r, cr1))
daverageC :: Bag Z → Δ(Bag Z) → AverageC → (ΔZ, AverageC)
daverageC xs dxs (s, cs1, n, cn1, r, cr1) =
  let (ds, cs2) = dsumC xs dxs cs1
      (dn, cn2) = dlengthC xs dxs cn1
      (dr, cr2) = ddivC s ds n dn cr1
  in (dr, ((s ⊕ ds), cs2, (n ⊕ dn), cn2, (r ⊕ dr), cr2))

```

For each function f , we introduce a type FC for its cache, such that a CTS function fC has type $A \rightarrow (B, FC)$ and CTS derivative dfC has type $A \rightarrow \Delta A \rightarrow FC \rightarrow (\Delta B, FC)$. Crucially, CTS derivatives like $daverageC$ must return an updated cache to ensure correct incrementalization, so that application of further changes works correctly. In general, if $(y_1, c_1) = fC x_1$ and $(dy, c_2) = dfC x_1 dx c_1$, then $(y_1 \oplus dy, c_2)$ must equal the result of the base function fC applied to the updated input $x_1 \oplus dx$, that is $(y_1 \oplus dy, c_2) = fC (x_1 \oplus dx)$.

For CTS-converted functions, the cache type FC is a tuple of intermediate results and caches of subcalls. For primitive functions like div , the cache type $DivC$ could contain information needed for efficient computation of output changes. In the case of div , no additional information is needed. The definition of $divC$ uses div and produces an empty cache, and the definition of $ddivC$ follows the earlier definition for $ddiv$, except that we now pass along an empty cache.

```

data DivC = DivC
divC :: Z → Z → (Z, DivC)
divC a b = (div a b, DivC)
ddivC :: Z → ΔZ → Z → ΔZ → DivC → (ΔZ, DivC)
ddivC a da b db DivC = (div (a ⊕ da) (b ⊕ db) - div a b, DivC)

```

Finally, we can rewrite $average_3$ to incrementally compute y_2 and y_3 :

```

ctsIncrementalAverage3 :: (Z, Z, Z)
ctsIncrementalAverage3 =
  let (y1, c1) = averageC xs1; (dy1, c2) = daverageC xs1 dxs1 c1
      y2 = y1 ⊕ dy1; (dy2, c3) = daverageC xs2 dxs2 c2
      y3 = y2 ⊕ dy2
  in (y1, y2, y3)

```

Since functions of the same type translate to functions of different types, the translation does not preserve well-typedness in a higher-order language in general, but it works well in our case studies (Sec. 4); Sec. 4.1 shows how to map such functions.

3 Formalization

	Terms			Closed values
$a_t ::= \text{let } a_p = a_t \text{ in } a_t$	<i>Let</i>	$a_v ::= a_E[\lambda \overline{a_p}. a_t]$		<i>Closure</i>
a_T	<i>Tuple</i>	$(\overline{a_v})$		<i>Tuple</i>
$f \overline{x}$	<i>Application</i>	ℓ		<i>Literal</i>
	Nested tuples	\mathbf{P}		<i>Primitive</i>
$a_T ::= x$	<i>Variable</i>	$0_{\mathbf{P}}$		<i>Nil change for primitive</i>
$x \oplus dx$	<i>Update</i>	$!a_v$		<i>Replacement change</i>
$(\overline{a_T})$	<i>Tuple</i>			Value environments
	Patterns	$a_E ::= \bullet$		<i>Empty</i>
$a_p ::= x$	<i>Variable</i>	$a_E; x = a_v$		<i>Value binding</i>
$(\overline{a_p})$	<i>Tuple</i>	$j, k, n \in \mathbb{N}$		Step indexes

Fig. 1: Our language λ_L of lambda-lifted programs. Tuples can be nullary.

We now formalize CTS-differentiation for an untyped Turing-complete λ -calculus, and formally prove it sound with respect to differentiation. We also give a novel proof of correctness for differentiation itself, since we cannot simply adapt Cai et al. [7]’s proof to the new syntax: Our language is untyped and Turing-complete, while Cai et al. [7]’s proof assumed a strongly normalizing simply-typed λ -calculus and relied on its naive set-theoretic denotational semantics. Our entire formalization is mechanized using Coq [26]. For reasons of space, some details are deferred to the appendix.

Transformations We introduce and prove sound three term transformations, namely differentiation, CTS translation and CTS differentiation, that take a function to its corresponding (non-CTS) derivative, CTS function and CTS derivative. Each CTS function produces a base output and a cache from a base input, while each CTS derivative produces an output change and an updated cache from an input, an input change and a base cache.

Proof technique To show soundness, we prove that CTS functions and derivatives simulate respectively non-CTS functions and derivatives. In turn, we formalize (non-CTS) differentiation as well, and we prove differentiation sound with respect to non-incremental evaluation. Overall, this shows that CTS functions and derivatives are sound relatively to non-incremental evaluation. Our presentation proceeds in the converse order: first, we present differentiation, formulated as a variant of Cai et al. [7]’s definition; then, we study CTS differentiation.

By using logical relations, we simplify significantly the setup of Cai et al. [7]. To handle an untyped language, we employ *step-indexed* logical relations. Besides, we conduct our development with big-step operational semantics because that choice simplifies the correctness proof for CTS conversion.

Structure of the formalization Sec. 3.1 introduces the syntax of the language λ_L we consider in this development, and introduces its four sublanguages λ_{AL} , λ_{IAL} , λ_{CAL} and λ_{ICAL} . Sec. 3.2 presents the syntax and the semantics of λ_{AL} , the source language for our transformations. Sec. 3.3 defines differentiation and its target language λ_{IAL} , and Sec. 3.4 proves differentiation correct. Sec. 3.5 defines CTS conversion, comprising CTS translation and CTS differentiation, and their target languages λ_{CAL} and λ_{ICAL} . Sec. 3.6 presents the semantics of λ_{CAL} . Finally, Sec. 3.7 proves CTS conversion correct.

Notations We write \overline{X} for a sequence of X of some unspecified length X_1, \dots, X_m .

3.1 Syntax for λ_L

A superlanguage To simplify our transformations, we require input programs to have been lambda-lifted [15] and converted to A'-normal form (A'NF). Lambda-lifted programs are convenient because they allow us to avoid a specific treatment for free variables in transformations. A'NF is a minor variant of ANF [24], where every result is bound to a variable before use; unlike ANF, we also bind the result of the tail call. Thus, every result can thus be stored in a cache by CTS conversion and reused later (as described in Sec. 2). This requirement is not onerous: A'NF is a minimal variant of ANF, and lambda-lifting and ANF conversion are routine in compilers for functional languages. Most examples we show are in this form.

In contrast, our transformation's outputs are lambda-lifted but not in A'NF. For instance, we restrict base functions to take exactly one argument—a base input. As shown in Sec. 2.1, CTS functions take instead two arguments — a base input and a cache — and CTS derivatives take three arguments — an input, an input change, and a cache. We could normalize transformation outputs to inhabit the source language and follow the same invariants, but this would complicate our proofs for little benefit. Hence, we do not *prescribe* transformation outputs to satisfy the same invariants, and we rather *describe* transformation outputs through separate grammars.

As a result of this design choice, we consider languages for base programs, derivatives, CTS programs and CTS derivatives. In our Coq mechanization, we formalize those as four separate languages, saving us many proof steps to check the validity of required structural invariants. For simplicity, in the paper we define a single language called λ_L (for λ -Lifted). This language satisfies invariants common to all these languages (including some of the A'NF invariants). Then, we define λ_L *sublanguages*. We do not formalize the semantics of λ_L itself, preferring to describe it informally; we only formalize the semantics of its sublanguages.

Syntax for terms The λ_L language is a relatively conventional lambda-lifted λ -calculus with a limited form of pattern matching on tuples. The syntax for terms and values is presented in Fig. 1. We separate terms and values in two distinct syntactic classes because we use big-step operational semantics. Our **let**-bindings are non-recursive as usual, and support shadowing. Terms cannot contain λ -expressions directly, but only refer to closures through the environment, and similarly for literals and primitives; we elaborate on this in Sec. 3.2. We do not introduce case expressions, but only bindings that destructure tuples, both in **let**-bindings and λ -expressions of closures. Our semantics does not assign meaning to match failures, but pattern-matchings are only used in generated programs and our correctness proofs ensure that the matches always succeed. We allow tuples to contain terms of form $x \oplus dx$, which update base values x with changes in dx , because A'NF-converting these updates is not necessary to the transformations. We often inspect the result of a function call “ $f x$ ”, which is not a valid term in our syntax. Hence, we write “ $@(f, x)$ ” as a syntactic sugar for “**let** $y = f x$ **in** y ” with y chosen fresh.

Syntax for closed values A closed value is either a closure, a tuple of values, a literal, a primitive, a nil change for a primitive or a replacement change. A closure is a pair of an evaluation environment E and a λ -abstraction closed with respect to E . The set of available literals ℓ is left abstract. It may contain usual first-order literals like integers. We also leave abstract the primitives \mathbf{p} like **if-then-else** or projections of tuple components. Each primitive \mathbf{p} comes with a nil change, which is its derivative as explained in Sec. 2. A change value can also represent a replacement by some closed value a_v . Replacement changes are not produced by static differentiation but are useful for clients of derivatives: we include them in the formalization to make sure that they are not incompatible with our system. As usual, environments E map variables to closed values.

Sublanguages of λ_L The source language for all our transformations is a sublanguage of λ_L named λ_{AL} , where A stands for A'NF. To each transformation we associate a target language, which matches the transformation image. The target language for CTS conversion is named λ_{CAL} , where “C” stands for CTS. The target languages of differentiation and CTS differentiation are called, respectively, λ_{IAL} and λ_{ICAL} , where the “I” stands for incremental.

3.2 The source language λ_{AL}

We show the syntax of λ_{AL} in Fig. 2. As said above, λ_{AL} is a sublanguage of λ_L denoting lambda-lifted base terms in A'NF. With no loss of generality, we assume that all bound variables in λ_{AL} programs and closures are distinct. The step-indexed big-step semantics (Fig. 3) for base terms is defined by the judgment written $E \vdash t \Downarrow_n v$ (where n can be omitted) and pronounced “Under environment E , base term t evaluates to closed value v in n steps.” Intuitively, our

<p>Term differentiation $\boxed{dt = \mathcal{D}^t(t)}$</p> $\mathcal{D}^t(x) = dx$ $\mathcal{D}^t(\mathbf{let} \ y = f \ x \ \mathbf{in} \ t) =$ $\mathbf{let} \ y = f \ x, \ dy = df \ x \ dx \ \mathbf{in} \ \mathcal{D}^t(t)$ $\mathcal{D}^t(\mathbf{let} \ y = (\bar{x}) \ \mathbf{in} \ t) =$ $\mathbf{let} \ y = (\bar{x}), \ dy = (\overline{dx}) \ \mathbf{in} \ \mathcal{D}^t(t)$ <p>Value differentiation $\boxed{dv = \mathcal{D}^t(v)}$</p> $\mathcal{D}^t((\bar{v})) = (\overline{\mathcal{D}^t(v)})$ $\mathcal{D}^t(E_f[\lambda x. t]) = \mathcal{D}^t(E_f)[\lambda x \ dx. \ \mathcal{D}^t(t)]$ $\mathcal{D}^t(\ell) = \mathbf{nil} \ \ell$ $\mathcal{D}^t(\mathbf{p}) = 0_{\mathbf{p}}$ <p>Environment differentiation $\boxed{dE = \mathcal{D}^t(E)}$</p> $\mathcal{D}^t(\bullet) = \bullet$ $\mathcal{D}^t(E; x = v) = \mathcal{D}^t(E); x = v; dx = \mathcal{D}^t(v)$ <p>Base/updated environment $\boxed{E = [dE]_i}$</p> $[\bullet]_i = \bullet \quad i = 1, 2$ $[dE; x = v; dx = dv]_i = [dE]_1; x = v' \quad v' = v \text{ if } i = 1 \text{ or } v' = v \oplus dv \text{ if } i = 2$	<p>λ_{IAL} change terms</p> $dt ::= dx$ $ \ \mathbf{let} \ y = f \ x, \ dy = df \ x \ dx \ \mathbf{in} \ dt$ $ \ \mathbf{let} \ y = (\bar{x}), \ dy = (\overline{dx}) \ \mathbf{in} \ dt$ <p>λ_{IAL} change values</p> $dv ::= (\overline{dv}) \ \ dE[\lambda x \ dx. \ dt] \ \ d\ell \ \ 0_{\mathbf{p}} \ \ !v$ <p>λ_{IAL} change environments</p> $dE ::= \bullet \ \ dE; x = v; dx = dv$ <p>λ_{AL} base terms</p> $t ::= x \ \ \mathbf{let} \ y = f \ x \ \mathbf{in} \ t \ \ \mathbf{let} \ y = (\bar{x}) \ \mathbf{in} \ t$ <p>λ_{AL} closed values</p> $v ::= (\bar{v}) \ \ E[\lambda x. t] \ \ \ell \ \ \mathbf{p}$ <p>λ_{AL} value environments</p> $E ::= \bullet \ \ E; x = v$
---	--

Fig. 2: Static differentiation $\mathcal{D}^t(-)$; syntax of its target language λ_{IAL} , tailored to the output of differentiation; syntax of its source language λ_{AL} . We assume that in λ_{IAL} the same **let** binds both y and dy and that α -renaming preserves this invariant. We also define the *base environment* $[dE]_1$ and the *updated environment* $[dE]_2$ of a change environment dE .

step-indexes count the number of “nodes” of a big-step derivation.⁴ As they are relatively standard, we defer the explanations of these rules to the appendix B.

Expressiveness A closure in the base environment can be used to represent a top-level definition. Since environment entries can point to primitives, we need no syntax to directly represent calls of primitives in the syntax of base terms. To encode in our syntax a program with top-level definitions and a term to be evaluated representing the entry point, one can produce a term t representing the entry point together with an environment E containing as values any top-level definitions, primitives and literals used in the program. Semi-formally, given an environment E_0 mentioning needed primitives and literals, and a list of top-level function definitions $D = \overline{f = \lambda x. t}$ defined in terms of E_0 , we can produce a base environment $E = \mathcal{L}(D)$, with \mathcal{L} defined by:

$$\mathcal{L}(\bullet) = E_0 \text{ and } \mathcal{L}(D, f = \lambda x. t) = E, f = E[\lambda x. t] \text{ where } \mathcal{L}(D) = E$$

⁴ Instead, Ahmed [4] and Acar et al. [3] count the number of steps that small-step evaluation would take, but our choice simplifies some proof steps and makes a minor difference in others.

$$\begin{array}{c}
\text{[S\textsc{VAR}]} \\
\frac{}{E \vdash x \Downarrow_1 E(x)} \\
\text{[S\textsc{TUPLE}]} \\
\frac{E; y = (E(\bar{x})) \vdash t \Downarrow_n v}{E \vdash \mathbf{let } y = (\bar{x}) \mathbf{ in } t \Downarrow_{n+1} v} \\
\text{[S\textsc{PRIMITIVECALL}]} \\
\frac{E(f) = \mathbf{p} \quad E; y = \delta_{\mathbf{p}}(E(x)) \vdash t \Downarrow_n v}{E \vdash \mathbf{let } y = f \mathbf{ x in } t \Downarrow_{n+1} v} \\
\text{[S\textsc{CLOSURECALL}]} \\
\frac{E(f) = E_f[\lambda x. t_f] \quad E_f; x = E(x) \vdash t_f \Downarrow_m v_y \quad E; y = v_y \vdash t \Downarrow_n v}{E \vdash \mathbf{let } y = f \mathbf{ x in } t \Downarrow_{m+n+1} v}
\end{array}$$

Fig. 3: Step-indexed big-step semantics for base terms of source language λ_{AL} .

Correspondingly, we extend all our term transformations to values and environments to transform such encoded top-level definitions.

Our mechanization can encode n -ary functions “ $\lambda(x_1, x_2, \dots, x_n). t$ ” through unary functions that accept tuples; we encode partial application using a **curry** primitive such that, essentially, **curry** $f x y = f(x, y)$; suspended partial applications are represented as closures. This encoding does not support currying efficiently, we further discuss this limitation in Sec. 4.4.

Control operators, like recursion combinators or branching, can be introduced as primitive operations as well. If the branching condition changes, expressing the output change in general requires replacement changes. Similarly to branching we can add tagged unions.

To check the assertions of the last two paragraphs, the Coq development contains the definition of a **curry** primitive as well as a primitive for a fixpoint combinator, allowing general recursion and recursive data structures as well.

3.3 Static differentiation from λ_{AL} to λ_{IAL}

Previous work [7] defines static differentiation for simply-typed λ -calculus terms. Fig. 2 transposes differentiation as a transformation from λ_{AL} to λ_{IAL} and defines λ_{IAL} ’s syntax.

Differentiating a base term t produces a change term $\mathcal{D}^t(t)$, its *derivative*. Differentiating final result variable x produces its change variable dx . Differentiation copies each binding of an intermediate result y to the output and adds a new binding for its change dy . If y is bound to tuple (\bar{x}) , then dy will be bound to the change tuple (\overline{dx}) . If y is bound to function application “ $f x$ ”, then dy will be bound to the application of function change df to input x and its change dx . We explain differentiation of environments $\mathcal{D}^t(E)$ later in this section.

Evaluating $\mathcal{D}^t(t)$ recomputes all intermediate results computed by t . This recomputation will be avoided through cache-transfer style in Sec. 3.5. A comparison with the original static differentiation [7] can be found in Appendix A.

Semantics for λ_{IAL} We move on to define how λ_{IAL} change terms evaluate to change values. We start by defining necessary definitions and operations on

$$\begin{array}{c}
\text{[SDVARIABLE]} \\
\frac{dE \vdash dx \Downarrow_1 dE(dx)}{dE \vdash dx \Downarrow_1 dE(dx)} \\
\\
\text{[SDTUPLE]} \\
\frac{dE(\bar{x}, \overline{dx}) = \bar{v}_x, \overline{dv}_x \quad dE; y = (\bar{v}_x); dy = (\overline{dv}_x) \vdash dt \Downarrow_n dv}{dE \vdash \mathbf{let} y = (\bar{x}), dy = (\overline{dx}) \mathbf{in} dt \Downarrow_{n+1} dv} \\
\\
\text{[SDREPLACECALL]} \\
\frac{\begin{array}{c} [dE]_1 \vdash @ (f, x) \Downarrow_m v_y \quad [dE]_2 \vdash @ (f, x) \Downarrow_n v_y' \\ dE(df) = !v_f \quad dE; y = v_y; dy = !v_y' \vdash dt \Downarrow_p dv \end{array}}{dE \vdash \mathbf{let} y = f x, dy = df x dx \mathbf{in} dt \Downarrow_{m+n+p+1} dv} \\
\\
\text{[SDPRIMITIVE NIL]} \\
\frac{dE(f, df) = \mathbf{p}, \mathbf{0}_{\mathbf{p}} \quad dE(x, dx) = v_x, dv_x \quad dE; y = \delta_{\mathbf{p}}(v_x); dy = \Delta_{\mathbf{p}}(v_x, dv_x) \vdash dt \Downarrow_n dv}{dE \vdash \mathbf{let} y = f x, dy = df x dx \mathbf{in} dt \Downarrow_{n+1} dv} \\
\\
\text{[SDCLOSURECHANGE]} \\
\frac{dE(f, df) = E_f[\lambda x. t_f], dE_f[\lambda x dx. dt_f] \quad dE(x, dx) = v_x, dv_x \quad E_f; x = v_x \vdash t_f \Downarrow_m v_y \quad dE_f; x = v_x; dx = dv_x \vdash dt_f \Downarrow_n dv_y \quad dE; y = v_y; dy = dv_y \vdash dt \Downarrow_p dv}{dE \vdash \mathbf{let} y = f x, dy = df x dx \mathbf{in} dt \Downarrow_{m+n+p+1} dv}
\end{array}$$

Fig. 4: Step-indexed big-step semantics for the change terms of λ_{IAL} .

changes, such as define *change values* dv , *change environments* dE , and the *update operator* \oplus .

Closed change values dv are particular λ_L values a_v . They are either a closure change, a tuple change, a literal change, a replacement change or a primitive nil change. A closure change is a closure containing a change environment dE and a λ -abstraction expecting a value and a change value as arguments to evaluate a change term into an output change value. An evaluation environment dE follows the same structure as **let**-bindings of change terms: it binds variables to closed values and each variable x is immediately followed by a binding for its associated change variable dx . As with **let**-bindings of change terms, α -renamings in an environment dE must rename dx into dy if x is renamed into y . We define the *update operator* \oplus to update a value with a change. This operator is a partial function written “ $v \oplus dv$ ”, defined as follows:

$$\begin{array}{ll}
v_1 \oplus !v_2 & = v_2 \\
\ell \oplus d\ell & = \delta_{\oplus}(\ell, d\ell) \\
E[\lambda x. t] \oplus dE[\lambda x dx. dt] & = (E \oplus dE)[\lambda x. t] \\
(v_1, \dots, v_n) \oplus (dv_1, \dots, dv_n) & = (v_1 \oplus dv_1, \dots, v_n \oplus dv_n) \\
\mathbf{p} \oplus \mathbf{0}_{\mathbf{p}} & = \mathbf{p}
\end{array}$$

where $(E; x = v) \oplus (dE; x = v; dx = dv) = ((E \oplus dE); x = (v \oplus dv))$.

Replacement changes can be used to update all values (literals, tuples, primitives and closures), while tuple changes can only update tuples, literal changes can only update literals, primitive nil can only update primitives and closure changes can only update closures. A replacement change overrides the current value v with a new one v' . On literals, \oplus is defined via some interpretation function δ_{\oplus} , which takes a literal and a literal change to produce an updated literal. Change update for a closure ignores dt instead of computing something like $dE[t \oplus dt]$. This may seem surprising, but we only need \oplus to behave well for valid changes (as shown by Theorem 3.1): for valid closure changes, dt must behave anyway similarly to $\mathcal{D}^c(t)$, which Cai et al. [7] show to be a nil change. Hence, $t \oplus \mathcal{D}^c(t)$ and $t \oplus dt$ both behave like t , so \oplus can ignore dt and only consider environment updates. This definition also avoids having to modify terms at runtime, which would be difficult to implement safely. We could also implement $f \oplus df$ as a function that invokes both f and df on its argument, as done by Cai et al. [7], but we believe that would be less efficient when \oplus is used at runtime. As we discuss in Sec. 3.4, we restrict validity to avoid this runtime overhead.

Having given these definitions, we show in Fig. 4 a step-indexed big-step semantics for change terms, defined through judgment $dE \vdash dt \Downarrow_n dv$ (where n can be omitted). This judgment is pronounced “Under the environment dE , the change term dt evaluates into the closed change value dv in n steps.” Rules [SDVAR] and [SDTUPLE] are unsurprising. To evaluate function calls in **let**-bindings “**let** $y = f x, dy = df x dx$ **in** dt ” we have three rules, depending on the shape of $dE(df)$. These rules all recompute the value v_y of y in the original environment, but compute differently the change dy to y . If $dE(df)$ replaces the value of f , [SDREPLACECALL] recomputes $v'_y = f x$ from scratch in the new environment, and bind dy to $!v'_y$ when evaluating the **let** body. If $dE(df)$ is the nil change for primitive \mathbf{p} , [SDPRIMITIVE NIL] computes dy by running \mathbf{p} ’s derivative through function $\Delta_{\mathbf{p}}(-)$. If $dE(df)$ is a closure change, [SDCLOSURECHANGE] invokes it normally to compute its change dv_y . As we show, if the closure change is valid, its body behaves like f ’s derivative, hence incrementalizes f correctly.

Closure changes with non-nil environment changes represent partial application of derivatives to non-nil changes; for instance, if f takes a pair and dx is a non-nil change, $0_{\text{curry}} f df x dx$ constructs a closure change containing dx , using the derivative of **curry** mentioned in Sec. 3.2. In general, such closure changes do not arise from the rules we show, only from derivatives of primitives.

3.4 A new soundness proof for static differentiation

In this section, we show that static differentiation is sound (Theorem 3.3) and that Eq. (1) holds:

$$f a_2 = f a_1 \oplus \mathcal{D}^c(f) a_1 da \tag{1}$$

whenever da is a valid change from a_1 to a_2 (as defined later). One might want to prove this equation assuming only that $a_1 \oplus da = a_2$, but this is false in general. A direct proof by induction on terms fails in the case for application (ultimately because $f_1 \oplus df = f_2$ and $a_1 \oplus da = a_2$ do not imply that $f_1 a_1 \oplus df a_1 da = f_2 a_2$).

- $d\ell \triangleright_n \ell \hookrightarrow \delta_{\oplus}(\ell, d\ell)$ • $!v_2 \triangleright_n v_1 \hookrightarrow v_2$ • $0_{\mathbf{p}} \triangleright_n \mathbf{p} \hookrightarrow \mathbf{p}$
- $(dv_1, \dots, dv_m) \triangleright_n (v_1, \dots, v_m) \hookrightarrow (v'_1, \dots, v'_m)$
 if and only if $(v_1, \dots, v_m) \oplus (dv_1, \dots, dv_m) = (v'_1, \dots, v'_m)$
 and $\forall k < n, \forall i \in [1 \dots m], dv_i \triangleright_k v_i \hookrightarrow v'_i$
- $dE[\lambda x dx. dt] \triangleright_n E_1[\lambda x. t] \hookrightarrow E_2[\lambda x. t]$
 if and only if $E_2 = E_1 \oplus dE$ and
 $\forall k < n, v_1, dv, v_2,$
 if $dv \triangleright_k v_1 \hookrightarrow v_2$ then $(dE; x = v_1; dx = dv \vdash dt) \blacktriangleright_k (E_1; x = v_1 \vdash t) \hookrightarrow (E_2; x = v_2 \vdash t)$
- $(dE \vdash dt) \blacktriangleright_n (E_1 \vdash t_1) \hookrightarrow (E_2 \vdash t_2)$
 if and only if $\forall k < n, v_1, v_2,$
 $E_1 \vdash t_1 \Downarrow_k v_1$ and $E_2 \vdash t_2 \Downarrow v_2$ implies that $\exists dv, dE \vdash dt \Downarrow dv \wedge dv \triangleright_{n-k} v_1 \hookrightarrow v_2$

Fig. 5: Step-indexed validity, through judgments for values and for terms.

As usual, this can be fixed by introducing a logical relation. We call ours *validity*: a function change is valid if it turns valid input changes into valid output changes.

Static differentiation is only sound on input changes that are *valid*. Cai et al. [7] show this for a strongly normalizing simply-typed λ -calculus using denotational semantics. Using an operational semantics, we generalize this result to an untyped and Turing-complete language, so we must turn to a *step-indexed* logical relation [4, 3].

Alert readers might wonder why we use a big-step semantics for a Turing-complete language. We indeed focus on incrementalizing computations that terminate on both old and new inputs. Our choice follows Acar et al. [3] (compared with in Sec. 5).

Validity as a step-indexed logical relation We say that “ dv is a valid change from v_1 to v_2 , up to k steps” and write

$$dv \triangleright_k v_1 \hookrightarrow v_2$$

to mean that dv is a change from v_1 to v_2 and that dv is a *valid* description of the differences between v_1 and v_2 , with validity tested with up to k steps. This relation *approximates* validity; if a change dv is valid at all approximations, it is simply valid (between v_1 and v_2); we write then $dv \triangleright v_1 \hookrightarrow v_2$ (omitting the step-index k) to mean that validity holds at all step-indexes. We similarly omit step-indexes k from other step-indexed relations when they hold for all k .

To justify this intuition of validity, we show that a valid change from v_1 to v_2 goes indeed from v_1 to v_2 (Theorem 3.1), and that if a change is valid up to k steps, it is also valid up to fewer steps (Lemma 3.2).

Theorem 3.1 (\oplus agrees with validity).

If $dv \triangleright_k v_1 \hookrightarrow v_2$ holds for all $k > 0$, then $v_1 \oplus dv = v_2$.

Lemma 3.2 (Downward-closure).

If $N \geq n$, then $dv \triangleright_N v_1 \hookrightarrow v_2$ implies $dv \triangleright_n v_1 \hookrightarrow v_2$.

Crucially, Theorem 3.1 enables (a) computing v_2 from a valid change and its source, and (b) showing Eq. (1) through validity. As discussed, \oplus ignores changes to closure bodies to be faster, which is only sound if those changes are nil; to ensure Theorem 3.1 still holds, validity on closure changes must be adapted accordingly and forbid non-nil changes to closure bodies. This choice, while unusual, does not affect our results: if input changes do not modify closure bodies, intermediate changes will not modify closure bodies either. Logical relation experts might regard this as a domain-specific invariant we add to our relation. Alternatives are discussed by Giarrusso [10, App. C].

As usual with step-indexing, validity is defined by well-founded induction over naturals ordered by $<$; to show well-foundedness we observe that evaluation always takes at least one step.

Validity for values, terms and environments is formally defined by cases in Fig. 5. First, a literal change $d\ell$ is a valid change from ℓ to $\ell \oplus d\ell = \delta_{\oplus}(\ell, d\ell)$. Since the function δ_{\oplus} is partial, the relation only holds for the literal changes $d\ell$ which are valid changes for ℓ . Second, a replacement change $!v_2$ is always a valid change from any value v_1 to v_2 . Third, a primitive nil change is a valid change between any primitive and itself. Fourth, a tuple change is valid up to step n , if each of its components is valid up to any step strictly less than n . Fifth, we define validity for closure changes. Roughly speaking, this statement means that a closure change is valid if (i) its environment change dE is valid for the original closure environment E_1 and for the new closure environment E_2 ; and (ii) when applied to related values, the closure *bodies* t are related by dt , as defined by the auxiliary judgment $(dE \vdash dt) \blacktriangleright_n (E_1 \vdash t_1) \hookrightarrow (E_2 \vdash t_2)$ for validity between terms under related environments (defined in C). As usual with step-indexed logical relations, in the definition for this judgment about terms, the number k of steps required to evaluate the term t_1 is subtracted from the number of steps n that can be used to relate the outcomes of the term evaluations.

Soundness of differentiation We can state a soundness theorem for differentiation without mentioning step-indexes; thanks to this theorem, we can compute the updated result v_2 not by rerunning a computation, but by updating the base result v_1 with the result change dv that we compute through a derivative on the input change. A corollary shows Eq. (1).

Theorem 3.3 (Soundness of differentiation in λ_{AL}). *If dE is a valid change environment from base environment E_1 to updated environment E_2 , that is $dE \triangleright E_1 \hookrightarrow E_2$, and if t converges both in the base and updated environment, that is $E_1 \vdash t \Downarrow v_1$ and $E_2 \vdash t \Downarrow v_2$, then $\mathcal{D}^t(t)$ evaluates under the change environment dE to a valid change dv between base result v_1 and updated result v_2 , that is $dE \vdash \mathcal{D}^t(t) \Downarrow dv$, $dv \triangleright v_1 \hookrightarrow v_2$ and $v_1 \oplus dv = v_2$.*

We must first show that derivatives map input changes valid up to k steps to output changes valid up to k steps, that is, the *fundamental property* of our step-indexed logical relation:

Lemma 3.4 (Fundamental Property).

For each n , if $dE \triangleright_n E_1 \hookrightarrow E_2$ then $(dE \vdash \mathcal{D}^t(t)) \blacktriangleright_n (E_1 \vdash t) \hookrightarrow (E_2 \vdash t)$.

<p>Translation of terms $\boxed{M = \mathcal{T}_t(t')}$</p> $\mathcal{T}_t(\mathbf{let } y = f x \mathbf{ in } t') = \mathbf{let } y, c_{fx}^y = f x \mathbf{ in } \mathcal{T}_t(t')$ $\mathcal{T}_t(\mathbf{let } y = (\bar{x}) \mathbf{ in } t') = \mathbf{let } y = (\bar{x}) \mathbf{ in } \mathcal{T}_t(t')$ $\mathcal{T}_t(x) = (x, \mathcal{C}(t))$	<p>Base terms</p> $M ::= \mathbf{let } y, c_{fx}^y = f x \mathbf{ in } M$ $ \mathbf{let } y = (\bar{x}) \mathbf{ in } M$ $ (x, C)$
<p>Cache of a term $\boxed{C = \mathcal{C}(t)}$</p> $\mathcal{C}(\mathbf{let } y = f x \mathbf{ in } t) = ((\mathcal{C}(t), y), c_{fx}^y)$ $\mathcal{C}(\mathbf{let } y = (\bar{x}) \mathbf{ in } t) = (\mathcal{C}(t), y)$ $\mathcal{C}(x) = ()$	<p>Cache terms/patterns</p> $C ::= (C, c_{fx}^y) (C, x) ()$
<p>Translation of values $\boxed{V = \mathcal{T}(v)}$</p> $\mathcal{T}((\bar{v})) = (\overline{\mathcal{T}(v)})$ $\mathcal{T}(E[\lambda x. t]) = \mathcal{T}(E)[\lambda x. \mathcal{T}_t(t)]$ $\mathcal{T}(\ell) = \ell$ $\mathcal{T}(\mathbf{p}) = \mathbf{p}$	<p>Closed values</p> $V ::= (\bar{V}) F[\lambda x. M] \ell \mathbf{p}$ <p>Cache values</p> $V_c ::= () (V_c, V_c) (V_c, V)$ <p>Evaluation environments</p> $F ::= \bullet F; D_v$ <p>Base environment entries</p> $D_v ::= x = V c_{fx}^y = V_c$

Fig. 6: Cache-Transfer Style translation and syntax of its target language λ_{CAL} .

3.5 CTS conversion

Fig. 6 and 7 define both the syntax of λ_{CAL} and λ_{ICAL} and CTS conversion. The latter comprises CTS differentiation $\mathcal{D}(-)$, from λ_{AL} to λ_{ICAL} , and CTS translation $\mathcal{T}(-)$, from λ_{AL} to λ_{CAL} .

Syntax definitions for the target languages λ_{CAL} and λ_{ICAL} Terms of λ_{CAL} follow again λ -lifted A'NF, like λ_{AL} , except that a **let**-binding for a function application “ $f x$ ” now binds an extra *cache identifier* c_{fx}^y besides output y . Cache identifiers have non-standard syntax: it can be seen as a triple that refers to the value identifiers f , x and y . Hence, an α -renaming of one of these three identifiers must refresh the cache identifier accordingly. Result terms explicitly return cache C through syntax (x, C) . Caches are encoded through nested tuples, but they are in fact a tree-like data structure that is isomorphic to an execution trace. This trace contains both immediate values and the execution traces of nested function calls.

The syntax for λ_{ICAL} matches the image of the CTS derivative and witnesses the CTS discipline followed by the derivatives: to determine dy , the derivative of f evaluated at point x with change dx expects the cache produced by evaluating y in the base term. The derivative returns the updated cache which contains the intermediate results that would be gathered by the evaluation of $f(x \oplus dx)$. The result term of every change term returns the computed change and a cache update dC , where each value identifier x of the input cache is updated with its corresponding change dx .

<p>Differentiation of terms $\boxed{dM = \mathcal{D}_t(t')}$</p> $\mathcal{D}_t(\mathbf{let} \ y = f \ x \ \mathbf{in} \ t') = \mathbf{let} \ dy, c_{fx}^y = df \ x \ dx \ c_{fx}^y \ \mathbf{in} \ \mathcal{D}_t(t')$ $\mathcal{D}_t(\mathbf{let} \ y = (\bar{x}) \ \mathbf{in} \ t') = \mathbf{let} \ dy = (\overline{dx}) \ \mathbf{in} \ \mathcal{D}_t(M')$ $\mathcal{D}_t(x) = (dx, \mathcal{U}(t))$ <p>Cache update of a term $\boxed{dC = \mathcal{U}(t)}$</p> $\mathcal{U}(\mathbf{let} \ y = f \ x \ \mathbf{in} \ t) = ((\mathcal{U}(t), y \oplus dy), c_{fx}^y)$ $\mathcal{U}(\mathbf{let} \ y = (\bar{x}) \ \mathbf{in} \ t) = (\mathcal{U}(t), y \oplus dy)$ $\mathcal{U}(x) = ()$ <p>Differentiation of change values $\boxed{dV = \mathcal{T}(dv)}$</p> $\mathcal{T}((\overline{dv})) = (\mathcal{T}(\overline{dv}))$ $\mathcal{T}(dE[\lambda x \ dx. \mathcal{D}^t(t)]) = \mathcal{T}(dE)[\lambda x \ dx \ C(t)]. \mathcal{D}_t(t)$ $\mathcal{T}(!v) = !\mathcal{T}(v)$ $\mathcal{T}(d\ell) = d\ell$ $\mathcal{T}(0_{\mathbf{p}}) = 0_{\mathbf{p}}$	<p>Change terms</p> $dM ::= \mathbf{let} \ dy, c_{fx}^y = df \ x \ dx \ c_{fx}^y \ \mathbf{in} \ dM$ $\quad \ \mathbf{let} \ dy = (\overline{dx}) \ \mathbf{in} \ dM$ $\quad \ (dx, dC)$ <p>Cache updates</p> $dC ::= (dC, c_{fx}^y) \mid (dC, x \oplus dx) \mid ()$ <p>Change values</p> $dV ::= (\overline{dV}) \mid dF[\lambda x \ dx \ C. dM] \mid d\ell \mid 0_{\mathbf{p}} \mid !V$ <p>Change environments</p> $dF ::= \bullet \mid dF; dD_v$ <p>Change environment entries</p> $dD_v ::= D_v \mid dx = dV$
---	--

Fig. 7: CTS differentiation and syntax of its target language λ_{ICAL} . Beware $\mathcal{T}(dE[\lambda x \ dx. \mathcal{D}^t(t)])$ applies a left-inverse of $\mathcal{D}^t(t)$ during pattern matching.

CTS conversion and differentiation These translations use two auxiliary functions: $\mathcal{C}(t)$ which computes the cache term of a λ_{AL} term t , and $\mathcal{U}(t)$, which computes the cache update of t 's derivative.

CTS translation on terms, $\mathcal{T}_t(t')$, accepts as inputs a *global* term t and a subterm t' of t . In tail position ($t' = x$), the translation generates code to return both the result x and the cache $\mathcal{C}(t)$ of the global term t . When the transformation visits **let**-bindings, it outputs extra bindings for caches c_{fx}^y on function calls and visits the **let**-body.

Similarly to $\mathcal{T}_t(t')$, CTS derivation $\mathcal{D}_t(t')$ accepts a global term t and a subterm t' of t . In tail position, the translation returns both the result change dx and the cache update $\mathcal{U}(t)$. On **let**-bindings, it *does not* output bindings for y but for dy , it outputs extra bindings for c_{fx}^y as in the previous case and visits the **let**-body.

To handle function definitions, we transform the base environment E through $\mathcal{T}(E)$ and $\mathcal{T}(\mathcal{D}^t(E))$ (translations of environments are done pointwise, see **D**). Since $\mathcal{D}^t(E)$ includes E , we describe $\mathcal{T}(\mathcal{D}^t(E))$ to also cover $\mathcal{T}(E)$. Overall, $\mathcal{T}(\mathcal{D}^t(E))$ CTS-converts each source closure $f = E[\lambda x. t]$ to a CTS-translated function, with body $\mathcal{T}_t(t)$, and to the CTS derivative df of f . This CTS derivative pattern matches on its input cache using cache pattern $\mathcal{C}(t)$. That way, we make sure that the shape of the cache expected by df is consistent with the shape of the cache produced by f . The body of derivative df is computed by CTS-deriving f 's body via $\mathcal{D}_t(t)$.

3.6 Semantics of λ_{CAL} and λ_{ICAL}

An evaluation environment F of λ_{CAL} contains both values and cache values. Values V resemble λ_{AL} values v , cache values V_c match cache terms C and change values dV match λ_{ICAL} change values dv . Evaluation environments dF for change terms must also bind change values, so functions in change closures take not just a base input x and an input change dx , like in λ_{ICAL} , but also an input cache C . By abuse of notation, we reuse the same syntax C to both deconstruct and construct caches.

Base terms of the language are evaluated using a conventional big-step semantics, consisting of two judgments. Judgment “ $F \vdash M \Downarrow (V, V_c)$ ” is read “Under evaluation environment F , base term M evaluates to value V and cache V_c ”. The semantics follows the one of λ_{AL} ; since terms include extra code to produce and carry caches along the computation, the semantics evaluates that code as well. For space reasons, we defer semantic rules to the appendix. Auxiliary judgment “ $F \vdash C \Downarrow V_c$ ” evaluates cache terms into cache values: It traverses a cache term and looks up the environment for the values to be cached.

Change terms of λ_{ICAL} are also evaluated using a big-step semantics, which resembles the semantics of λ_{ICAL} and λ_{CAL} . Unlike those semantics, evaluating cache updates ($dC, x \oplus dx$) is evaluated using the \oplus operator (overloaded on λ_{CAL} values and λ_{ICAL} changes). By lack of space, its rules are deferred to Appendix E. This semantics relies on three judgments. Judgment “ $dF \vdash dM \Downarrow (dV, V_c)$ ” is read “Under evaluation environment F , change term dM evaluates to change value dV and updated cache V_c ”. The first auxiliary judgment “ $dF \vdash dC \Downarrow V_c$ ” defines evaluation of cache update terms. The final auxiliary judgment “ $V_c \sim C \rightarrow dF$ ” describes a limited form of pattern matching used by CTS derivatives: namely, how a cache pattern C matches a cache value V_c to produce a change environment dF .

3.7 Soundness of CTS conversion

The proof is based on a simulation in lock-step, but two subtle points emerge. First, we must relate λ_{AL} environments that do not contain caches, with λ_{CAL} environments that do. Second, while evaluating CTS derivatives, the evaluation environment mixes caches from the base computation and updated caches computed by the derivatives.

Theorem 3.7 follows because differentiation is sound (Theorem 3.3) and evaluation commutes with CTS conversion; this last point requires two lemmas. First, CTS translation of base terms commutes with our semantics:

Lemma 3.5 (Commutation for base evaluations).

For all E, t and v , if $E \vdash t \Downarrow v$, there exists $V_c, \mathcal{T}(E) \vdash \mathcal{T}_t(t) \Downarrow (\mathcal{T}(v), V_c)$.

Second, we need a corresponding lemma for CTS translation of differentiation results: intuitively, evaluating a derivative and CTS translating the resulting change value must give the same result as evaluating the CTS derivative. But

to formalize this, we must specify which environments are used for evaluation, and this requires two technicalities.

Assume derivative $\mathcal{D}^i(t)$ evaluates correctly in some environment dE . Evaluating CTS derivative $\mathcal{D}_t(t)$ requires cache values from the base computation, but they are not in $\mathcal{T}(dE)$! Therefore, we must introduce a judgment to complete a CTS-translated environment with the appropriate caches (see Appendix F).

Next, consider evaluating a change term of the form $dM = \mathbb{C}[dM']$, where \mathbb{C} is a standard single-hole change-term context—that is, for λ_{ICAL} , a sequence of **let**-bindings. When evaluating dM , we eventually evaluate dM' in a change environment dF updated by \mathbb{C} : the change environment dF contains both the updated caches coming from the evaluation of \mathbb{C} and the caches coming from the base computation (which will be updated by the evaluation of dM). Again, a new judgment, given in Appendix F, is required to model this process.

With these two judgments, the second key Lemma stating the commutation between evaluation of derivatives and evaluation of CTS derivatives can be stated. We give here an informal version of this Lemma, the actual formal version can be found in Appendix F.

Lemma 3.6 (Commutation for derivatives evaluation).

If the evaluation of $\mathcal{D}^i(t)$ leads to an environment dE_0 when it reaches the differentiated context $\mathcal{D}^i(\mathbb{C})$ where $t = \mathbb{C}[t']$, and if the CTS conversion of t under this environment completed with base (resp. changed) caches evaluates into a base value $\mathcal{T}(v)$ (resp. a changed value $\mathcal{T}(v')$) and a base cache value V_c (resp. an updated cache value V'_c), then under an environment containing the caches already updated by the evaluation of $\mathcal{D}^i(\mathbb{C})$ and the base caches to be updated, the CTS derivative of t' evaluates to $\mathcal{T}(dv)$ such that $v \oplus dv = v'$ and to the updated cache V'_c .

Finally, we can state soundness of CTS differentiation. This theorem says that CTS derivatives not only produce valid changes for incrementalization but that they also correctly consume and update caches.

Theorem 3.7 (Soundness of CTS differentiation).

If the following hypotheses hold:

1. $dE \triangleright E \leftrightarrow E'$
2. $E \vdash t \Downarrow v$
3. $E' \vdash t \Downarrow v'$

then there exists dv , V_c , V'_c and F_0 such that:

1. $\mathcal{T}(E) \vdash \mathcal{T}(t) \Downarrow (\mathcal{T}(v), V_c)$
2. $\mathcal{T}(E') \vdash \mathcal{T}(t) \Downarrow (\mathcal{T}(v'), V'_c)$
3. $\mathcal{C}(t) \sim V_c \rightarrow F_0$
4. $\mathcal{T}(dE); F_0 \vdash \mathcal{D}_t(t) \Downarrow (\mathcal{T}(dv), V'_c)$
5. $v \oplus dv = v'$

4 Incrementalization case studies

In this section, we investigate two questions: whether our transformations can target a typed language like Haskell and whether automatically transformed programs can perform well. We implement by hand primitives on sequences, bags and maps in Haskell. The input terms in all case studies are written in a deep embedding of λ_{AL} into Haskell. The transformations generate Haskell code that uses our primitives and their derivatives.

We run the transformations on three case studies: a computation of the average value of a bag of integers, a nested loop over two sequences and a more involved example inspired by Koch et al. [17]’s work on incrementalizing database queries. For each case study, we make sure that results are consistent between from scratch recomputation and incremental evaluation; we measure the execution time for from scratch recomputation and incremental computation as well as the space consumption of caches. We obtain efficient incremental programs, that is ones for which incremental computation is faster than from scratch recomputation. The measurements indicate that we do get the expected asymptotic improvement in time of incremental computation over from scratch recomputation by a linear factor while the caches grows in a similar linear factor.

Our benchmarks were compiled by GHC 8.2.2 and run on a 2.20GHz hexa core Intel(R) Xeon(R) CPU E5-2420 v2 with 32GB of RAM running Ubuntu 14.04. We use the *criterion* [21] benchmarking library.

4.1 Averaging bags of integers

Sec. 2.1 motivates our transformation with a running example of computing the average over a bag of integers. We represent bags as maps from elements to (possibly negative) multiplicities. Earlier work [7, 17] represents bag changes as bags of removed and added elements. We use a different representation of bag changes that takes advantage of the changes to elements and provide primitives on bags and their derivatives. The CTS variant of *map*, that we call *mapC*, takes a function *fC* in CTS and a bag *as* and produces a bag and a cache. The cache stores for each invocation of *fC*, and therefore for each distinct element in *as*, the result of *fC* of type *b* and the cache of type *c*.

Inspired by Rossberg et al. [23], all higher-order functions (and typically, also their caches) are parametric over cache types of their function arguments. Here, functions *mapC* and *dmapC* and cache type *MapC* are parametric over the cache type *c* of *fC* and *dfC*.

```
map :: (a -> b) -> Bag a -> Bag b
data MapC a b c = MapC (Map a (b, c))
mapC :: (a -> (b, c)) -> Bag a -> (Bag b, MapC a b c)
dmapC :: (a -> (b, c)) -> (a -> Δa -> c -> (Δb, c)) -> Bag a -> Δ(Bag a) ->
    MapC a b c -> (Δ(Bag b), MapC a b c)
```

We wrote the *length* and *sum* functions used in our benchmarks in terms of primitives *map* and *foldGroup* and had their CTS function and CTS derivative generated automatically.

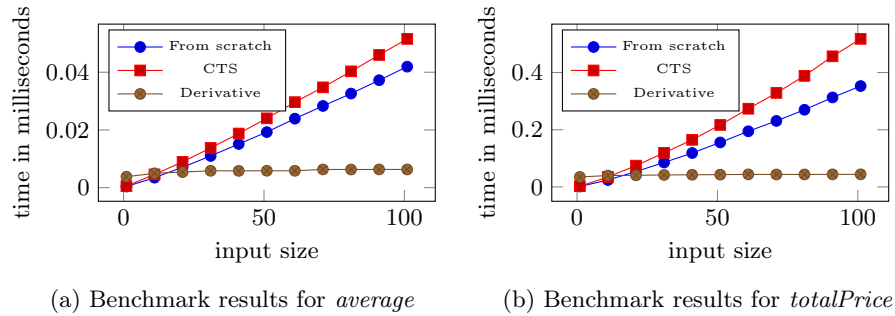


Fig. 8: Benchmark results for *average* and *totalPrice*

We evaluate whether we can produce an updated result with *daverageC* shown in Sec. 2.1 faster than by from scratch recomputation with *average*. We expect the speedup of *daverageC* to depend on the size of the input bag n . We fix an input bag of size n as the bag containing the numbers from 1 to n . We define a change that inserts the integer 1 into the bag. To measure execution time of from scratch recomputation, we apply *average* to the input bag updated with the change. To measure execution time of the CTS function *averageC*, we apply *averageC* to the input bag updated with the change. To measure execution time of the CTS derivative *daverageC*, we apply *daverageC* to the input bag, the change and the cache produced by *averageC* when applied to the input bag. In all three cases we ensure that all results and caches are fully forced so as to not hide any computational cost behind laziness.

The plot in Fig. 8a shows execution time versus the size n of the base input. To produce the base result and cache, the CTS transformed function *averageC* takes longer than the original *average* function takes to produce just the result. Producing the updated result incrementally is slower than from scratch recomputation for small input sizes, but because of the difference in time complexity becomes faster as the input size grows. The size of the cache grows linearly with the size of the input, which is not optimal for this example. We leave optimizing the space usage of examples like this to future work.

4.2 Nested loops over two sequences

Next, we consider CTS differentiation on a higher-order example. To incrementalize this example efficiently, we have to enable detecting nil function changes at runtime by representing function changes as closures that can be inspected by incremental programs. Our example here is the Cartesian product of two sequences computed in terms of functions *map* and *concat*.

```

cartesianProduct :: Sequence a → Sequence b → Sequence (a, b)
cartesianProduct xs ys = concatMap (λx → map (λy → (x, y)) ys) xs
concatMap :: (a → Sequence b) → Sequence a → Sequence b
concatMap f xs = concat (map f xs)

```

We implemented incremental sequences and related primitives following Firsov and Jeltsch [9]: our change operations and first-order operations (such as *concat*) reuse their implementation. On the other hand, we must extend higher-order operations such as *map* to handle non-nil function changes and caching. A correct and efficient CTS derivative *dmapC* has to work differently depending on whether the given function change is nil or not: For a non-nil function change it has to go over the input sequence; for a nil function change it has to avoid that.

Cai et al. [7] use static analysis to conservatively approximate nil function changes as changes to terms that are closed in the original program. But in this example the function argument $(\lambda y \rightarrow (x, y))$ to *map* in *cartesianProduct* is not a closed term. It is, however, crucial for the asymptotic improvement that we avoid looping over the inner sequence when the change to the free variable x in the change environment is $\mathbf{0}_x$.

To enable runtime nil change detection, we apply closure conversion to the original program and explicitly construct closures and changes to closures. While the only valid change for closed functions is their nil change, for closures we can have non-nil function changes. A function change *df*, represented as a closure change, is nil exactly when all changes it closes over are nil.

We represent closed functions and closures as variants of the same type. Correspondingly we represent changes to a closed function and changes to a closure as variants of the same type of function changes. We inspect this representation at runtime to find out if a function change is a nil change.

```


data Fun a b c where
  Closed :: (a → (b, c)) → Fun a b c
  Closure :: (e → a → (b, c)) → e → Fun a b c

data  $\Delta$ (Fun a b c) where
  DClosed :: (a →  $\Delta$ a → c → ( $\Delta$ b, c)) →  $\Delta$ (Fun a b c)
  DClosure :: (e →  $\Delta$ e → a →  $\Delta$ a → c → ( $\Delta$ b, c)) → e →  $\Delta$ e →  $\Delta$ (Fun a b c)

```

We use the same benchmark setup as in the benchmark for the average computation on bags. The input of size n is a pair of sequences (xs, ys) . Each sequence initially contains the integers from 1 to n . Updating the result in reaction to a change *dxs* to the outer sequence xs takes less time than updating the result in reaction to a change *dys* to the inner sequence ys . While a change to the outer sequence xs results in an easily located change in the output sequence, a change for the inner sequence ys results in a change that needs a lot more calculation to find the elements it affects. We benchmark changes to the outer sequence xs and the inner sequence ys separately where the change to one sequence is the insertion of a single integer 1 at position 1 and the change for the other one is the nil change.

Fig. 9 shows execution time versus input size. In this example again preparing the cache takes longer than from scratch recomputation alone. The speedup of incremental computation over from scratch recomputation increases with the size of the base input sequences because of the difference in time complexity. Eventually we do get speedups for both kinds of changes (to the inner and to the outer sequence), but for changes to the outer sequence we get a speedup



(a) Benchmark results for Cartesian product changing *inner* sequence. (b) Benchmark results for Cartesian product changing *outer* sequence.

Fig. 9: Benchmark results for *cartesianProduct*

earlier, at a smaller input size. The size of the cache grows super linearly in this example.

4.3 Indexed joins of two bags

Our goal is to show that we can compose primitive functions into larger and more complex programs and apply CTS differentiation to get a fast incremental program. We use an example inspired from the DBToaster literature [17]. In this example we have a bag of orders and a bag of line items. An order is a pair of an order key and an exchange rate. A line item is a pair of an order key and a price. We build an index mapping each order key to the sum of all exchange rates of the orders with this key and an index from order key to the sum of the prices of all line items with this key. We then merge the two maps by key, multiplying corresponding sums of exchange rates and sums of prices. We compute the total price of the orders and line items as the sum of those products.

```

type Order = (Z, Z)
type LineItem = (Z, Z)
totalPrice :: Bag Order → Bag LineItem → Z
totalPrice orders lineItems = let
  orderIndex = groupBy fst orders
  orderSumIndex = Map.map (Bag.foldMapGroup snd) orderIndex
  lineItemIndex = groupBy fst lineItems
  lineItemSumIndex = Map.map (Bag.foldMapGroup snd) lineItemIndex
  merged = Map.merge orderSumIndex lineItemSumIndex
  total = Map.foldMapGroup multiply merged
in total

groupBy :: (a → k) → Bag a → Map k (Bag a)
groupBy keyOf bag =
  Bag.foldMapGroup (λa → Map.singleton (keyOf a) (Bag.singleton a)) bag

```

Unlike DBToaster, we assume our program is already transformed to explicitly use indexes, as above. Because our indexes are maps, we implemented a change structure, CTS primitives and their CTS derivatives for maps.

To build the indexes, we use a *groupBy* function built from primitive functions *foldMapGroup* on bags and *singleton* for bags and maps respectively. The CTS function *groupByC* and the CTS derivative *dgroupByC* are automatically generated. While computing the indexes with *groupBy* is self-maintainable, merging them is not. We need to cache and incrementally update the intermediately created indexes to avoid recomputing them.

We evaluate the performance in the same way we did in the other case studies. The input of size n is a pair of bags where both contain the pairs (i, i) for i between 1 and n . The change is an insertion of the order $(1, 1)$ into the orders bag. For sufficiently large inputs, our CTS derivative of the original program produces updated results much faster than from scratch recomputation, again because of a difference in time complexity as indicated by Fig. 8b. The size of the cache grows linearly with the size of the input in this example. This is unavoidable, because we need to keep the indexes.

4.4 Limitations and future work

Typing of CTS programs Functions of the same type $f_1, f_2 :: A \rightarrow B$ can be transformed to CTS functions $f_1 :: A \rightarrow (B, C_1), f_2 :: A \rightarrow (B, C_2)$ with different cache types C_1, C_2 , since cache types depend on the implementation. This heterogeneous typing of translated functions poses difficult typing issues, e.g. what is the translated type of a *list* $(A \rightarrow B)$? We cannot hide cache types behind existential quantifiers because they would be too abstract for derivatives, which only work on very specific cache types. We can fix this problem with some runtime overhead by using a single type *Cache*, defined as a tagged union of all cache types or, maybe with more sophisticated type systems — like first-class translucent sums, open existentials or Typed Adaption’s refinement types [12] — that could be able to correctly track down cache types properly.

In any case, we believe that these machineries would add a lot of complexity without helping much with the proof of correctness. Indeed, the simulation relation is more handy here because it maintains a global invariant about the whole evaluations (typically the consistency of cache types between base computations and derivatives), not many local invariants about values as types would.

One might wonder why caches could not be totally hidden from the programmer by embedding them in the derivatives themselves; or in other words, why we did not simply translate functions of type $A \rightarrow B$ into functions of type $A \rightarrow B \times (\Delta A \rightarrow \Delta B)$. We tried this as well; but unlike automatic differentiation, we must remember and update caches according to input changes (especially when receiving a sequence of such changes as in Sec. 2.1). Returning the updated cache to the caller works; we tried closing over the caches in the derivative, but this ultimately fails (because we could receive function changes to the original function, but those would need access to such caches).

Comprehensive performance evaluation This paper focuses on theory and we leave benchmarking in comparison to other implementations of incremental computation to future work. The examples in our case study were rather simple (except perhaps for the indexed join). Nevertheless, the results were encouraging and we expect them to carry over to more complex examples, but not to all programs. A comparison to other work would also include a comparison of space usage for auxiliary data structure, in our case the caches.

Cache pruning via absence analysis To reduce memory usage and runtime overhead, it should be possible to automatically remove from transformed programs any caches or cache fragments that are not used (directly or indirectly) to compute outputs. Liu [19] performs this transformation on CTS programs by using *absence analysis*, which was later extended to higher-order languages by Sergey et al. [25]. In lazy languages, absence analysis removes thunks that are not needed to compute the output. We conjecture that the analysis could remove unused caches or inputs, if it is extended to *not* treat caches as part of the output.

Unary vs n -ary abstraction We only show our transformation correct for unary functions and tuples. But many languages provide efficient support for applying curried functions such as $div :: \mathbb{Z} \rightarrow \mathbb{Z} \rightarrow \mathbb{Z}$. Naively transforming such a curried function to CTS would produce a function $divC$ of type $\mathbb{Z} \rightarrow (\mathbb{Z} \rightarrow (\mathbb{Z}, DivC_2)), DivC_1$ with $DivC_1 = ()$, which adds excessive overhead. In Sec. 2 and our evaluation we use curried functions and never need to use this naive encoding, but only because we always invoke functions of known arity.

5 Related work

Cache-transfer-style Liu [19]’s work has been the fundamental inspiration to this work, but her approach has no correctness proof and is restricted to a first-order untyped language. Moreover, while the idea of cache-transfer-style is similar, it’s unclear if her approach to incrementalization would extend to higher-order programs. Firsov and Jeltsch [9] also approach incrementalization by code transformation, but their approach does not deal with changes to functions. Instead of transforming functions written in terms of primitives, they provide combinators to write CTS functions and derivatives together. On the other hand, they extend their approach to support mutable caches, while restricting to immutable ones as we do might lead to a logarithmic slowdown.

Finite differencing Incremental computation on collections or databases by finite differencing has a long tradition [22, 6]. The most recent and impressive line of work is the one on DBToaster [16, 17], which is a highly efficient approach to incrementalize queries over bags by combining iterated finite differencing with other program transformations. They show asymptotic speedups both in theory and through experimental evaluations. Changes are only allowed for datatypes that form groups (such as bags or certain maps), but not for instance for lists or

sets. Similar ideas were recently extended to higher-order and nested computation [18], though only for datatypes that can be turned into groups. Koch et al. [18] emphasize that iterated differentiation is necessary to obtain efficient derivatives; however, ANF conversion and remembering intermediate results appear to address the same problem, similarly to the field of automatic differentiation [27].

Logical relations To study correctness of incremental programs we use a logical relation among base values v_1 , updated values v_2 and changes dv . To define a logical relation for an untyped λ -calculus we use a *step-indexed* logical relation, following Appel and McAllester [5], Ahmed [4]; in particular, our definitions are closest to the ones by Acar et al. [3], who also work with an untyped language, big-step semantics and (a different form of) incremental computation. However, they do not consider first-class changes. Technically, we use environments rather than substitution, and index our big-step semantics differently.

Dynamic incrementalization The approaches to incremental computation with the widest applicability are in the family of self-adjusting computation [1, 2], including its descendant Adapton [14]. These approaches incrementalize programs by combining memoization and change propagation: after creating a trace of base computations, updated inputs are compared with old ones in $O(1)$ to find corresponding outputs, which are updated to account for input modifications. Compared to self-adjusting computation, Adapton only updates results that are demanded. As usual, incrementalization is not efficient on arbitrary programs, but only on programs designed so that input changes produce small changes to the computation trace; refinement type systems have been designed to assist in this task [8, 12]. To identify matching inputs, Nominal Adapton [13] replaces input comparisons by pointer equality with first-class labels, enabling more reuse.

6 Conclusion

We have presented a program transformation which turns a functional program into its derivative and efficiently shares redundant computations between them thanks to a statically computed cache. This work has been mechanized using Coq and the corresponding proof is available online, together with the case study material.⁵

Although our first practical case studies show promising results, this paper focused on putting CTS differentiation on solid theoretical ground. For the moment, we only have scratched the surface of the incrementalization opportunities opened by CTS primitives and their CTS derivatives: in our opinion, exploring the design space for cache data structures will lead to interesting new results in purely functional incremental programming.

Acknowledgments We are grateful to reviewers of previous versions of this paper: they made important suggestions to help us improve our technical presentation. This work was partially supported by DFG project 282458149.

⁵ <https://github.com/yurug/cts>

Bibliography

- [1] Acar, U.A.: Self-Adjusting Computation. Ph.D. thesis, Carnegie Mellon University (2005)
- [2] Acar, U.A.: Self-adjusting computation: (an overview). In: PEPM. pp. 1–6. ACM (2009)
- [3] Acar, U.A., Ahmed, A., Blume, M.: Imperative self-adjusting computation. In: Proceedings of the 35th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages. pp. 309–322. POPL '08, ACM, New York, NY, USA (2008), <http://doi.acm.org/10.1145/1328438.1328476>
- [4] Ahmed, A.: Step-indexed syntactic logical relations for recursive and quantified types. In: Programming Languages and Systems. pp. 69–83. Springer, Berlin, Heidelberg (Mar 2006)
- [5] Appel, A.W., McAllester, D.: An indexed model of recursive types for foundational proof-carrying code. *ACM Trans. Program. Lang. Syst.* 23(5), 657–683 (Sep 2001), <http://doi.acm.org/10.1145/504709.504712>
- [6] Blakeley, J.A., Larson, P.A., Tompa, F.W.: Efficiently updating materialized views. In: SIGMOD. pp. 61–71. ACM (1986)
- [7] Cai, Y., Giarrusso, P.G., Rendel, T., Ostermann, K.: A theory of changes for higher-order languages — incrementalizing λ -calculi by static differentiation. In: Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation. pp. 145–155. PLDI '14, ACM, New York, NY, USA (2014), <http://doi.acm.org/10.1145/2594291.2594304>
- [8] Çiçek, E., Paraskevopoulou, Z., Garg, D.: A type theory for incremental computational complexity with control flow changes. In: Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming. pp. 132–145. ICFP 2016, ACM, New York, NY, USA (2016)
- [9] Firsov, D., Jeltsch, W.: Purely functional incremental computing. In: Castor, F., Liu, Y.D. (eds.) *Programming Languages*, pp. 62–77. No. 9889 in *Lecture Notes in Computer Science*, Springer International Publishing (Sep 2016)
- [10] Giarrusso, P.G.: Optimizing and Incrementalizing Higher-order Collection Queries by AST Transformation. Ph.D. thesis, University of Tübingen (2018), defended. Draft available at <http://inc-lc.github.io/>
- [11] Gupta, A., Mumick, I.S.: Maintenance of materialized views: problems, techniques, and applications. In: Gupta, A., Mumick, I.S. (eds.) *Materialized views*, pp. 145–157. MIT Press (1999)
- [12] Hammer, M.A., Dunfield, J., Economou, D.J., Narasimhamurthy, M.: Typed Adaption: refinement types for incremental computations with precise names. *arXiv:1610.00097 [cs]* (Oct 2016)
- [13] Hammer, M.A., Dunfield, J., Headley, K., Labich, N., Foster, J.S., Hicks, M., Van Horn, D.: Incremental computation with names. In: Proceedings of the 2015 ACM SIGPLAN International Conference on Object-Oriented

- Programming, Systems, Languages, and Applications. pp. 748–766. OOPSLA 2015, ACM, New York, NY, USA (2015), <http://doi.acm.org/10.1145/2814270.2814305>
- [14] Hammer, M.A., Phang, K.Y., Hicks, M., Foster, J.S.: Adapton: composable, demand-driven incremental computation. In: Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation. pp. 156–166. PLDI '14, ACM, New York, NY, USA (2014)
 - [15] Johnsson, T.: Lambda lifting: transforming programs to recursive equations. In: Jouannaud, J.P. (ed.) Functional Programming Languages and Computer Architecture, pp. 190–203. No. 201 in Lecture Notes in Computer Science, Springer Berlin Heidelberg (Sep 1985)
 - [16] Koch, C.: Incremental query evaluation in a ring of databases. In: Symp. Principles of Database Systems (PODS). pp. 87–98. ACM (2010)
 - [17] Koch, C., Ahmad, Y., Kennedy, O., Nikolic, M., Nötzli, A., Lupei, D., Shaikhha, A.: DBToaster: higher-order delta processing for dynamic, frequently fresh views. *The VLDB Journal* 23(2), 253–278 (2014), <http://dx.doi.org/10.1007/s00778-013-0348-4>
 - [18] Koch, C., Lupei, D., Tannen, V.: Incremental view maintenance for collection programming. In: Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems. pp. 75–90. PODS '16, ACM, New York, NY, USA (2016)
 - [19] Liu, Y.A.: Efficiency by incrementalization: an introduction. *HOSC* 13(4), 289–313 (2000)
 - [20] Liu, Y.A., Teitelbaum, T.: Caching intermediate results for program improvement. In: Proceedings of the 1995 ACM SIGPLAN symposium on Partial evaluation and semantics-based program manipulation. pp. 190–201. PEPM '95, ACM, New York, NY, USA (1995), <http://doi.acm.org/10.1145/215465.215590>
 - [21] O'Sullivan, B.: criterion: a Haskell microbenchmarking library (2014), <http://www.serpentine.com/criterion/>
 - [22] Paige, R., Koenig, S.: Finite differencing of computable expressions. *TOPLAS* 4(3), 402–454 (Jul 1982)
 - [23] Rossberg, A., Russo, C.V., Dreyer, D.: F-ing modules. In: Proceedings of the 5th ACM SIGPLAN Workshop on Types in Language Design and Implementation. pp. 89–102. TLDI '10, ACM, New York, NY, USA (2010)
 - [24] Sabry, A., Felleisen, M.: Reasoning about programs in continuation-passing style. *Lisp and symbolic computation* 6(3-4), 289–360 (1993)
 - [25] Sergey, I., Vytiniotis, D., Peyton Jones, S.: Modular, higher-order cardinality analysis in theory and practice. In: Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. pp. 335–347. POPL '14, ACM, New York, NY, USA (2014)
 - [26] The Coq development team: The Coq proof assistant reference manual (2018), <http://coq.inria.fr>, version 8.8
 - [27] Wang, F., Wu, X., Essertel, G., Decker, J., Rompf, T.: Demystifying differentiable programming: Shift/reset the penultimate backpropagator (????), <https://arxiv.org/abs/1803.10228>

A Comparison with original static differentiation

Cai et al. [7]’s static differentiation for λ -terms [7] is defined instead as follows:

$$\begin{aligned} \text{Derive}(x) &= dx \\ \text{Derive}(t\ u) &= (\text{Derive}(t))\ u\ (\text{Derive}(u)) \\ \text{Derive}(\lambda x. t) &= \lambda x\ dx. \text{Derive}(t) \end{aligned}$$

Even though the first two cases of Cai et al. [7]’s differentiation map into the two cases of our differentiation variant, one may ask where the third case is realized now. Actually, this third case occurs while we transform the base environment E . Indeed, we will assume that the closures of the environment of the source program have been adjoined a nil change or derivative, as defined by $\mathcal{D}^\ell(E)$. Change value $\mathcal{D}^\ell(v)$ represents a nil change for v ; if v is a closure, its nil change is also its derivative [7]. Computing $\mathcal{D}^\ell(v)$ requires a meta-level function **nil** which computes a **nil** change for every literal ℓ . The existence of such a function is generally assumed in previous work [7] as it is always valid to use the change $!\ell$. Yet, since this triggers recomputation, it can be more efficient to pick nil changes appropriate to the domain and to the definition of \oplus .

B Description of the semantic rules for λ_{AL}

Rule [SVAR] looks variable x up in environment E . Other rules evaluate **let**-binding “**let** $y = \dots$ **in** t ” in environment E : Each rule computes y ’s new value v_y (taking m steps) and evaluates in n steps the body t to a value v , using environment E extended by binding y to v_y . The overall **let**-binding evaluates to v in $m + n + 1$ steps. But different rules compute the value of y differently. [STUPLE] looks each variable in \bar{x} up in E to evaluate tuple (\bar{x}) (in $m = 0$ steps). [SPRIMITIVECALL] evaluates function calls where variable f is bound in E to a primitive \mathbf{p} , evaluated as specified by a function $\delta_{\mathbf{p}}(-)$ from closed values to closed values. To evaluate such a primitive call, this rule applies $\delta_{\mathbf{p}}(-)$ to x ’s value (in $m = 0$ steps). [SCLOSURECALL] evaluates function calls where variable f is bound in E to closure $E_f[\lambda x. t_f]$: this rule evaluates the closure body t_f in m steps, using the closure environment E_f extended with the value of x in E .

C Validity relation over environments

$$\bullet \triangleright_k \bullet \bullet \hookrightarrow \bullet \quad \frac{dE \triangleright_k E_1 \hookrightarrow E_2 \quad dv \triangleright_k v_1 \hookrightarrow v_2}{dE; x = v_1; dx = dv \triangleright_k E_1; x = v_1 \hookrightarrow E_2; x = v_2}$$

D CTS translation of environments

Translation of value environments $\boxed{F = \mathcal{T}(E)}$

$$\begin{aligned} \mathcal{T}(\bullet) &= \bullet \\ \mathcal{T}(E; x = v) &= \mathcal{T}(E); x = \mathcal{T}(v) \end{aligned}$$

Differentiation of change environments $\boxed{dF = \mathcal{T}(dE)}$

$$\begin{aligned} \mathcal{T}(\bullet) &= \bullet \\ \mathcal{T}(dE; x = v, dx = dv) &= \mathcal{T}(dE); x = \mathcal{T}(v), dx = \mathcal{T}(dv) \end{aligned}$$

E Target language semantics

Evaluation of base terms $\boxed{F \vdash M \Downarrow (V, V_c)}$

$$\begin{array}{c} \text{[TRESULT]} \\ \frac{F(x) = V \quad F \vdash C \Downarrow V_c}{F \vdash (x, C) \Downarrow (V, V_c)} \end{array} \qquad \begin{array}{c} \text{[TTUPLE]} \\ \frac{F; y = F(\bar{x}) \vdash M \Downarrow (V, V_c)}{F \vdash \mathbf{let } y = (\bar{x}) \mathbf{ in } M \Downarrow (V, V_c)} \end{array}$$

$$\begin{array}{c} \text{[TCLOSURECALL]} \\ \frac{\begin{array}{c} F(f) = F'[\lambda x'. M'] \\ F'; x' = F(x) \vdash M' \Downarrow (V', V'_c) \\ F; y = V'; c_{fx}^y = V'_c \vdash M \Downarrow (V, V_c) \end{array}}{F \vdash \mathbf{let } y, c_{fx}^y = f \mathbf{ in } M \Downarrow (V, V_c)} \end{array}$$

$$\begin{array}{c} \text{[TPRIMITIVECALL]} \\ \frac{F(f) = \mathbf{p} \quad \delta_{\mathbf{p}}(F(x)) = (V', V'_c) \quad F; y = V'; c_{fx}^y = V'_c \vdash M \Downarrow v}{F \vdash \mathbf{let } y, c_{fx}^y = f \mathbf{ in } M \Downarrow (V, V_c)} \end{array}$$

Evaluation of caches $\boxed{F \vdash C \Downarrow V_c}$

$$\begin{array}{c} \text{[TEEMPTYCACHE]} \\ \frac{}{F \vdash () \Downarrow ()} \end{array} \qquad \begin{array}{c} \text{[TCACHEVAR]} \\ \frac{F(x) = V \quad F \vdash C \Downarrow V_c}{F \vdash (C, x) \Downarrow (V_c, V)} \end{array} \qquad \begin{array}{c} \text{[TCACHESUBCACHE]} \\ \frac{F(c_{fx}^y) = V'_c \quad F \vdash C \Downarrow V_c}{F \vdash (C, c_{fx}^y) \Downarrow (V_c, V'_c)} \end{array}$$

Fig. 10: Target language λ_{CAL} (semantics of base terms and caches).

Fig. 10 and 11 have been omitted from the paper: they contains the operational semantics for λ_{CAL} and λ_{ICAL} .

The rules for the evaluation of λ_{CAL} change terms are similar to the evaluation rules of λ_{AL} change terms except that (i) caches are carried along the computation and updated at return sites; (ii) there is no recomputation of base

Evaluation of change terms

$$\boxed{dF \vdash dM \Downarrow (dV, V_c)}$$

$$\frac{[\text{TDRRESULT}] \quad dF(dx) = dV \quad dF \vdash dC \Downarrow V_c}{dF \vdash (dx, dC) \Downarrow (dV, V_c)}$$

$$\frac{[\text{TDTUPLE}] \quad dF; dy = dF(\overline{dx}) \vdash dM \Downarrow (dV, V_c)}{dF \vdash \mathbf{let} \, dy = (\overline{dx}) \, \mathbf{in} \, dM \Downarrow (dV, V_c)}$$

$$\frac{[\text{TDRPLACECALL}] \quad \begin{array}{c} dF(df) = !V_f \\ [dF]_2 \vdash @ (f, x) \Downarrow (V', V_c) \\ dF; dy = !V'; c_{fx}^y = V_c' \vdash dM \Downarrow (dV, V_c) \end{array}}{dF \vdash \mathbf{let} \, dy, c_{fx}^y = df \, dx \, c_{fx}^y \, \mathbf{in} \, dM \Downarrow (dV, V_c)}$$

$$\frac{[\text{TDPRIMITIVE NIL}] \quad \begin{array}{c} dF(f, df) = \mathbf{p}, 0_{\mathbf{p}} \\ dF(x, dx) = V_x, dV_x \quad dE; dy, c_{fx}^y = \Delta_{\mathbf{p}}(V_x, dV_x, dF(c_{fx}^y)) \vdash dM \Downarrow (dV, V_c) \end{array}}{dF \vdash \mathbf{let} \, dy, c_{fx}^y = df \, x \, dx \, c_{fx}^y \, \mathbf{in} \, dM \Downarrow (dV, V_c)}$$

$$\frac{[\text{TDCLOSURECHANGE}] \quad \begin{array}{c} dF(df) = dF_f[\lambda x \, dx \, C. dM_f] \\ dF(c_{fx}^y) \sim C \rightarrow dF' \quad dF_f; x = dF(x); dx = dF(dx); dF' \vdash dM_f \Downarrow (dV_y, V_c') \\ dF; dy = dV_y, c_{fx}^y = V_c' \vdash dM \Downarrow (dV, V_c) \end{array}}{dF \vdash \mathbf{let} \, dy, c_{fx}^y = df \, x \, dx \, c_{fx}^y \, \mathbf{in} \, dM \Downarrow (dV, V_c)}$$

Binding of caches

$$\boxed{V_c \sim C \rightarrow dF}$$

$$\frac{[\text{TMATCHEMPTYCACHE}] \quad () \sim () \rightarrow ()}{() \sim () \rightarrow ()} \quad \frac{[\text{TMATCHCACHEDVALUE}] \quad V_c \sim C \rightarrow dF}{(V_c, V) \sim (C, x) \rightarrow dF; (x = V)}$$

$$\frac{[\text{TMATCHSUBCACHE}] \quad V_c \sim C \rightarrow dF}{(V_c, V_c') \sim (C, c_{fx}^y) \rightarrow dF; (c_{fx}^y = V_c')}$$

Fig. 11: Target language λ_{ICAL} (semantics of change terms and cache updates).

values since they are now extracted from the caches. Rule [TDRESULT] returns the final change value of a computation as well as a updated cache resulting from the evaluation of the cache update term dC . Rule [TDTUPLE] resembles its counterpart in the source language, but the tuple for y is not built as it has already been pushed in the environment by the cache. As for λ_{AL} , there are three rules to deal with **let**-bindings depending on the shape of the change bound to df in the environment:

- If df is bound to a replacement, the rule [TDREPLACECALL] applies. In that case, we reevaluate the function call in the updated environment $[dF]_2$ (defined similarly as in the source language). This evaluation leads to a new value V' which replaces the original one as well as an updated cache for c_{fx}^y .
- If df is bound to a nil change and f is bound to primitive p , the rule [TD-PRIMITIVE NIL] applies. The derivative of \mathbf{p} is invoked with the value of x , its change value and the cache of the original call to \mathbf{p} . The semantics of \mathbf{p} 's derivative is given by builtin function $\Delta_{\mathbf{p}}(-)$, as in the source language.
- If df is bound to a closure change and f is bound to a closure, the rule [TD-CLOSURECHANGE] applies. The body of the closure change is evaluated under the closure change environment extended with the value of the formal argument x , its change dx and with the environment resulting from the binding of the original cache value to the variables occurring in the cache C . This evaluation leads to a change and an updated cache bound in the environment to continue with the evaluation of the rest of the term.

F Soundness proof of CTS conversion

This section gives complementary technical details about the homonymous section of the paper.

The proof requires a judgment to relate source change environments with CTS change environments that appear during the evaluation of λ_{CAL} change terms *under context*. Assuming that dF' is a change environment with no cache, this judgment is written $dF \vdash dF_0 \uparrow \mathbb{C} \rightarrow_k dF_k$ and is read “Under the change environment dF , the change environment with no cache dF is completed into the change environment dF_k with the caches produced by \mathbb{C} in the base computation if $k = 1$ or in the updated computation if $k = 2$ ”.

The rules for this judgment are given in Fig. 12. [COMPLETE NIL] is a straightforward base case. [COMPLETE TUPLE] does not introduce any cache since only function calls do produce caches. [COMPLETE WITH BASE CALL] (resp. [COMPLETE WITH UPDATED CALL]) computes a cache V_c by evaluating $@(f, x)$ under the base (resp. updated) environment $[dF]_1$ (resp. $[dF]_2$). In [COMPLETE WITH UPDATED CALL], one might wonder why the recursive call which evaluates the remaining context \mathbb{C} is not performed under an environment where $y = V'$. Actually, this is consistent with the fact that in the CTS derivatives base values are only updated at the end of the computation when the updated cache is returned. This is necessary since derivatives calls in \mathbb{C} still need values from the base computation, not the ones from the updated computation.

$$\begin{array}{c}
\text{[COMPLETENIL]} \\
\frac{}{dF \vdash \bullet \uparrow [] \rightarrow_k \bullet} \\
\\
\text{[COMPLETETUPLE]} \\
\frac{dF; y = (\overline{V}); dy = (\overline{dV}) \vdash dF_0 \uparrow \mathbb{C} \rightarrow_k dF_k}{dF \vdash dF_0; y = (\overline{V}); dy = (\overline{dV}) \uparrow \mathbf{let } y = (\overline{x}) \mathbf{in } \mathbb{C} \rightarrow_k dF_k; y = (\overline{V}); dy = (\overline{dV})} \\
\\
\text{[COMPLETEWITHBASECALL]} \\
\frac{[dF]_1 \vdash @(f, x) \Downarrow (V, V_c) \quad dF; y = V; dy = dV; c_{fx}^y = V_c \vdash dF_0 \uparrow \mathbb{C} \rightarrow_1 dF_1}{dF \vdash dF_0; y = V; dy = dV \uparrow \mathbf{let } y = f \mathbf{in } \mathbb{C} \rightarrow_1 dF_1; y = V; dy = dV; c_{fx}^y = V_c} \\
\\
\text{[COMPLETEWITHUPDATEDCALL]} \\
\frac{[dF]_2 \vdash @(f, x) \Downarrow (V', V_c) \quad V \oplus dV = V' \quad dF; y = V; dy = dV; c_{fx}^y = V_c \vdash dF_0 \uparrow \mathbb{C} \rightarrow_2 dF_2}{dF \vdash dF_0; y = V; dy = dV \uparrow \mathbf{let } y = f \mathbf{in } \mathbb{C} \rightarrow_2 dF_2; y = V; dy = dV; c_{fx}^y = V_c}
\end{array}$$

Fig. 12: Completion of a change environment with the caches produced under a context.

Thanks to this judgment, we can state the following key technical lemma whose proof represents a large part of our Coq formalization.

Lemma F.1 (Evaluation of derivatives commutes with evaluation of CTS derivatives).

Let $t = \mathbb{C}[t']$ be a term and dE be a change environment.

If the following hypotheses hold:

1. $dE \vdash \mathcal{D}^t(t) \Downarrow dv$
2. $dE \vdash \mathcal{D}^t(\mathbb{C}) \Downarrow dE_0$
3. $[\mathcal{T}(dE)]_1 \vdash \mathcal{T}(dE_0) \uparrow \mathbb{C} \rightarrow_1 dF_1$
4. $[\mathcal{T}(dE)]_2 \vdash \mathcal{T}(dE_0) \uparrow \mathbb{C} \rightarrow_2 dF_2$
5. $[\mathcal{T}(dE)]_1; dF_1 \vdash \mathcal{T}_t(t') \Downarrow (\mathcal{T}(v), V_c)$
6. $[\mathcal{T}(dE)]_2; dF_2 \vdash \mathcal{T}_t(t') \Downarrow (\mathcal{T}(v'), V'_c)$
7. $\mathcal{C}(t') \sim V_c^{\uparrow|\mathbb{C}|} \rightarrow F_0$

Then the following conclusion holds:

$$\mathcal{T}(dE); F_0; dF_2 \vdash \mathcal{D}_t(t') \Downarrow (\mathcal{T}(dv), V'_c)$$

Let us explain this statement in detail. The hypothesis (1) simply states the existence of a converging computation for the standard derivative of t . The hypothesis (2) makes explicit the change environment dE_0 under which t' is evaluated by using an evaluation judgment for the derivative of \mathbb{C} .⁶ The hypothesis (3) (resp. (4)) completes the CTS version of dE_0 with the caches of the

⁶ The definition of this judgment (resp. the differentiation of term contexts) is omitted, but it mimicks the evaluation of terms (resp. the differentiation of terms).

base (resp. updated) computations. The hypothesis (5) (resp. (6)) evaluates t' under the base (resp. updated) change environment equipped with caches to get the cached values for the whole base (resp. updated) computation. Finally, the hypothesis (7) extracts from V_c only the base cached values that are needed to evaluate $\mathcal{D}_t(t')$: the operation $V_c^{\uparrow|\mathbb{C}|}$ removes the $|\mathbb{C}|$ first elements from the cache V_c . Indeed, these cached values have already been updated by the evaluation of \mathbb{C} and they are present in dF_2 . The conclusion of this lemma states that $\mathcal{T}(dE)$ (the CTS version of dE), extended with the F_0 , the not-yet-updated cache values from the base computation, and with the updated change environment dF_2 , is a sound environment to evaluate the CTS derivative to the CTS version of dv and to the updated cache V'_c .

G Space usage measurements




(a) Benchmark results for space usage of `|average|`.(b) Benchmark results for space usage of `|totalPrice|`.(c) Benchmark results for space usage of `|cartesianProduct|`.

Fig. 13: Benchmark results for space usage.