

Differential Execution with Lexical Tracing

SEBASTIAN ERDWEG, KIT, Germany

RUNQING XU, KIT, Germany

MO BITAR, KIT, Germany

Incremental computing promises large speed-ups after small input edits. Yet, most incrementality approaches merely skip unchanged work and recompute the remaining sub-computations, even when the inputs change only slightly. Differential execution avoids this by propagating data changes (i.e., deltas), and prior work has shown how to develop a provably correct differential big-step semantics. Unfortunately, that semantics must still replay the original computation at every step, squandering much of the potential gain of incrementalization. While the semantics clearly needs caching to avoid recomputations, a sound and efficient caching discipline is challenging. First, each execution step must be uniquely identified; second, the identifier must remain stable even when the preceding control flow changes. To this end, we develop *lexical tracing*, which identifies execution steps through their path in the derivation tree of the big-step semantics. We then extend differential execution with lexical tracing and caching to deliver, for the first time, a formally verified, asymptotically efficient account of differential execution for imperative languages. In particular, we developed a novel mechanized theory of cache stability for lexical traces and their semantic rules, which was essential in proving the differential caching semantics correct and complete in Rocq.

CCS Concepts: • **Theory of computation** → **Design and analysis of algorithms; Semantics and reasoning.**

Additional Key Words and Phrases: incremental computing, operational semantics, mechanized metatheory

ACM Reference Format:

Sebastian Erdweg, Runqing Xu, and Mo Bitar. 2026. Differential Execution with Lexical Tracing. *Proc. ACM Program. Lang.* 10, OOPSLA1, Article 153 (April 2026), 25 pages. <https://doi.org/10.1145/3798261>

1 Introduction

Incremental computations save execution time and energy by reacting to input changes. Most approaches to incremental computing rely on *selective recomputation* [19]: They rerun subcomputations that transitively depend on the changed input, but skip subcomputations whose inputs are unchanged. For example, build systems apply selective recomputation to react to file changes [13], and self-adjusting computations popularized selective recomputation in the PL community [3, 11]. Unfortunately, selective recomputation has one severe problem: Whether an input change is large or small, the computation has to be rerun and cannot exploit algebraic properties of the changed data. For example, if we compute $\max [1, 2, 3, 4, 5]$ and reverse the input list $\max [5, 4, 3, 2, 1]$, selective recomputation has to do a full rerun, wasting time and energy.

An alternative approach to incremental computing is *differential execution* [14], where changes are processed using a dedicated *differential semantics*. Consider an imperative language with a big-step base semantics $s, st \Rightarrow st'$. The corresponding differential semantics $s, \Delta st \Rightarrow_{\Delta} \Delta st'$ executes s relative to a store of value changes and tries to invest minimal effort while being correct: The

Authors' Contact Information: [Sebastian Erdweg](#), KIT, Karlsruhe, Germany; [Runqing Xu](#), KIT, Karlsruhe, Germany; [Mo Bitar](#), KIT, Karlsruhe, Germany.



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

© 2026 Copyright held by the owner/author(s).

ACM 2475-1421/2026/4-ART153

<https://doi.org/10.1145/3798261>

resulting differential store $\Delta st'$ must be such that $s, st \oplus \Delta st \Rightarrow st' \oplus \Delta st'$, which is known as from-scratch consistency. Differential execution provides efficient incrementality because it propagates value changes to primitive operators, which can exploit algebraic properties of the underlying data types. For example, consider a simple program with free variable xs , and an initial and incremental run:

# program uses free variable xs	# initial run	# incremental run
# which must be bound in initial store	$xs = [1,2,3,4,5]$	$\Delta xs = \text{reverse}$
$m := \max xs$	$m = 5$	$\Delta m = \text{noc}$
$x := \text{head } xs$	$x = 1$	$\Delta x = +4$
if $x == m$:	$\rightarrow \text{false}$	$\rightarrow \text{true}$
$ys := \text{tail } xs$		
else :		
$ys := xs$	$ys = [1,2,3,4,5]$	$\Delta ys = \text{drop } 1 \circ \text{reverse}$

In the initial run of the program, we set $xs = [1,2,3,4,5]$ and then compute its maximum and head element before selecting the else branch. More interesting is the incremental run using differential execution, where we set Δxs to be the reverse of xs . This information is stored in the differential store Δst and retrieved whenever we access variable xs . In particular, $\max xs$ yields *noc* (no change) when xs is reversed, because the order of list elements does not affect the maximum. The primitive differential maximum operator can exploit this property because it knows *how* its input changed. The *head* operator needs to compare the old head element to the new head element, and finds that it increased by 4. And if the equation comparator remembers the difference between the old x and m (which was -4), it can quickly determine that the new head element is equal to the unchanged maximum. Hence, the differential run will choose the then branch, whose result differs from the original ys by first reversing it and then dropping one element. We can then compositionally feed this change of ys to subsequent computations to continue differential execution.

Differential execution can yield asymptotic speedups compared to a from-scratch computation. However, we skipped over a crucial detail: caching. Caching is needed in differential execution for two reasons. First, the primitive operators often need to remember information about past executions. We saw this in the *head* and equality operators, which needed to know the past head element and the difference of the previously compared values. And even though *max* could react to *reverse* without prior information, inserting a list element requires *max* to compare the new element to the previous maximum. Xu and Erdweg [22] provide a detailed account of stateful differential operators, yet they do not embed them in a language and do not explain how the cache can be managed during execution. Second, control structures need to remember information about the past execution state. In our example, the conditional performed a branch switch in the incremental run: We switched to the then branch, whose body we ran in the state that was valid before the original conditional. Similarly, the differential execution of loops requires cached information. Indeed, differential execution without caching is pointless.

While caching is crucial for the correctness and efficiency of differential execution, there is no formalization of it yet. In this paper, we formalize the caching behavior of differential execution by solving two key technical challenges. The first challenge is how to identify cache entries in a way that every operator and control structure obtains their distinct entry, which is critical for correctness. The second challenge is how to ensure cache entries remain as stable as possible when processing changes, which is critical for efficiency. To this end, we develop the theory of lexical traces, which uniquely identify points in the program execution in a stable fashion, as we explain in the next section. We use lexical traces to develop a tracing semantics for the original and differential execution, where a cache is propagated in store-passing style. We then provide a new cache-aware

correctness criterion for differential execution:

$$\begin{aligned} s, st, ch_0 @ tr &\Rightarrow st', ch_1 \rightarrow \\ s, \Delta st, ch_1 @ tr &\Rightarrow_{\Delta} \Delta st', ch_2 \rightarrow \\ s, st \oplus \Delta st, ch_0 @ tr &\Rightarrow st' \oplus \Delta st', ch_2 \end{aligned}$$

Essentially, this is from scratch consistency: If we rerun s with the patched input store $st \oplus \Delta st$, we get the same result as if we had patched the output store st' using the differential result $\Delta st'$. However, we include lexical traces tr and caches ch to model efficient differential execution. To prove correctness, then, we will also need to reason about cache equivalence and cache stability extensively. We formalized lexical tracing and differential caching execution for a small imperative language in Rocq and provide a mechanized proof of its correctness and completeness.

In summary, we make the following contributions:

- We present lexical tracing, an approach to uniquely identify points in the program execution that are stable under subcomputation changes.
- We develop caching semantics for the original and differential execution using lexical traces.
- We develop a theory of local cache equality and cache stability in Rocq, which is elemental for proving the correctness and completeness of the caching differential semantics.
- We implement a differential interpreter based on our principled caching discipline and demonstrate that its performance is on par with a prior ad-hoc implementation.


2 Incremental Caching and Lexical Tracing

Differential execution must cache information from the previous run to react to input changes correctly and efficiently. For example, the differential multiplication operator is an instructive example, even though it may be faster to simply recompute multiplications. When we change the operands of a multiplication $x*y$ with $\Delta x = +a$ and $\Delta y = +b$, the output change is $+(a*y+a*b+x*b)$, which we can derive if we remember x and y . However, different instances of the same operator must each remember their own operands. For example, the cache may not confuse the multiplications in s_1 and s_5 :

# program with	# initial run / cache	# incremental run / cache
# free variable input	input = -2	Δ input = +3
s_1 : $x := 2 * \text{input}$	$x = -4 / (2, -2)$	$\Delta x = +2*3 = +6 / (2, 1)$
s_2 : if $x > 0$:	\rightarrow false	\rightarrow true
s_3 : $y := x + 1$		$\Delta y = +5$
else :		
s_4 : $y := x + 2$	$y = -2$	
s_5 : $r := x * y$	$r = 8 / (-4, -2)$	$\Delta r = +(6*-2 + 6*5 + 5*-4) = -2 / (2, 3)$

Say, we run this program with `input = -2` originally. The multiplication in s_1 must remember its operands 2 and -2, whereas the multiplication in s_5 must remember -4 and -2. Only then will both multiplications be able to react to changes to its respective operands. When we increase `input` by +3, the operands in s_1 change by +0 and +3. Since the old left-operand was 2, the multiplication result increases by $+2*3=+6$ from -4 to 2. Subsequently, this leads to changes of the operands in s_5 by +6 and +5. Since the cached operands are -4 and -2, the result changes by -2 following the formula from above. Note that we also updated cache entries in the incremental run, so that we can process subsequent changes if needed.

We need to uniquely identify operators during the execution of the program. Unfortunately, the position of the operator is insufficient, because operators can occur in loops and in recursive functions, yet each application of an operator needs a distinct cache entry. One way to identify operator applications is by using the program trace that led to the operator's execution. For example,



Examples of lexical traces and the nodes they resolve to:

$\text{seq}_L \cdot \text{assign}_{\text{rhs}}$	\mapsto	$2 * \text{input}$
$\text{seq}_R \cdot \text{seq}_L \cdot \text{if}_{\text{cond}}$	\mapsto	$x > 0$
$\text{seq}_R \cdot \text{seq}_L \cdot \text{if}_{\text{else}}$	\mapsto	$y := x+2$
$\text{seq}_R \cdot \text{seq}_R \cdot \text{assign}_{\text{rhs}}$	\mapsto	$x * y$

Fig. 1. An execution tree gives rise to lexical traces for each individual node.

in the initial run, the multiplication in s_5 is reached through trace $s_1 \cdot s_2 \cdot s_4 \cdot s_5$. Indeed, this trace clearly distinguishes the operators in s_1 and s_5 . However, the trace lacks stability and leads to a huge number of cache misses. For example, in the incremental run, we reach s_5 through $s_1 \cdot s_2 \cdot s_3 \cdot s_5$. Since the two traces are different, we would not find the cache entry and would need to perform a recomputation for the second execution of s_5 . Indeed, any change to the control flow preceding a statement will affect the statement's trace. Clearly, this fails our objectives for incremental computing. We propose lexical tracing to regain trace stability.

2.1 Execution Trees and Lexical Tracing

Traditional tracing records the sequence of executed statements. The key idea of lexical tracing is to instead record the path in the big-step derivation tree that leads from the root to a given statement. For our previous example, the following derivation tree models the initial run in a big-step semantics:


$$\begin{array}{c}
 \frac{2 * \text{input}, st_0 \Rightarrow -4}{x := 2 * \text{input}, st_0 \Rightarrow st_1} \quad \frac{\frac{x > 0, st_1 \Rightarrow \text{false} \quad \frac{x+2, st_1 \Rightarrow -2}{y := x+2, st_1 \Rightarrow st_2}}{\text{if } (x > 0) (y := x+1) (y := x+2), st_1 \Rightarrow st_2} \quad \frac{x * y, st_2 \Rightarrow 8}{r := x * y, st_2 \Rightarrow st_3}}{\text{if } (x > 0) (y := x+1) (y := x+2); r := x * y, st_1 \Rightarrow st_3} \\
 \hline
 x := 2 * \text{input}; \text{if } (x > 0) (y := x+1) (y := x+2); r := x * y, st_0 \Rightarrow st_3
 \end{array}$$

where $st_1 = st_0; x = -4 \quad st_2 = st_1; y = -2 \quad st_3 = st_2; r = 8$

While we can extract lexical traces from the derivation, it contains a lot of information irrelevant to tracing. Therefore, it is instructive to simplify the derivation tree into what we call an execution tree: A tree that records the execution's structure but ignores the results (here: the values and stores). We represent execution trees in a style similar to abstract syntax trees, with the starting point at the top. However, the execution tree is a projection of the derivation tree and thus only contains the code that was reached during execution.

For our example, we show the execution tree of the initial run in Figure 1. Note that the execution tree only includes the else branch but omits the unreached then branch. The execution tree makes it easy to uniquely identify executed statements and expressions through their path from the root node. For our example, we show a few paths in Figure 1 with the nodes they resolve to. These are the *lexical traces* of the program.

Lexical traces are partial: They do not record the full history of execution, but only identify the current statement or expression. This is exactly what we need when using lexical traces for caching. Specifically, the key advantage of lexical traces is that they are stable with respect to the execution of non-overlapping subcomputations. For example, the trace $\text{seq}_R \cdot \text{seq}_R$ for the assignment of r is independent of the execution of the preceding statements at seq_L and $\text{seq}_R \cdot \text{seq}_L$. Even in the



Examples of lexical traces and the nodes they resolve to:

$$\begin{aligned} \text{seq}_R \cdot \text{rep}_{\text{body}} &\mapsto r := r * x \\ \text{seq}_R \cdot \text{rep}_{\text{iter}} \cdot \text{rep}_{\text{body}} &\mapsto r := r * x \\ \text{seq}_R \cdot \text{rep}_{\text{iter}}^k &\mapsto \text{repeat } (3 - k) \end{aligned}$$

Fig. 2. The execution tree for a loop witnesses the loop unfolding as done by the big-step semantics.

incremental run, when the execution of the if statement at $\text{seq}_R \cdot \text{seq}_L$ changes from the else to the then branch, the trace $\text{seq}_R \cdot \text{seq}_R$ still identifies the assignment of r . This stability of lexical traces is elemental for using them to identify nodes in a cache.

Before moving on, we should take a brief moment to consider lexical traces for programs that have loops. When a loop body executes multiple times, we need to distinguish operator applications in each iteration. Lexical traces provide just that: When we record the path in the big-step derivation or execution tree, we obtain distinct lexical traces for each loop iteration. For example, we can compute the third power of x using $r := 1; \text{repeat } 3 (r := r * x)$, which produces the execution tree shown in Figure 2. The execution tree shows the unfolded loop, providing distinct lexical traces to reach each execution of the loop body. This way, we can cache separate entries for all three applications of the multiplication operator in $r := r * x$. Here and in the remainder of this paper we write step^k for the k -fold repetition of step.

2.2 A Semantics with Lexical Tracing

Before turning to differential execution and its caching, we want to show how lexical tracing can be incorporated into a language's semantics. To this end, we provide a big-step semantics for a simple imperative language with lexical tracing. We will later add caching to the semantics, where we use the lexical traces to identify cache entries.

We show the syntax and semantics of the simple imperative language IMP in Figure 3. IMP features mutable variables, conditional statements, and repeat loops ($\text{rep } e \text{ s}$). Note that we employ an auxiliary syntactic form $\text{rep}^n s$ for loops with n remaining iterations (cf. Section 3). We use this information later as part of our caching strategy for differential execution. A lexical trace tr is a sequence of trace links l , describing paths in the implicit execution tree. The semantics of expressions e , $st @ tr \Rightarrow v$ and statements s , $st @ tr \Rightarrow st'$ expects the current trace tr as an additional input. Other than that, the semantics is completely standard. The only novelty is that the semantics extends the trace as it recurses. Thus, the current trace tr always explains how execution proceeded from the root of the program to the current expression or statement.

Let us go through a few exemplary rules. For binary operators $e_1 \text{ op } e_2$, we evaluate e_1 under $tr \cdot \text{op}_L$ and e_2 under $tr \cdot \text{op}_R$. The same happens for sequential statements $s_1 ; s_2$. For conditionals $\text{if } e \text{ s}_1 \text{ s}_2$, we evaluate e under $tr \cdot \text{if}_{\text{cond}}$ and then decide which branch to execute. The branch selection is evident in the sub-trace used recursively. For loops $\text{rep } e \text{ s}$, we evaluate e to determine how many iterations n are necessary. Then we execute $\text{rep}^n s$, which repeats the body of the loop n times using rep_{iter} to distinguish loop iterations in the trace.

We will later reason about traces to prove that independent subcomputations leave their respective cache entries intact. A key property that we will use is trace independence.

Definition 2.1 (Trace independence $tr_1 \# tr_2$). Lexical traces tr_1 and tr_2 are independent $tr_1 \# tr_2$ if neither trace is a prefix of the other: $tr_1 \cdot tr \neq tr_2$ and $tr_1 \neq tr_2 \cdot tr$ for all tr .

(values)	$v ::= \text{true} \mid \text{false} \mid n$
(expressions)	$e ::= v \mid x \mid e + e \mid e * e \mid e > e$
(statements)	$s ::= \text{skip} \mid x := e \mid s ; s \mid \text{if } e \text{ } s \text{ } s \mid \text{rep } e \text{ } s \mid \text{rep}^n s$
(stores)	$st \in x \rightarrow v$
(trace links)	$l ::= +_L \mid +_R \mid *_L \mid *_R \mid >_L \mid >_R$ $\mid \text{assign}_{\text{rhs}} \mid \text{seq}_L \mid \text{seq}_R \mid \text{if}_{\text{cond}} \mid \text{if}_{\text{then}} \mid \text{if}_{\text{else}}$ $\mid \text{rep}_{\text{count}} \mid \text{rep}_{\text{aux}} \mid \text{rep}_{\text{body}} \mid \text{rep}_{\text{iter}}$
(lexical traces)	$tr ::= l \cdots l$

$$\begin{array}{c}
\frac{}{x, st @ tr \Rightarrow st x} \quad \frac{e_1, st @ tr \cdot \text{op}_L \Rightarrow v_1 \quad e_2, st @ tr \cdot \text{op}_R \Rightarrow v_2}{e_1 \text{ op } e_2, st @ tr \Rightarrow v_1 \text{ op } v_2} \\
\\
\frac{}{\text{skip}, st @ tr \Rightarrow st} \quad \frac{e, st @ tr \cdot \text{assign}_{\text{rhs}} \Rightarrow v}{x := e, st @ tr \Rightarrow st[x \mapsto v]} \quad \frac{s_1, st @ tr \cdot \text{seq}_L \Rightarrow st_1 \quad s_2, st_1 @ tr \cdot \text{seq}_R \Rightarrow st_2}{s_1 ; s_2, st @ tr \Rightarrow st_2} \\
\\
\frac{e, st @ tr \cdot \text{if}_{\text{cond}} \Rightarrow \text{true} \quad s_1, st @ tr \cdot \text{if}_{\text{then}} \Rightarrow st_1}{\text{if } e \text{ } s_1 \text{ } s_2, st @ tr \Rightarrow st_1} \quad \frac{e, st @ tr \cdot \text{if}_{\text{cond}} \Rightarrow \text{false} \quad s_2, st @ tr \cdot \text{if}_{\text{else}} \Rightarrow st_2}{\text{if } e \text{ } s_1 \text{ } s_2, st @ tr \Rightarrow st_2} \\
\\
\frac{e, st @ tr \cdot \text{rep}_{\text{count}} \Rightarrow n \quad \text{rep}^n s, st @ tr \cdot \text{rep}_{\text{aux}} \Rightarrow st'}{\text{rep } e \text{ } s, st @ tr \Rightarrow st'} \quad \frac{}{\text{rep}^0 s, st @ tr \Rightarrow st} \quad \frac{s, st @ tr \cdot \text{rep}_{\text{body}} \Rightarrow st' \quad \text{rep}^n s, st' @ tr \cdot \text{rep}_{\text{iter}} \Rightarrow st''}{\text{rep}^{n+1} s, st @ tr \Rightarrow st''}
\end{array}$$

Fig. 3. Big-step semantics for IMP expressions and statements with lexical tracing.

For example, we have $\text{seq}_R \cdot \text{seq}_L \# \text{seq}_R \cdot \text{seq}_R$ in our running example from Figure 1. We will later see that this ensures that the (differential) execution of the if statement at $\text{seq}_R \cdot \text{seq}_L$ cannot affect the cached entries for the assignment at $\text{seq}_R \cdot \text{seq}_R$. Similarly, for the repeat example from Figure 2, we have $\text{seq}_R \cdot \text{rep}_{\text{iter}}^k \cdot \text{rep}_{\text{body}} \# \text{seq}_R \cdot \text{rep}_{\text{iter}}^m \cdot \text{rep}_{\text{body}}$ for any $k \neq m$. This ensures that each loop iteration can use the cache safely. We are now ready to formalize the differential semantics of IMP with caching.

3 Differential Semantics with Caching

Differential semantics was recently proposed as a new approach for incremental computing [14]. A differential semantics executes a program and processes input changes to produce output changes. For example, the differential semantics for IMP expressions should take the form $e, \Delta st \Rightarrow_{\Delta} \Delta v$, where $\Delta st : \text{var} \rightarrow \Delta v$ maps variables to their value change. However, we cannot possibly implement a differential semantics of this form, because we need information about the original run. To illustrate, consider again the multiplication expression and its standard semantics:

$$\frac{e_1, st \Rightarrow m \quad e_2, st \Rightarrow n}{e_1 * e_2, st \Rightarrow m * n}$$

If we attempt a differential semantics for multiplication without caching, we necessarily fail:

$$\frac{e_1, \Delta st \Rightarrow_{\Delta} +a \quad e_2, \Delta st \Rightarrow_{\Delta} +b}{e_1 * e_2, \Delta st \Rightarrow_{\Delta} +(a * n + a * b + m * b)}$$

$$\begin{array}{c}
\frac{}{n, \Delta st, ch @ tr \Rightarrow_{\Delta} +0, ch} \quad (\Delta \text{NUM}) \qquad \frac{b \in \{\text{true}, \text{false}\}}{b, \Delta st, ch @ tr \Rightarrow_{\Delta} \text{noc}, ch} \quad (\Delta \text{BOOL}) \\
\\
\frac{}{x, \Delta st, ch @ tr \Rightarrow_{\Delta} \Delta st x, ch} \quad (\Delta \text{VAR}) \qquad \frac{e_1, \Delta st, ch_0 @ tr \cdot +_L \Rightarrow_{\Delta} +a, ch_1 \quad e_2, \Delta st, ch_1 @ tr \cdot +_R \Rightarrow_{\Delta} +b, ch_2}{e_1 + e_2, \Delta st, ch_0 @ tr \Rightarrow_{\Delta} +(a + b), ch_2} \quad (\Delta \text{ADD}) \\
\\
\frac{e_1, \Delta st, ch_0 @ tr \cdot *_L \Rightarrow_{\Delta} +a, ch_1 \quad e_2, \Delta st, ch_1 @ tr \cdot *_R \Rightarrow_{\Delta} +b, ch_2 \quad ch_0 tr = \text{mul } m n \quad ch_3 = ch_2[tr \mapsto \text{mul}(m + a)(n + b)]}{e_1 * e_2, \Delta st, ch_0 @ tr \Rightarrow_{\Delta} +(a * n + a * b + m * b), ch_3} \quad (\Delta \text{MUL}) \\
\\
\frac{e_1, \Delta st, ch_0 @ tr \cdot >_L \Rightarrow_{\Delta} +a, ch_1 \quad e_2, \Delta st, ch_1 @ tr \cdot >_R \Rightarrow_{\Delta} +b, ch_2 \quad ch_0 tr = \text{gt } d \quad ch_3 = ch_2[tr \mapsto \text{gt}(d + a - b)]}{e_1 > e_2, \Delta st, ch_0 @ tr \Rightarrow_{\Delta} (d + a - b > 0) \ominus (d > 0), ch_3} \quad (\Delta \text{GT})
\end{array}$$

Fig. 4. The differential caching semantics for IMP expressions with threaded caches ch .

This differentially evaluates e_1 and e_2 to changes $+a$ and $+b$. The problem is that we cannot describe how the multiplication result changes without referring to the previous operands m and n . In the original introduction of differential semantics, Kumar et al. [14] sidestepped this issue by providing the original store to the differential semantics. Their rule for multiplication reads:

$$\frac{e_1, st \Rightarrow m \quad e_2, st \Rightarrow n \quad e_1, \Delta st \Rightarrow_{\Delta} +a \quad e_2, \Delta st \Rightarrow_{\Delta} +b}{e_1 * e_2, st, \Delta st \Rightarrow_{\Delta} +(a * n + a * b + m * b)}$$

This rule is correct and executable, but very inefficient because it reevaluates e_1 and e_2 given the original store st . Indeed, almost all reduction rules of Kumar et al. [14] have to do some recomputation, because their semantics lacks caching. But caching is difficult to incorporate: The caching semantics is complex, it requires stable cache keys, and reasoning about cache correctness requires a lot of proof engineering. In the remainder of this section, we present our differential caching semantics based on lexical tracing.

3.1 Differential Caching Semantics for IMP Expressions

The differential caching semantics for IMP expressions takes the form

$$e, \Delta st, ch @ tr \Rightarrow_{\Delta} \Delta v, ch'.$$

That is, it takes an expression, a differential store, a cache ch , and a trace tr , and it produces a change value Δv and a possibly updated cache ch' . Caches $ch : tr \rightarrow \eta$ map traces to cache entries η , which is a variant type for all different kinds of cache entries required by IMP. For IMP expressions, we need variants mul and gt to support differential multiplication and greater-than.

We show the differential caching semantics in Figure 4, which uses the same tracing as in Section 2.2. The differential evaluation of numeric and Boolean constants yields no change. For variables, (ΔVAR) looks up the change in the differential store $\Delta st \in x \rightarrow \Delta v$. Additions evaluate as usual, but (ΔADD) makes sure to propagate the cache in store-passing style. The remaining two rules are the interesting ones, because we need caching for multiplication and greater-than.

The differential evaluation of multiplications is handled by (ΔMUL) . As usual, (ΔMUL) recursively evaluates the operands e_1 and e_2 to obtain their changes $+a$ and $+b$. But to produce the result

change $+(a * n + a * b + m * b)$, we also need to know the value of the original operands m and n . This is where the cache comes into play: We retrieve m and n from the cache, which must provide a `mul` entry for the multiplication expression $e_1 * e_2$. We use the lexical trace tr to identify $e_1 * e_2$ and distinguish it from other multiplication expressions in the program. Finally, we must also update the cache to prepare for any subsequent input changes that may occur. To this end, we store the updated operands in ch_3 and yield it as part of the output.

The differential evaluation of grater-than expressions in (ΔGT) has a similar form, but differs in its details. First, for differential grater-than, it suffices to remember the integer distance $d = m - n$ between the original operands, rather than remembering the operands themselves. Second, we update the cache entry to reflect the distance of the changed operands: $(m + a) - (n + b) = (m - n) + (a - b) = d + a - b$. Third, we compute the output change by diffing the Boolean values $new \ominus old$, which yields a Boolean change `noc` or `neg`. For example, $5 > 3$ induces $d = 2$, and $d > 0$ is true. If we now get operand changes $a = +1$ and $b = +4$, then $d + a - b = -1$ and the new result is false. Differential evaluation then yields the result change $false \ominus true = \text{neg}$.

The careful reader may have noticed that the differential evaluation is stuck unless the input cache is already initialized. Therefore, we need a specialized initialization semantics, which behaves like the standard big-step semantics, but also initializes the cache for subsequent differential runs. Indeed, this initialization semantics is equivalent to the tracing from Figure 3, but additionally threads a cache ch in store-passing style. Moreover, when the differential semantics requires cache entries, the initializing semantics must provide them. For expressions, the evaluation rules for multiplication and greater-than expressions must initialize the cache:

$$\frac{\begin{array}{l} e_1, st, ch_0 @ tr \cdot *_L \Rightarrow m, ch_1 \\ e_2, st, ch_1 @ tr \cdot *_R \Rightarrow n, ch_2 \\ ch_3 = ch_2 [tr \mapsto \text{mul } m n] \end{array}}{e_1 * e_2, st, ch_0 @ tr \Rightarrow m * n, ch_3} \text{ (MUL)} \quad \frac{\begin{array}{l} e_1, st, ch_0 @ tr \cdot >_L \Rightarrow m, ch_1 \\ e_2, st, ch_1 @ tr \cdot >_R \Rightarrow n, ch_2 \\ ch_3 = ch_2 [tr \mapsto \text{gt } (m - n)] \end{array}}{e_1 > e_2, st, ch_0 @ tr \Rightarrow (m - n > 0), ch_3} \text{ (GT)}$$

We will later prove in Section 4 that differential evaluation of IMP expressions is complete and correct. Completeness ensures that differential evaluation succeeds whenever the initial evaluation succeeds. In particular, the initial evaluation provides cache entries whenever the differential evaluation requires them. Correctness means that differential evaluation yields results that match a from-scratch re-evaluation of the program with updated inputs. In particular, this guarantees that the initial evaluation provides correct cache entries. However, the actual proofs are complex and require theories of cache equivalence and cache stability; we dedicate Section 4 to the metatheory.

3.2 Differential Caching Semantics for Assignments and Conditionals

The differential semantics for statements is more involved, because control structures are inherently difficult to handle incrementally. The fundamental problem is *branch switching*: When an input change triggers the execution of new code that was not originally executed. In particular, when an input change flips the condition of an if statement, we must switch from one branch to the other. Here, differential execution is out of options and the only way forward is a from-scratch execution of the new code. This requires caching: We must remember the original execution state so that we can restore it for the from-scratch execution of the new branch. A similar problem occurs for loops, where we may have to execute more or less iterations, which also requires caching.

We show the differential caching semantics for statements in Figure 5. The differential execution judgment takes the form

$$s, \Delta st, ch @ tr \Rightarrow_{\Delta} \Delta st', ch'$$

$$\begin{array}{c}
\frac{}{\text{skip}, \Delta st, ch @ tr \Rightarrow_{\Delta} \Delta st, ch} \quad (\Delta\text{SKIP}) \qquad \frac{e, \Delta st, ch @ tr \cdot \text{assign}_{\text{rhs}} \Rightarrow_{\Delta} \Delta v, ch'}{x := e, \Delta st, ch @ tr \Rightarrow_{\Delta} \Delta st[x \mapsto \Delta v], ch'} \quad (\Delta\text{ASSIGN}) \\
\\
\frac{s_1, \Delta st, ch @ tr \cdot \text{seq}_L \Rightarrow_{\Delta} \Delta st_1, ch_1 \quad s_2, \Delta st_1, ch_1 @ tr \cdot \text{seq}_R \Rightarrow_{\Delta} \Delta st_2, ch_2}{s_1 ; s_2, \Delta st, ch @ tr \Rightarrow_{\Delta} \Delta st_2, ch_2} \quad (\Delta\text{SEQ}) \\
\\
\begin{array}{c}
ch_0 tr = \text{if true } st \ st' \\
e, \Delta st, ch_0 @ tr \cdot \text{if}_{\text{cond}} \Rightarrow_{\Delta} \text{noc}, ch_1 \\
s_1, \Delta st, ch_1 @ tr \cdot \text{if}_{\text{then}} \Rightarrow_{\Delta} \Delta st', ch_2 \\
ch_3 = ch_2[tr \mapsto \text{if true } (st \oplus \Delta st) (st' \oplus \Delta st')]
\end{array} \qquad \begin{array}{c}
ch_0 tr = \text{if false } st \ st' \\
e, \Delta st, ch_0 @ tr \cdot \text{if}_{\text{cond}} \Rightarrow_{\Delta} \text{noc}, ch_1 \\
s_2, \Delta st, ch_1 @ tr \cdot \text{if}_{\text{else}} \Rightarrow_{\Delta} \Delta st', ch_2 \\
ch_3 = ch_2[tr \mapsto \text{if false } (st \oplus \Delta st) (st' \oplus \Delta st')]
\end{array} \\
\hline
\text{if } e \ s_1 \ s_2, \Delta st, ch_0 @ tr \Rightarrow_{\Delta} \Delta st', ch_3 \quad (\Delta\text{IF-TT}) \qquad \text{if } e \ s_1 \ s_2, \Delta st, ch_0 @ tr \Rightarrow_{\Delta} \Delta st', ch_3 \quad (\Delta\text{IF-FF}) \\
\\
\begin{array}{c}
ch_0 tr = \text{if true } st \ st' \\
e, \Delta st, ch_0 @ tr \cdot \text{if}_{\text{cond}} \Rightarrow_{\Delta} \text{neg}, ch_1 \\
s_2, st \oplus \Delta st, ch_1 @ tr \cdot \text{if}_{\text{else}} \Rightarrow_{\Delta} st'', ch_2 \\
\Delta st' = st'' \ominus st' \\
ch_3 = ch_2[tr \mapsto \text{if false } (st \oplus \Delta st) st'']
\end{array} \qquad \begin{array}{c}
ch_0 tr = \text{if false } st \ st' \\
e, \Delta st, ch_0 @ tr \cdot \text{if}_{\text{cond}} \Rightarrow_{\Delta} \text{neg}, ch_1 \\
s_1, st \oplus \Delta st, ch_1 @ tr \cdot \text{if}_{\text{then}} \Rightarrow_{\Delta} st'', ch_2 \\
\Delta st' = st'' \ominus st' \\
ch_3 = ch_2[tr \mapsto \text{if true } (st \oplus \Delta st) st'']
\end{array} \\
\hline
\text{if } e \ s_1 \ s_2, \Delta st, ch_0 @ tr \Rightarrow_{\Delta} \Delta st', ch_3 \quad (\Delta\text{IF-TF}) \qquad \text{if } e \ s_1 \ s_2, \Delta st, ch_0 @ tr \Rightarrow_{\Delta} \Delta st', ch_3 \quad (\Delta\text{IF-FT})
\end{array}$$

Fig. 5. The differential semantics for IMP statements uses the cache for conditionals. We highlight key case distinctions and mark premises that rely on from-scratch execution.

For each statement and differential store, the semantics computes a differential store that describes how the original output changes. Skip statements have no effect on the output and hence do not change it. For assignments $x := e$, (ΔASSIGN) uses differential evaluation of expressions to determine how e changes, and then puts this change in the differential store. For example, when x is increased by one, $y := 2 * x$ differentially evaluates independent of the original store:

$$y := 2 * x, (x \mapsto +1), ch @ tr \Rightarrow_{\Delta} (x \mapsto +1, y \mapsto +2), ch'.$$

That is, when x increases by 1, y increases by 2. When statements occur in sequence, rule (ΔSEQ) executes them subsequently and threads the differential store and cache as usual. Indeed, the first three rules do not involve conditional control structures and are simple adaptations from the tracing semantics of IMP from Section 2.2. The next rules are more interesting.

For if statements, we have to distinguish whether a branch switch occurs or not. For example, when x changes from true to false in (if x ($z := y + 1$) ($z := y * 2$)), we must switch from the then to the else branch. To this end, we must remember the store before the if statement to execute the else branch from scratch. Moreover, to compute how the else-branch result differs from the previous then-branch result, we also need to remember the store after the if statement, so we can compare the two resulting stores. Our differential caching semantics makes this behavior explicit and, of course, also updates the cache to keep it up-to-date.

In fact, we have four rules to deal with if statements, depending on which branch was originally executed and whether a branch switch occurs. The cache entry (if b st st') indicates whether the condition was previously $b = \text{true}$ or $b = \text{false}$, and which store we had before (st) and after (st') executing the original if statement. When the differential evaluation of the condition yields $e, \dots \Rightarrow_{\Delta} \text{noc}$ (no change), we remain at the same branch as before. For $b = \text{true}$, rule ($\Delta\text{IF-TT}$)

remains at the then branch and uses differential execution of s_1 to obtain the output change $\Delta st'$. Conversely, for $b = \text{false}$, rule $(\Delta\text{IF-FF})$ remains at the else branch and uses differential execution of s_2 . But before we can yield the output change, we must restore the consistency of the cache by updating the store before and after the execution of the if statement. We need them when a branch switch occurs.

A branch switch occurs when the differential evaluation of the condition yields $e, \dots \Rightarrow_{\Delta} \text{neg}$ (negate). In rule $(\Delta\text{IF-TF})$, we switch from $b = \text{true}$ to false, whereas $(\Delta\text{IF-FT})$ switches the other way around. As differential execution switches from one branch to the other, we have to execute the new branch from scratch, because differential execution requires a previous initializing execution. Indeed, rule $(\Delta\text{IF-TF})$ executes else branch s_2 with original store st from scratch, and $(\Delta\text{IF-FT})$ executes s_1 from scratch, as highlighted in Figure 5. But rather than a differential store, this yields an output store st'' . Therefore, we diff (\ominus) the new store st'' and the original output store st' to obtain a differential store, which describes the observable changes of the if statement. Lastly, we update the cache, making sure to flip the condition b and update the stores.

The differential semantics requires the cache to be initialized for conditionals. Therefore, the initial run of the program must prepare the cache, which is straightforward:

$$\frac{\begin{array}{l} e, st, ch_0 @ tr \cdot \text{if}_{\text{cond}} \Rightarrow \text{true}, ch_1 \\ s_1, st, ch_1 @ tr \cdot \text{if}_{\text{then}} \Rightarrow st_1, ch_2 \\ ch_3 = ch_2 [tr \mapsto \text{if true } st \ st_1] \end{array}}{\text{if } e \ s_1 \ s_2, st, ch_0 @ tr \Rightarrow st_1, ch_3} \quad \frac{\begin{array}{l} e, st, ch_0 @ tr \cdot \text{if}_{\text{cond}} \Rightarrow \text{false}, ch_1 \\ s_2, st, ch_1 @ tr \cdot \text{if}_{\text{else}} \Rightarrow st_2, ch_2 \\ ch_3 = ch_2 [tr \mapsto \text{if false } st \ st_2] \end{array}}{\text{if } e \ s_1 \ s_2, st, ch_0 @ tr \Rightarrow st_2, ch_3}$$

The following program demonstrates the differential execution of conditionals:

```
# free variable n          n = 3          ch = { ifcond ↦ gt 3,
if n > 0                   ifthen · assignrhs ↦ mul 5 3
  r := 5 * n                ε ↦ if true {n ↦ 3} {n ↦ 3, r ↦ 15} }
else:
  r := 2 * n                Δn = -5      ch' = { ifcond ↦ gt (-2),
                             ifthen · assignrhs ↦ mul 5 3
                             ifelse · assignrhs ↦ mul 2 (-2)
                             ε ↦ if false {n ↦ -2} {n ↦ -2, r ↦ -4}
```

The initial run uses $n = 3$. We show the cache ch after the initial run, which contains three entries: (i) `gt` for the if condition, which is true, (ii) `mul` for the multiplication in the then branch, and (iii) `if` with the stores before and after the conditional. Now, consider we make a change $\Delta st = \{\Delta n \mapsto -5\}$ and run it through the differential semantics. The differential semantics updates the cached entry for `gt` and negates the if condition. Hence, we trigger $(\Delta\text{IF-TF})$ and re-execute the else branch using the cached store, which introduces a new cache entry for its own multiplication at `ifelse · assignrhs`. Finally, we update the cache and diff the resulting store $\{r \mapsto (-4)\}$ with the previous store $\{r \mapsto 15\}$ to deduce $\Delta r = -19$. Note that we currently do not clear cache entries, but could have removed the entry for `ifthen · assignrhs` after the branch switch.

3.3 Differential Caching Semantics for Loops

For loops, there are various possible caching strategies that affect the performance/memory trade-off of incremental computing. For this paper, we chose a memory-economical caching strategy that is efficient when the number of iterations stays the same or grows, but inefficient for loop rollbacks. Specifically, we only cache the store from before and after the loop, but no intermediate stores from individual loop iterations. Therefore, all cache handling occurs for the $(\text{rep } e \ s)$ construct,

$$\begin{array}{c}
\frac{}{\text{rep}^0 s, \Delta st, ch @ tr \Rightarrow_{\Delta} \Delta st, ch} \quad (\Delta\text{REP-0}) \qquad \frac{s, \Delta st_0, ch_0 @ tr \cdot \text{rep}_{\text{body}} \Rightarrow_{\Delta} \Delta st_1, ch_1 \quad \text{rep}^n s, \Delta st_1, ch_1 @ tr \cdot \text{rep}_{\text{iter}} \Rightarrow_{\Delta} \Delta st_2, ch_2}{\text{rep}^{n+1} s, \Delta st_0, ch_0 @ tr \Rightarrow_{\Delta} \Delta st_2, ch_2} \quad (\Delta\text{REP+1}) \\
\\
\frac{\begin{array}{l} ch_0 tr = \text{rep } n \text{ st } st' \\ e, \Delta st, ch_0 @ tr \cdot \text{rep}_{\text{count}} \Rightarrow_{\Delta} \pm 0, ch_1 \\ \text{rep}^n s, \Delta st, ch_1 @ tr \cdot \text{rep}_{\text{aux}} \Rightarrow_{\Delta} \Delta st', ch_2 \\ ch_3 = ch_2[tr \mapsto \text{rep } n (st \oplus \Delta st) (st' \oplus \Delta st')] \end{array}}{\text{rep } e s, \Delta st, ch_0 @ tr \Rightarrow_{\Delta} \Delta st', ch_3} \quad (\Delta\text{REP-NOC}) \qquad \frac{\begin{array}{l} ch_0 tr = \text{rep } n \text{ st } st' \\ e, \Delta st, ch_0 @ tr \cdot \text{rep}_{\text{count}} \Rightarrow_{\Delta} -m, ch_1 \\ \text{rep}^{n-m} s, st \oplus \Delta st, ch_1 @ tr \cdot \text{rep}_{\text{aux}} \Rightarrow_{\Delta} st'', ch_2 \\ \Delta st' = st'' \ominus st' \\ ch_3 = ch_2[tr \mapsto \text{rep } (n-m) (st \oplus \Delta st) st''] \end{array}}{\text{rep } e s, \Delta st, ch_0 @ tr \Rightarrow_{\Delta} \Delta st', ch_3} \quad (\Delta\text{REP-SUB}) \\
\\
\frac{\begin{array}{l} ch_0 tr = \text{rep } n \text{ st } st' \\ e, \Delta st, ch_0 @ tr \cdot \text{rep}_{\text{count}} \Rightarrow_{\Delta} +m, ch_1 \\ \text{rep}^n s, \Delta st, ch_1 @ tr \cdot \text{rep}_{\text{aux}} \Rightarrow_{\Delta} \Delta st', ch_2 \\ \text{rep}^m s, (st' \oplus \Delta st'), ch_2 @ tr \cdot \text{rep}_{\text{aux}} \cdot \text{rep}_{\text{iter}}^n \Rightarrow_{\Delta} st'', ch_3 \\ \Delta st' = st'' \ominus st' \\ ch_4 = ch_3[tr \mapsto \text{rep } (n+m) (st \oplus \Delta st) st''] \end{array}}{\text{rep } e s, \Delta st, ch_0 @ tr \Rightarrow_{\Delta} \Delta st', ch_4} \quad (\Delta\text{REP-ADD})
\end{array}$$

Fig. 6. For loops, the differential semantics inspects how the iteration count changes.

whereas the differential semantics of $(\text{rep}^n s)$ does not interact with the cache. Indeed, as shown in Figure 6, rules $(\Delta\text{REP-0})$ and $(\Delta\text{REP+1})$ bear no surprises.

When the iteration count does not change $(\Delta\text{REP-NOC})$, the control flow of the program remains the same. Therefore, it is sufficient to differentially execute $(\text{rep}^n s)$ to propagate the changes. Note that the number of iterations n done by the original loop is retrieved from the cache. When instead the number of iterations is reduced $(\Delta\text{REP-SUB})$, we execute the loop from scratch. This is the only option, because we cannot execute loop bodies backward and chose to not cache intermediate results. We obtain the resulting store change by diffing the original (and cached) result st' with the new result st'' .

Finally, when the iteration count grows $(\Delta\text{REP-ADD})$, we compute the output change in two steps. First, we propagate the input change through the original n iterations incrementally. Second, we execute the m new iterations from scratch. Since $(\text{rep}^m s)$ is a continuation of $(\text{rep}^n s)$, we have to produce a trace that simulates this behavior by n -fold repetition of rep_{iter} . Only then will the two subcomputations have isolated cache entries, which is required for correctness. When $m \ll n$, $(\Delta\text{REP-ADD})$ yields significant speedups compared to a full recomputation [14].

For example, consider the following program that computes the factorial of n :

<pre># free variable n i := 2; r := 1 repeat n - 1: r := r * i i := i + 1</pre>	<p>$n = 3$</p> <p>$\Delta n = +2$</p>	$ \begin{aligned} ch &= \{ \text{seq}_R \cdot \text{seq}_R \cdot \text{rep}_{\text{aux}} \cdot \text{rep}_{\text{body}} \mapsto \text{mul } 1 \ 2, \\ &\quad \text{seq}_R \cdot \text{seq}_R \cdot \text{rep}_{\text{aux}} \cdot \text{rep}_{\text{iter}}^1 \cdot \text{rep}_{\text{body}} \mapsto \text{mul } 2 \ 3, \\ &\quad \text{seq}_R \cdot \text{seq}_R \mapsto \text{rep } 2 \{ n \mapsto 3, i \mapsto 2, r \mapsto 1 \} \\ &\quad \quad \quad \{ n \mapsto 3, i \mapsto 4, r \mapsto 6 \} \} \\ ch' &= \{ \text{seq}_R \cdot \text{seq}_R \cdot \text{rep}_{\text{aux}} \cdot \text{rep}_{\text{body}} \mapsto \text{mul } 1 \ 2, \\ &\quad \text{seq}_R \cdot \text{seq}_R \cdot \text{rep}_{\text{aux}} \cdot \text{rep}_{\text{iter}}^1 \cdot \text{rep}_{\text{body}} \mapsto \text{mul } 2 \ 3, \\ &\quad \text{seq}_R \cdot \text{seq}_R \cdot \text{rep}_{\text{aux}} \cdot \text{rep}_{\text{iter}}^2 \cdot \text{rep}_{\text{body}} \mapsto \text{mul } 6 \ 4, \\ &\quad \text{seq}_R \cdot \text{seq}_R \cdot \text{rep}_{\text{aux}} \cdot \text{rep}_{\text{iter}}^3 \cdot \text{rep}_{\text{body}} \mapsto \text{mul } 24 \ 5, \\ &\quad \text{seq}_R \cdot \text{seq}_R \mapsto \text{rep } 4 \{ n \mapsto 5, i \mapsto 2, r \mapsto 1 \} \\ &\quad \quad \quad \{ n \mapsto 5, i \mapsto 6, r \mapsto 120 \} \} \end{aligned} $
---	---	--

We start with $n = 3$, which yields $r = 6$ after two iterations. After the initial run, the cache contains three entries: a distinct entry for each iteration's multiplication, and a single entry to cache the result of the loop. Now, consider we make a change $\Delta st = \{\Delta n \mapsto +2\}$. Since the loop's counter expression $n - 1$ evaluates to $+2$ under Δst , we trigger (Δ REP-ADD). This rule first propagates Δst through the original 2 iterations, which yields $\Delta st' = \Delta st$, that is, there is no further change. Indeed, it is safe to skip the original iterations because the changed Δn does not occur in the body of the loop (Kumar et al. [14] call this the short-circuiting optimization). Next, we run the 2 new iterations in $st' \oplus \Delta st' = \{n \mapsto 5, i \mapsto 4, r \mapsto 6\}$, which yields $st'' = \{n \mapsto 5, i \mapsto 6, r \mapsto 120\}$ and inserts two additional cache entries for multiplications. Lastly, we update the cache entry for the loop and derive the output change $st'' \ominus st' = \{i \mapsto +2, r \mapsto +114\}$.

This concludes our differential caching semantics. We have explained the key ideas of lexical tracing and incremental caching with a focus on illustration and intuition. The next section studies the theory of our semantics and proves it correct.

4 Cache Equivalence, Cache Stability, and Differential Correctness

We want to prove in Rocq that the differential semantics yields valid and correct changes. This ensures we can safely use the differential semantics in place of a recomputation. Given the original execution $s, st, ch_0 \Rightarrow st', ch_1$ and a subsequent differential execution $s, \Delta st, ch_1 \Rightarrow \Delta st', ch_2$, we must prove:

Validity The output change must be valid for the original output $st' \Vdash \Delta st'$, which ensures patching $st' \oplus \Delta st'$ succeeds. For example, $\{l \mapsto [2, 4, 6]\} \Vdash \{l \mapsto \text{drop } 1\}$ holds, whereas $\{l \mapsto [2, 4, 6]\} \Vdash \{l \mapsto \text{drop } 5\}$ does not hold, because the list only contains 3 elements.

Correctness The output change st' must reflect the difference between the original computation and a from-scratch recomputation. That is, we require $s, st \oplus \Delta st, ch_0 \Rightarrow_{\Delta} st' \oplus \Delta st', ch_2$.

Completeness Lastly, we show that the differential semantics can react to any valid input change and never gets stuck. That is, $s, \Delta st, ch_1 \Rightarrow_{\Delta} \Delta st', ch_2$ exists whenever $s, st, ch_0 \Rightarrow st', ch_1$. In particular, the cache is always well-shaped to satisfy the cache requirements of differential execution.

Unfortunately, proving these properties is not simple at all. A key challenge is that we need to reason about the original and the differential execution together, and the original output cache becomes the differential input cache. Consider the case for sequences $s_1 ; s_2$ and its rules:

$$\frac{s_1, st, ch_0 @ tr \cdot \text{seq}_L \Rightarrow st_a, ch_a \quad s_2, st_a, ch_a @ tr \cdot \text{seq}_R \Rightarrow st', ch_1}{s_1 ; s_2, st, ch_0 @ tr \Rightarrow st', ch_1} \text{ (SEQ)} \quad \frac{s_1, \Delta st, ch_1 @ tr \cdot \text{seq}_L \Rightarrow_{\Delta} \Delta st_b, ch_b \quad s_2, \Delta st_b, ch_b @ tr \cdot \text{seq}_R \Rightarrow_{\Delta} \Delta st', ch_2}{s_1 ; s_2, \Delta st, ch_1 @ tr \Rightarrow_{\Delta} \Delta st', ch_2} \text{ (\Delta SEQ)}$$

The problem is that we cannot use the induction hypothesis for s_1 or s_2 . To use an induction hypothesis, the output cache of the initial computation $s_1, st, ch_0 @ tr \cdot \text{seq}_L \Rightarrow st_a, ch_a$ needs to be the same as the input cache of the differential computation $s_1, \Delta st, ch_1 @ tr \cdot \text{seq}_L \Rightarrow_{\Delta} \Delta st_b, ch_b$. However, $ch_a = ch_1$ does not hold in general.

For example, consider $x := 3 ; y := 2 * x$. The original execution of $x := 3$ does not cache anything, hence $ch_a = \{\}$. However, the subsequent execution of $y := 2 * x$ caches the multiplication, hence $ch_1 = \{\text{seq}_R \mapsto \text{mul } 23\}$. Since the differential computation starts with ch_1 , we have $ch_a \neq ch_1$, which shows why we cannot use the induction hypothesis for s_1 . But, there is a way forward: When executing s_1 , we only access the cache for traces starting with seq_L , whereas when executing s_2 , we only access the cache for traces starting with seq_R . Therefore, ch_1 is equal to ch_a for all traces starting with seq_L . Indeed, since $tr \cdot \text{seq}_L \neq tr \cdot \text{seq}_R$, the executions of s_1 and s_2 operate on independent parts of the cache. We need to incorporate this observation into our theory.

$$\begin{array}{c}
\frac{}{s \succ \varepsilon} \quad \frac{e_1 \succ tr}{e_1 \text{ op } e_2 \succ \text{op}_L \cdot tr} \quad \frac{e_2 \succ tr}{e_1 \text{ op } e_2 \succ \text{op}_R \cdot tr} \\
\\
\frac{}{s \succ \varepsilon} \quad \frac{e \succ tr}{x := e \succ \text{assign}_{\text{rhs}} \cdot tr} \quad \frac{s_1 \succ tr}{s_1 ; s_2 \succ \text{seq}_L \cdot tr} \quad \frac{s_2 \succ tr}{s_1 ; s_2 \succ \text{seq}_R \cdot tr} \\
\\
\frac{e \succ tr}{\text{if } e \ s_1 \ s_2 \succ \text{if}_{\text{cond}} \cdot tr} \quad \frac{s_1 \succ tr}{\text{if } e \ s_1 \ s_2 \succ \text{if}_{\text{then}} \cdot tr} \quad \frac{s_2 \succ tr}{\text{if } e \ s_1 \ s_2 \succ \text{if}_{\text{else}} \cdot tr} \\
\\
\frac{e \succ tr}{\text{rep } e \ s \succ \text{rep}_{\text{count}} \cdot tr} \quad \frac{\text{rep}^n s \succ tr}{\text{rep } e \ s \succ \text{rep}_{\text{aux}} \cdot tr} \quad \frac{s \succ tr}{\text{rep}^n s \succ \text{rep}_{\text{body}} \cdot tr} \quad \frac{\text{rep}^n s \succ tr}{\text{rep}^n s \succ \text{rep}_{\text{iter}} \cdot tr}
\end{array}$$

Fig. 7. Trace reachability $s \succ tr$ enumerates all possible trace suffixes that can occur during the execution of s .

Note. Except for trace reachability, we only show the metatheory for IMP statements and their execution. All theorems we show have an analogous counterpart for the expressions of IMP. Of course, our Rocq formalization contains proofs for all of the metatheory: for IMP expressions and statements. However, since the metatheory for expressions does not reveal any additional insight, we decided to omit it from the presentation. We also generally assume all programs and stores are well-typed, that is, we only consider executions where $\Gamma \vdash st$ and $\Gamma \vdash s \text{ ok}$ for some context Γ .

4.1 Cache Equivalence and Trace Reachability

To strengthen the induction hypothesis, we need to generalize the properties we are proving. Specifically, we may not request that the cache produced by the initializing execution is used as input for the differential execution. Rather, we only require these caches to be *observably equivalent*. Two caches are observably equivalent if they provide equal cache entries for all *reachable traces*.

Definition 4.1 (Trace reachability). Given a statement s , only reachable traces $s \succ tr$ can occur during the execution of s . We define the set of reachable traces inductively in Figure 7.

To define cache equivalence, we additionally consider the trace tr_0 that records the execution from the program's main entry point to the current statement s . Two caches are equivalent for executing s if they have identical entries for all traces $(tr_0 \cdot tr)$, where tr is reachable from s :

Definition 4.2 (Cache equivalence). Given a statement s and a trace tr_0 , caches ch_1 and ch_2 are observably equivalent $s @ tr_0 \vdash ch_1 \sim ch_2$ if for all $s \succ tr$, we have $ch_1(tr_0 \cdot tr) = ch_2(tr_0 \cdot tr)$.

We formalized trace reachability and cache equivalence in Rocq and proved a few interesting properties. First, we show that we indeed defined an equivalence relation, which is trivial to prove.

Lemma 4.3. For any s and tr , cache equivalence $s @ tr \vdash ch_1 \sim ch_2$ is an equivalence relation.

A more interesting property is that cache updates are sometimes irrelevant. Specifically, a cache update at tr' is irrelevant for any cache equivalence rooted at tr if tr is not a prefix of tr' . This is because cache equivalence at tr only concerns cache entries for reachable continuations of tr , which does not include tr' .

Lemma 4.4. For any tr not prefix of tr' , $s @ tr \vdash ch_1 \sim ch_2$ if and only if $s @ tr \vdash ch_1 \sim ch_2[tr' \mapsto e]$.

PROOF. Since $tr \cdot tr_0 \neq tr'$ for any tr_0 , the entry for tr' is irrelevant for cache equivalence. \square

Lastly, it is interesting to note that trace reachability over-approximates the actual traces inhabited by the execution. This is observable for conditionals, where both branches are reachable independent

of the condition. For loops, this is even more remarkable, because trace reachability and cache equivalence make no prediction about the number of iterations a loop may take.

Lemma 4.5. For any m and n , if $\text{rep}^m s @ tr \vdash ch_1 \sim ch_2$, then $\text{rep}^n s @ tr \vdash ch_1 \sim ch_2$.

PROOF. By induction on tr , $\text{rep}^m s \succ tr$ implies $\text{rep}^n s \succ tr$. It follows that the iteration count is irrelevant for cache equivalence, because the set of relevant traces is independent of the count. \square

Cache equivalence forms the basis for reasoning about the differential execution inductively. But before we can prove validity, correctness, and completeness of differential execution, we first must observe and verify a number of key properties that connect execution and cache equivalence.

4.2 Cache Stability and Cache Invariance

We introduced lexical traces for two reasons: First, to ensure we can uniquely identify language constructs that need caching during the execution. For IMP, this was multiplication, comparisons, conditionals, and loops. The theorems in Section 4.3 ensure we succeeded.

The second reason we introduced lexical traces was to make cache entries stable during differential execution. For example, in $s_1 ; s_2$, changes to the control flow in s_1 do not proliferate the lexical traces that occur in s_2 . Specifically, we argued that independent subcomputations do not affect each others' cache entries. Recall from Definition 2.1 that two traces are independent $tr \# tr'$ if neither is a prefix of the other, which means tr and tr' root independent sub-computations. We thus formulate cache stability for the initializing execution.

Theorem 4.6 (Initial cache stability). Given independent traces $tr \# tr'$, an execution rooted at tr does not affect cache entries rooted at tr' : If $s, st, ch @ tr \Rightarrow st', ch'$, then $s' @ tr' \vdash ch \sim ch'$.

PROOF. By induction over the execution relation, using Lemmas 4.3 and 4.4. Note that s' is indeed independent of s , because the root traces are sufficient to distinguish the cache accesses. \square

The proof of this theorem is quite straightforward actually, although we programmed a number of tactics in Rocq to improve automation. Unfortunately, cache stability leaves one special case: When tr is not a prefix of tr' , but tr' may be a prefix of tr . In this case, we cannot rely on the reachable traces used in cache equivalence, because that could include trace continuations shared by tr and tr' . But, we can guarantee cache stability for the specific trace tr' , which we need later.

Lemma 4.7. For any tr not prefix of tr' , when $s, st, ch @ tr \Rightarrow st', ch'$, then $ch tr' = ch' tr'$.

PROOF. By induction over the execution relation, using Lemma 4.3. \square

With this, we can prove an important property about cache equivalence. Namely, equivalent caches lead to equivalent initializing executions, where the resulting stores are identical and the resulting caches are again equivalent. In other words, the initial execution is invariant to equivalent caches.

Theorem 4.8 (Cache invariance). Given caches such that $s @ tr_0 \vdash ch_1 \sim ch_2$. For any st , if $s, st, ch_1 @ tr_0 \Rightarrow st_1, ch'_1$ and $s, st, ch_2 @ tr_0 \Rightarrow st_2, ch'_2$, then $st_1 = st_2$ and $s @ tr_0 \vdash ch'_1 \sim ch'_2$.

PROOF. By induction over one of the execution relations, using Theorem 4.6 and all of the lemmas from above. \square

The proof of cache invariance is rather involved, even though we wrote tactics that generalize repeating proof patterns. But this property is elemental to proving differential correctness when branch switching occurs, where it ensures that the from-scratch execution and the branch-switch execution yield the same stores, even though their caches differ.

4.3 Differential Validity, Correctness, and Completeness

So far, we have reasoned about cache equivalence and have shown the initializing execution is well-behaved with respect to it. It is time to validate that the differential execution is also well-behaved. We start by proving that differential execution adheres to cache stability.

Theorem 4.9 (Differential cache stability). Given independent traces $tr \# tr'$, differential execution rooted at tr does not affect cache entries rooted at tr' . That is, if $s, \Delta st, ch @ tr \Rightarrow_{\Delta} \Delta st', ch'$, then $s' @ tr' \vdash ch \sim ch'$.

PROOF. By induction over the differential execution relation, using Theorem 4.6 and Lemmas 4.3 and 4.4. Again, s' is independent of s , because the root traces are sufficient to distinguish the cache accesses. \square

We now have the necessary building blocks to tackle our main theorems. Differential validity ensures that the changes produced by differential execution can be patched. Indeed, patching $st' \oplus \Delta st'$ is only defined when $st' \Vdash \Delta st'$.

Theorem 4.10 (Differential validity). Given a valid input change $st \Vdash \Delta st$ and equivalent caches $s @ tr \vdash ch_1 \sim ch_2$. If the initializing execution $s, st, ch_0 @ tr \Rightarrow st', ch_1$ is followed by a differential execution $s, \Delta st, ch_2 @ tr \Rightarrow_{\Delta} \Delta st', ch_3$, then $st' \Vdash \Delta st'$.

PROOF. By induction over the differential execution relation. Note that $st' \Vdash \Delta st'$ is defined pointwise and holds if for all x , $st' x \Vdash \Delta st' x$. For assignments, we hence rely on the validity of differential evaluation of the assignment's right-hand side. For branch switching, we exploit that diffing produces valid changes: $st_2 \Vdash st_1 \ominus st_2$. \square

Now that we know the differential result is valid, we can investigate the correctness of the resulting change store $\Delta st'$. Specifically, we compare the patched result $st' \oplus \Delta st'$ to a store st'' , which has been recomputed from scratch for the patched input.

Theorem 4.11 (Differential correctness). Given a valid input change $st \Vdash \Delta st$ and equivalent caches $s @ tr \vdash ch_1 \sim ch_2$. If the initializing execution $s, st, ch_0 @ tr \Rightarrow st', ch_1$ is followed by a differential execution $s, \Delta st, ch_2 @ tr \Rightarrow_{\Delta} \Delta st', ch_3$, then $st' \oplus \Delta st'$ yields the same result as a from-scratch execution. That is, for ch_4 equivalent to the initial cache $s @ tr \vdash ch_0 \sim ch_4$, then any from-scratch execution on the changed input $s, st \oplus \Delta st, ch_4 @ tr \Rightarrow st'', ch_5$ yields $st'' = st' \oplus \Delta st'$.

PROOF. By induction over the differential execution relation. The most complex case is (Δ REP-ADD), where we need to reason about rep^{n+m} s and show that we can split this into rep^n s followed by rep^m s . Moreover, we must prove that the starting trace $tr \cdot \text{rep}_{\text{aux}} \cdot \text{rep}_{\text{iter}}^n$ of rep^m s indeed corresponds to the trace used by the from-scratch execution, which allows us to derive the necessary cache equivalences. \square

Differential correctness guarantees that differential execution yields the same result as a from-scratch execution. However, correctness does not ensure that differential execution actually finds an output change. For example, a differential execution relation that is always stuck would satisfy the above correctness criterion. We need a completeness theorem to ensure that differential execution must deliver an output change.

Theorem 4.12 (Differential completeness). Given a valid input change $st \Vdash \Delta st$ and equivalent caches $s @ tr \vdash ch_1 \sim ch_2$ as well as $s @ tr \vdash ch_0 \sim ch_4$. If the initializing execution $s, st, ch_0 @ tr \Rightarrow st', ch_1$ and the from-scratch execution $s, st \oplus \Delta st, ch_4 @ tr \Rightarrow st'', ch_5$ succeed, then the differential execution succeeds $s, \Delta st, ch_2 @ tr \Rightarrow_{\Delta} \Delta st', ch_3$ for some $\Delta st'$ and ch_3 .

PROOF. By induction on the initializing execution relation. For each initializing execution rule, we need to find differential execution rules that can fire, which ensures the cache requirements of the differential semantics are always satisfied. Note that the premise about from-scratch execution is actually only needed when executions can diverge, which is not the case for IMP. \square

With respect to termination, our completeness theorem states that the differential execution must terminate whenever both the original and the patched execution terminate. This is a very strong property. In case either original or patched execution diverge, it is likely that the differential execution will do the same, but this is not covered by our theory. Together, correctness and completeness ensure we can use differential execution as a drop-in replacement for from-scratch executions.

4.4 Summary

We have stated and verified key properties about the metatheory of differential execution. The most important results are differential validity, correctness, and completeness, which together ensure we can use differential execution as a drop-in replacement for a from-scratch execution. To verify these properties, we had to develop a theory of local cache equivalence and prove that execution only affects cache entries local to the executed sub-program. Indeed, execution is invariant under equivalent caches, which is crucial for reasoning about branch switching within differential execution.

We have formalized the complete metatheory in the Rocq proof assistant. The mechanized metatheory comprises about 3600 SLoC and includes 146 lemmas and 26 tactics (small and large). We will submit the Rocq code for artifact evaluation and make it open source.

5 Extending the Expressivity of IMP

Our imperative language IMP, in its current form, cannot express general-purpose algorithms because it lacks syntax constructs for recursion and while loops. We now extend the language's syntax and semantics to incorporate these essential features.

5.1 Differential Execution of Recursive Functions

To support recursive functions, we extend the IMP syntax with constructs for function definitions and calls, as detailed in [Figure 8](#). A function definition takes the form $(\text{def } f(x) = s; e)$, which specifies the function's formal parameter x , its body (a statement s), and its return expression e . A function call is a statement of the form $(x := \text{call } f \ e)$, which invokes function f with the argument expression e and assigns the result to the variable x . To trace the a function call, we also introduce three corresponding trace links: call_{arg} for argument evaluation, $\text{call}_{\text{body}}$ for the execution of the function body, and $\text{call}_{\text{exit}}$ for the evaluation of the return expression.

With the additions of functions, the structure of the big-step and differential semantics is extended. We introduce a function environment env to map function names to definitions, as [highlighted in Figure 8](#). While all semantics rules in [Figure 3](#) must be updated to carry this environment, we omit them for brevity, with the understanding that env is simply propagated where it is not actively used. The (CALL) rule implements a call-by-value semantics where functions operate in a local scope. It evaluates the argument e , then executes the function's body s in a fresh store st_{empty} containing only the parameter binding. The function's final return value is then assigned to x in the original store st , ensuring the call has no other side effects on the caller's variables. Rule (ΔCALL) is structurally analogous to (CALL), but it threads a differential store and a cache through the execution. This allows previously computed results stored in the cache to be reused in subsequent computations.

$$\begin{array}{ll}
\text{(function definitions)} & d ::= \text{def } f(x) = s; e \\
\text{(statements)} & s ::= \dots \mid x := \text{call } f \ e \\
\text{(environments)} & \text{env} \in f \rightarrow d \\
\text{(trace links)} & l ::= \dots \mid \text{call}_{\text{arg}} \mid \text{call}_{\text{body}} \mid \text{call}_{\text{exit}}
\end{array}$$

$$\frac{
\begin{array}{l}
\text{env } f = \text{def } f(y) = s; r \\
s, \text{st}_{\text{empty}}[y \mapsto v_e], \text{ch}_1, \text{env} @ \text{tr} \cdot \text{call}_{\text{body}} \Rightarrow \text{st}_1, \text{ch}_2 \\
e, \text{st}, \text{ch}_0 @ \text{tr} \cdot \text{call}_{\text{arg}} \Rightarrow v_e, \text{ch}_1 \\
r, \text{st}_1, \text{ch}_2 @ \text{tr} \cdot \text{call}_{\text{exit}} \Rightarrow v_r, \text{ch}_3
\end{array}
}{
x := \text{call } f \ e, \text{st}, \text{ch}_0, \text{env} @ \text{tr} \Rightarrow \text{st}[x \mapsto v_r], \text{ch}_3 \quad (\text{CALL})
}$$

$$\frac{
\begin{array}{l}
\text{env } f = \text{def } f(y) = s; r \quad e, \Delta \text{st}, \text{ch}_0 @ \text{tr} \cdot \text{call}_{\text{arg}} \Rightarrow_{\Delta} \Delta v_e, \text{ch}_1 \\
s, \Delta \text{st}_{\text{empty}}[y \mapsto \Delta v_e], \text{ch}_1, \text{env} @ \text{tr} \cdot \text{call}_{\text{body}} \Rightarrow_{\Delta} \Delta \text{st}_1, \text{ch}_2 \\
r, \Delta \text{st}_1, \text{ch}_2 @ \text{tr} \cdot \text{call}_{\text{exit}} \Rightarrow_{\Delta} \Delta v_r, \text{ch}_3
\end{array}
}{
x := \text{call } f \ e, \Delta \text{st}, \text{ch}_0, \text{env} @ \text{tr} \Rightarrow_{\Delta} \Delta \text{st}[x \mapsto \Delta v_r], \text{ch}_3 \quad (\Delta \text{CALL})
}$$

Fig. 8. Big-step semantics and differential semantics for function calls and the **new function environment**.

Since function bodies can contain arbitrary statements, this extension also makes IMP Turing-complete. We have formally proven that differential validity ([Theorem 4.10](#)), correctness ([Theorem 4.11](#)), and completeness ([Theorem 4.12](#)) hold for the recursive semantics in Rocq.

5.2 Differential Execution of While Loops

IMP supports repeat loops, where the number of iterations is known before the loop starts. This was useful for the differential semantics, because it allowed us to predict if we need the same, more, or less iterations after an input change. In particular, our differential semantics for repeat loops bailed out when the number of iterations had shrunk ($\Delta_{\text{REP-SUB}}$), which required knowledge of the new number of iterations ahead of time. However, it is also possible to support generic while loops, as we will show here. Note that while loops are *not* part of our mechanized metatheory.

While loops continue to iterate until the while condition evaluates to `false`. Hence, we cannot predict the number of iterations a priori. The best we can do is to start iterating the loop to propagate the changes using differential execution until either the original or the new iterations stop, meaning their control flow diverges. Now, we have three scenarios: The differential execution of the loop may terminate earlier, later, or after the same number of iterations as the original execution. However, since we already committed to a differential execution of the loop, it would be wasteful to bail now and to start a from-scratch execution from the beginning. Instead, we choose to cache the current store at each loop iteration.

We introduce the following syntax to describe the semantics of while loops and their lexical traces.

$$\begin{array}{ll}
\text{(statements)} & s ::= \dots \mid \text{while } e \ s \mid \text{while}^n \ e \ s \\
\text{(trace links)} & l ::= \dots \mid \text{while}_{\text{aux}} \mid \text{while}_{\text{cond}} \mid \text{while}_{\text{body}} \mid \text{while}_{\text{iter}}
\end{array}$$

Besides the standard while loop `while e s`, we also add an auxiliary construct `whilen e s`. In the non-differential semantics, `whilen e s` counts upward how many iterations n of the loop have *already happened*. But in the differential semantics, `whilen e s` counts downward how many iterations n are left before the original execution stopped. This allows us to detect which execution stops a loop earlier.

We show the big-step semantics of while loops in [Figure 9](#). Rule (`WHILE`) initializes the iteration at `while0 e s` and retrieves the final iteration count i from the sub-computation. Indeed, as **highlighted**,

$$\begin{array}{c}
\frac{\text{while}^0 e s, st, ch_0 @ tr \cdot \text{while}_{\text{aux}} \Rightarrow st', ch_1, i}{\text{while } e s, st, ch_0 @ tr \Rightarrow st', ch_1 [tr \mapsto \text{count } i], 0} \quad (\text{WHILE}) \\
\\
\frac{e, st, ch_0 @ tr \cdot \text{while}_{\text{cond}} \Rightarrow \text{false}, ch_1}{\text{while}^n e s, st, ch_0 @ tr \Rightarrow st, ch_1 [tr \mapsto \text{step } st], n} \quad (\text{WHILE-FALSE}) \\
\\
\frac{\begin{array}{l} e, st, ch_0 @ tr \cdot \text{while}_{\text{cond}} \Rightarrow \text{true}, ch_1 \\ s, st, ch_1 [tr \mapsto \text{step } st] @ tr \cdot \text{while}_{\text{body}} \Rightarrow st', ch_2, i_0 \\ \text{while}^{n+1} e s, st', ch_2 @ tr \cdot \text{while}_{\text{iter}} \Rightarrow st'', ch_3, i \end{array}}{\text{while}^n e s, st, ch_0 @ tr \Rightarrow st'', ch_3, i} \quad (\text{WHILE-TRUE}) \\
\\
\frac{\begin{array}{l} ch_0 tr = \text{count } i \quad \text{while}^i e s, \Delta st, ch_0 @ tr \cdot \text{while}_{\text{aux}} \Rightarrow_{\Delta} \Delta st', ch_1, \Delta i \\ \text{while } e s, \Delta st, ch_0 @ tr \Rightarrow_{\Delta} \Delta st', ch_1 [tr \mapsto \text{count } i \oplus \Delta i], +0 \end{array}}{\text{while}^0 e s, \Delta st, ch_0 @ tr \Rightarrow_{\Delta} \Delta st, ch_1, +0} \quad (\Delta \text{WHILE}) \\
\\
\frac{e, \Delta st, ch_0 @ tr \cdot \text{while}_{\text{cond}} \Rightarrow_{\Delta} \text{noc}, ch_1}{\text{while}^0 e s, \Delta st, ch_0 @ tr \Rightarrow_{\Delta} \Delta st, ch_1, +0} \quad (\Delta \text{WHILE-BOTH-STOP}) \\
\\
\frac{\begin{array}{l} e, \Delta st, ch_0 @ tr \cdot \text{while}_{\text{cond}} \Rightarrow_{\Delta} \text{noc}, ch_1 \\ s, \Delta st, ch_1 @ tr \cdot \text{while}_{\text{body}} \Rightarrow_{\Delta} \Delta st', ch_2, \Delta i_0 \\ \text{while}^n e s, \Delta st', ch_2 @ tr \cdot \text{while}_{\text{iter}} \Rightarrow_{\Delta} \Delta st'', ch_3, \Delta i \end{array}}{\text{while}^{n+1} e s, \Delta st, ch_0 @ tr \Rightarrow_{\Delta} \Delta st'', ch_3, \Delta i} \quad (\Delta \text{WHILE-BOTH-STEP}) \\
\\
\frac{\begin{array}{l} e, \Delta st, ch_0 @ tr \cdot \text{while}_{\text{cond}} \Rightarrow_{\Delta} \text{neg}, ch_1 \\ ch_0 tr = \text{step } st \quad \text{while}^0 e s, st \oplus \Delta st, ch_1 @ tr \Rightarrow_{\Delta} st', ch_2, i \end{array}}{\text{while}^0 e s, \Delta st, ch_0 @ tr \Rightarrow_{\Delta} st' \ominus st, ch_2, +i} \quad (\Delta \text{WHILE-RUN-LONGER}) \\
\\
\frac{\begin{array}{l} e, \Delta st, ch_0 @ tr \cdot \text{while}_{\text{cond}} \Rightarrow_{\Delta} \text{neg}, ch_1 \\ ch_0 tr = \text{step } st \quad ch_0 (tr ++ \text{while}_{\text{iter}}^{n+1}) = \text{step } st' \quad ch_2 = ch_1 [tr \mapsto st \oplus \Delta st] \end{array}}{\text{while}^{n+1} e s, \Delta st, ch_0 @ tr \Rightarrow_{\Delta} (st \oplus \Delta st) \ominus st', ch_2, -(n+1)} \quad (\Delta \text{WHILE-STOP-EARLY})
\end{array}$$

Fig. 9. Big-step and differential semantics for while loops and the new iteration count.

we had to extend the big-step relation to produce the iteration count of $\text{while}^n e s$ loops, and the differential relation to compute how the iteration count changes.¹ The sole purpose of these counts is to allow caching at the end of the loop in (WHILE). Rule (WHILE-FALSE) stops the loop when the condition evaluates to false, and caches the current store. Rule (WHILE-TRUE) caches the store, runs the body, and continues the loop with an incremented iteration count.

Figure 9 also shows the differential execution rules for while loops. Rule (Δ WHILE) retrieves the original iteration count i from the cache and initializes the differential iteration at $\text{while}^i e s$. Once the iteration is finished, we update the cached iteration count using the iteration change

¹We could have introduced an auxiliary reduction relation, but opted for this simpler version encoding where all other constructs yield a dummy iteration count 0.

Δi . The next two rules ($\Delta_{\text{WHILE-BOTH-STOP}}$) and ($\Delta_{\text{WHILE-BOTH-STEP}}$) handle the cases, where original and differential execution agree on the iteration length: The evaluation of the condition e is \cdot . When we finished all original iterations and reach $\text{while}^0 e s$ in ($\Delta_{\text{WHILE-BOTH-STOP}}$), the iteration count remains the same. When instead there are original iteration left $\text{while}^{n+1} e s$ and the condition is unchanged (i.e., it is still true), we differentially run the body s once and continue with one less original iterations $\text{while}^n e s$.

But what if the iteration counts of the original and the differential execution disagree. Rule ($\Delta_{\text{WHILE-RUN-LONGER}}$) fires when we reach $\text{while}^0 e s$ (all original iterations are finished) but the condition's values is negated. This means the original condition was false, but now it is true. Hence, we must continue running the loop longer than before. To this end, we start a from-scratch execution after the original iterations, and mark the extra iterations i as a change to the iteration count. Rule ($\Delta_{\text{WHILE-STOP-EARLY}}$) handles the opposite case: We reach $\text{while}^{n+1} e s$ (there are more original iterations) but again the condition's values is negated. This means the original condition was true, but now it is false, and we must stop the loop earlier than before. Here we can use a small trick: We retrieve the store at the end of the original loop by calculating the associated lexical trace $tr \uparrow\uparrow \text{while}_{\text{iter}}^{n+1}$. This way, we can compute how the current up-to-date store $st \oplus \Delta st$ differs from the original store at the end of the loop st' . Lastly, we mark the surplus iterations $n + 1$ as a negative iteration count change.

In summary, the differential semantics of while loops uses a clever trick: In the initial semantics, $\text{while}^i e s$ counts how many iterations have already been done, that is, i is growing. In contrast, in the differential semantics, $\text{while}^i e s$ counts how many original iterations we can still match with the differential run, that is, i is shrinking. This allows for an efficient handling of while loops during differential execution. We have formalized the semantics of while loops in Rocq and have implemented and tested it in Scala. However, we were not able to prove the metatheoretical properties of differential execution for while loops. The problem is that we need to relate the differential run to the initial run in the validity and correctness theorems, but the iteration schemes misalign: The initial semantics moves from $\text{while}^i e s$ to $\text{while}^{i+1} e s$, whereas the differential semantics moves from $\text{while}^{i+1} e s$ to $\text{while}^i e s$. This makes it difficult to find a sufficiently strong induction hypothesis. It may be possible to design a less optimized semantics for while loops that avoids this problem, but we leave this for future work.

6 Implementation and Evaluation

We implemented the initializing and differential semantics of IMP as interpreters in Scala. The implementation follows the semantics very closely except for a few optimizations.

6.1 Optimizations in Implementing Differential Execution

The cache is passed around throughout the semantics, but we actually only ever access the latest version of the cache. Therefore, we can use a single shared cache with destructive updates instead. The implementation of the interpreters then follows a simple architecture:

```
class IncrementalInterpreter(prog: Program):
  val cache: Cache = new Cache
  val initialInterpreter = new InitialInterpreter(prog, cache)
  val differentialInterpreter = new DifferentialInterpreter(prog, cache, initialInterpreter)

  def runInitial(args: Seq[Value]): Value = initialInterpreter.run(args)
  def runDifferential(args: Seq[Value]): Value = differentialInterpreter.run(args)
```

```

enum Step:
  case Return; case SeqL; case SeqR(n: Int)
  case RepCount; case RepAux; case RepBody; case Iter(n: Int)
  ...
enum Trace:
  case Root
  case TraceStep(step: Step, tr: Trace)
  override val hashCode: Int = MurmurHash3.productHash(this)
  def push(s: Step): Trace = TraceStep(s, this)
  def pushIter(): Trace = this match
    case TraceStep(Step.Iter(n), tail) => TraceStep(Step.Iter(n + 1), tail)
    case _ => TraceStep(Step.Iter(1), this)
  def pushSeqR(): Trace = this match
    case TraceStep(Step.SeqR(n), tail) => TraceStep(Step.SeqR(n + 1), tail)
    case _ => TraceStep(Step.SeqR(1), this)

```

Fig. 10. Lexical traces in Scala. We compress subsequent occurrences of `repiter` and `seqr`.

Note that both interpreters receive access to the mutable cache. Moreover, the differential interpreter obtains a reference to the initial interpreter, as required by the semantics.

Traces are used as keys in the cache. We show our Scala encoding of lexical traces in [Figure 10](#). As traces can grow long, the time required to compute hash values matters. We applied two simple optimizations to reduce this time. First, we compress traces: When multiple `SeqR` steps occur in sequence, we represent them as a single `SeqR(n)` in our implementation. We do the same for sequences of `RepIter`. Second, we cache the hash for each trace, such that the hash of `step · tr` can be computed in constant time independent of the size of `tr`. We believe that further trace compressions and more compact trace representations are possible, but have not investigated this yet.

6.2 Performance Evaluation

We proposed a principled approach to caching for differential execution: lexical traces that provide unique and stable access to the cache. Our approach improves over prior work on differential execution [14] in that the caching behavior is formally specified and part of the mechanized theory. In contrast, Kumar et al. [14] used extensive recomputation in their semantics and only considered caching in their implementation, which masks the effective caching semantics. But does our principled approach to caching have a negative impact on the initial or differential performance of the interpreter? To answer this question, we repeated the experiment of Kumar et al. [14] using their and our implementation of a differential interpreter for IMP.

We evaluate the initialization time and the incremental update time of the differential interpreters and compare them to a non-incremental baseline interpreter. We evaluate the interpreters on an IMP program that implements the Bellman-Ford algorithm: a single-source shortest path algorithm that supports negative edges and runs in cubic time. We took the code of the IMP program from Kumar et al. [14] and reproduce it in [Figure 11](#). The Bellman-Ford implementation uses a global mutable `table` to obtain the initial graph, where `table[u,v]` provides the edge weight between nodes `u` and `v`, both of which are identified by positive numbers. The algorithm then performs dynamic programming: At any point, `table[0,v]` contains the shortest path from the source node to `v` found so far.

We have extended our IMP interpreter to support the mutable `table` needed by Bellman-Ford. Semantically, `table` behaves like a store, but it uses index pairs (u, v) rather than variable names. In particular, the differential interpreter accepts `table` changes, which we use to encode changes to

```

def bellmanFord(numNodes: Int, src: Int): Unit {
  decl i: Int, u: Int, v: Int, w: Int
  i = 1
  repeat numNodes {
    table[0, i] = Int.Max
    i = i + 1
  }
  table[0, src] = 0
  repeat numNodes - 1 {
    u = 1
    repeat numNodes {
      v = 1
      repeat numNodes {
        w = table[u, v]
        if (w != 0 && table[0, v] > table[0, u] + w) {
          table[0, v] = table[0, u] + w
        }
        v = v + 1 }
      u = u + 1 }}}

```

Fig. 11. We evaluate the differential interpreters on the Bellman-Ford algorithm implemented in IMP.

the initial graph given to Bellman-Ford. We adopt the initial graph and the graph changes used by Kumar et al. [14]: The initial graph has $N \in \{10, 20, \dots, 150\}$ nodes that form a single cycle, which means each edge is important for the final result. We then introduce new edges with high weight that do not affect the shortest paths of the source node, so that these changes are amenable to efficient incrementalization.

We measure the wall-clock running time of the baseline interpreter, the initial and differential interpreter by Kumar et al. [14], and our new initial and differential interpreter using lexical tracing. We repeat each measurement 10 times and report the peak performance. We ran all measurements on an Apple M2 chip with 24GB memory running OSX 15.5. The interpreters are implemented in Scala 3, executed with OpenJDK 19.0.1 that uses up to 16 GB of memory for the heap. The benchmarking itself is done using the JMH framework.

We show the results in Figure 12. On the left-hand side, we see the initialization time in relation to the non-incremental baseline interpreter. We can see that differential interpretation has a considerable initialization overhead, mostly because it needs to populate the cache. However, our principled approach based on lexical traces only takes slightly longer compared to the prior work by Kumar et al. [14]. On the right-hand side, we see the incremental updating time of the differential interpreters in comparison to a from-scratch computation using the baseline interpreter. Here, we can see that (i) differential interpretation indeed provides asymptotic speedups and (ii) our principled approach does not introduce relevant performance penalties. We conclude that lexical tracing enables provably correct caching for differential execution while retaining the asymptotic performance advantages of differential execution.

6.3 Discussion: Where Do Incremental Performance Benefits Come From?

The main goal of incremental computing is to provide an asymptotic speedup over from-scratch recomputation. As the reader may have noticed, the operators used in our benchmark are simple arithmetic operators with a differential semantics that is not more efficient than its standard counterpart. Therefore, the performance benefits mainly stem from the incremental loop behavior

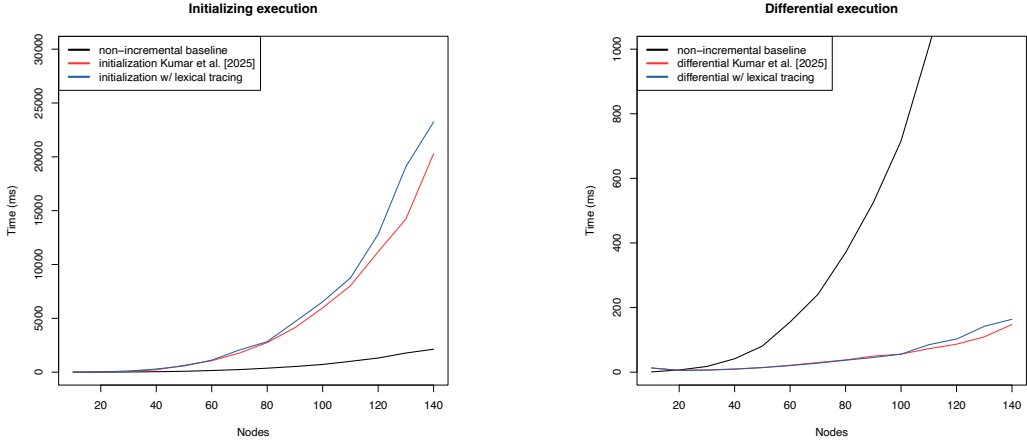


Fig. 12. Performance results. Left: Initialization time. Right: Differential updating time.

and the rules for short-circuiting introduced by Kumar et al. [14]. Short-circuiting allows us to skip the differential execution of statements that are not reading any changed variable.

For incremental loops, when the iteration count is unchanged, the loop only runs the differential semantics, yielding a speedup over the full recomputation. When the iteration count increases, only the additional iterations are executed with the standard semantics. On the other hand, when the iteration count decreases, we currently have to run all iterations with the standard semantics. However, this can be mitigated by caching the store after each iteration instead of only caching it at the end of the loop. This introduces a space-time tradeoff that we plan to investigate in our future work. Caching every n -th iteration as a checkpoint instead of caching every iteration is also a possible solution that requires working out efficient heuristics.

Although the differential version of the arithmetic operators is not more efficient than the standard one, falling back to their standard version during differential execution requires a transition to standard execution and then back to the differential one. Such a transition requires caching the operands and the result to be able to patch the operands with their delta values, execute the operator and diff the result afterwards to insert its delta value into the delta store. The overhead of such a transition is higher than the small overhead of the differential operator.

Nevertheless, more complex data structures like linked lists have differential operators with asymptotic speedups. As a simple example, consider the length operator on lists. As elements are inserted or removed from a list, a differential length operator can compute the length change in constant time $O(1)$, whereas a recomputation takes linear time $O(n)$. Since incorporating such operators would divert from our primary focus on caching and its correctness, we leave this to future and related work [22].

7 Related Work

Prior work by Kumar et al. [14] focuses on building a correct foundation for differential execution. They introduce a differential semantics with the properties of completeness, output validity, and from-scratch consistency. They also develop some optimization rules that significantly speed up the differential execution. Their semantics, however, recomputes the original store at every step, leaving caching as future work. In this paper, we build on the established correctness and provide efficiency by introducing provably correct caching, which is essential for gaining any performance

out of incremental computing and for preserving the correctness of the differential semantics. With incorrect caching, the differential semantics can get stuck because of missing entries or, even worse, can lead to invalid results that are hard to debug, for example when outdated cache entries are used from previous runs. Finally, our idea behind lexical tracing is general enough to implement incremental caching for other languages, which will be the case in our future work.

Xu and Erdweg [22] complement the work on differential execution by providing a theory for primitive differential operators that can be embedded into the differential semantics. Their differential operators are stateful: Each operator can define its own internal state paired with a state invariant. A stateful differential operator is correct if it maintains its state invariant and, given a valid state, it produces a from-scratch consistent output change. With our technique of lexical tracing, we can embed stateful differential operators into the differential semantics to associate each operator application with a unique cache key. The cache then includes the internal state of each differential operator application, enabling sound reuse of the operator state across different executions.

Dynamic-algorithm research maintains problem-specific data structures that let an update of size δ be handled in $\text{poly}(|\delta|)$ time—e.g. fully dynamic connectivity and minimal spanning trees [12], dynamic planar convex hulls [7], and decremental shortest paths. These solutions achieve optimal asymptotic bounds for their narrow domains, yet demand bespoke invariants and deep expertise. Our work follows a complementary strategy: by embedding incrementality into the language semantics itself, ordinary imperative code inherits efficient updates “for free,” avoiding the engineering effort of crafting special-purpose dynamic data structures.

Memoization systems such as lazy functional runtimes or memoization libraries reuse results at the granularity of pure function calls but collapse in the face of mutation [1]. Self-adjusting computation (SAC) generalises memoization with dynamic dependency graphs and change propagation [2–4]. Extensions add traceable data types [5] and demand-driven evaluation (Adapton) [11]. While SAC excels when control flow stays unchanged, branching or loop-bound edits often trigger whole-graph rebuilding. Our differential semantics bypasses dependency graphs altogether: change functions derived from the operational rules propagate updates directly, and lexical traces yield stable cache keys even when control flow diverges.

Memoization and self-adjusting computations fundamentally are based on selective recomputation [19]: a sub-computation is recomputed whenever any of its inputs changes in the slightest. Many approaches realize selective recomputing through dirt flagging: first mark all transitively affected inputs as dirty, then recompute all sub-computations with marked inputs. Selective-recomputation also underlies practical build systems and incremental compilers (e.g., make, Bazel, many IDE services), which use dependency graphs or timestamp comparisons to decide what to rebuild. Our semantics generalises those heuristics to *fine-grained*, language-level keys: each operation’s lexical trace plays the role of a build target, and cache entries record exactly the information needed to patch results after edits. And, of course, we perform differential computing and handle data changes (i.e., deltas) rather than changed data.

Semantics-based transformational incrementalization (STI) systematically rewrites a program f into an incremental variant f' able to compute $f(x + \delta x)$ from $f(x)$ and δx [15, 16]. STI embeds explicit *cache*, *prune*, and *discover* phases that (i) choose intermediate results worth retaining, (ii) discard entries no longer relevant, and (iii) introduce auxiliary state when it reduces update cost. Indeed, STI materializes the new full output in a single pass as an instance of selective recomputing, while we compute how the output changes, which composes with other differential approaches. Moreover, STI is applied at *compile time* and requires the developer to regenerate f' whenever significant structural changes occur, while our run-time differential semantics needs only the original program and adapts automatically through lexical traces.

Like STI, the incremental lambda calculus (ILC) also uses compile-time rewriting. However, ILC performs a type-directed source transformation of the original program to produce a derivative program [8, 9]. Like our approach and unlike STI, ILC's derivative program operates on changes: the derivative program maps input changes to output changes, given the original input. Since ILC derivatives also require the original input, the original input has to be patched at the end of each change-processing round to prepare for the next round. Moreover, the generated code often needs subsequent optimizations that have to be incorporated into the program generator. We instead execute the *same source* under an alternative semantics that uses caching: caching limits recomputed inputs to those rules that really need it and optimization opportunities live in the evaluator, not in extra generated code. Finally, ILC supports higher-order pure code, but does not explain how to handle imperative programming features.

Relational and data-flow systems employ delta rules to maintain materialized views [6, 10]. Differential Datalog [18] and Differential Dataflow [17] extend the idea to recursive and streaming settings with recursive aggregation [20, 21]. These approaches are domain-specific and only handle relational changes: inserted and deleted database tuples. We provide a general-purpose differential programming semantics and we mechanize its correctness and completeness theorems.

8 Conclusion

Starting from a big-step operational semantics, we (1) derived change-propagation rules that translate input deltas into output deltas, (2) proved from-scratch consistency and completeness in Rocq, and (3) implemented the semantics in an interpreter capable of transparently accelerating ordinary imperative programs. Specifically, this work defines a formally verified caching semantics in which every execution step is given a stable, unique key—its lexical trace—and shows that these traces are sufficient to guarantee from-scratch consistency under arbitrary control-flow variation. In proving correctness and completeness of the differential semantics, we established a novel theory about cache stability for lexical traces and their semantic rules. The resulting differential caching semantics is easy to implement, extensible, and yields the expected asymptotic improvements on incremental runs. This work forms the basis for the development of differential semantics for richer languages, where caching is an integral part of the formalization and correctness guarantees.

Acknowledgements

We thank the anonymous reviewers for their effort and helpful suggestions. This work is part of the AutoInc project and has been funded by the European Research Council (ERC) under the European Union's Horizon 2023 (Grant Agreement ID 101125325).

References

- [1] Martín Abadi, Butler W. Lampson, and Jean-Jacques Lévy. 1996. Analysis and Caching of Dependencies. In *Proceedings of the 1996 ACM SIGPLAN International Conference on Functional Programming, ICFP 1996, Philadelphia, Pennsylvania, USA, May 24-26, 1996*, Robert Harper and Richard L. Wexelblat (Eds.). ACM, 83–91. doi:10.1145/232627.232638
- [2] Umut A. Acar. 2005. *Self-Adjusting Computation*. Ph. D. Dissertation. Carnegie Mellon University.
- [3] Umut A. Acar, Amal Ahmed, and Matthias Blume. 2008. Imperative self-adjusting computation. In *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008*, George C. Necula and Philip Wadler (Eds.). ACM, 309–322. doi:10.1145/1328438.1328476
- [4] Umut A. Acar, Guy E. Blelloch, and Robert Harper. 2002. Adaptive functional programming. In *Conference Record of POPL 2002: The 29th SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Portland, OR, USA, January 16-18, 2002*, John Launchbury and John C. Mitchell (Eds.). ACM, 247–259. doi:10.1145/503272.503296
- [5] Umut A. Acar, Guy E. Blelloch, Ruy Ley-Wild, Kanat Tangwongsan, and Duru Türkoglu. 2010. Traceable data types for self-adjusting computation. In *Proceedings of the 2010 ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2010, Toronto, Ontario, Canada, June 5-10, 2010*, Benjamin G. Zorn and Alex Aiken (Eds.). ACM, 483–496. doi:10.1145/1806596.1806650

- [6] José A. Blakeley, Per-Åke Larson, and Frank Wm. Tompa. 1986. Efficiently Updating Materialized Views. In *Proceedings of the 1986 ACM SIGMOD International Conference on Management of Data, Washington, DC, USA, May 28-30, 1986*, Carlo Zaniolo (Ed.). ACM Press, 61–71. doi:10.1145/16894.16861
- [7] Gerth Støtting Brodal and Riko Jacob. 2002. Dynamic Planar Convex Hull. In *43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings*. IEEE Computer Society, 617–626. doi:10.1109/SFCS.2002.1181985
- [8] Yufei Cai, Paolo G. Giarrusso, Tillmann Rendel, and Klaus Ostermann. 2014. A theory of changes for higher-order languages: incrementalizing λ -calculi by static differentiation. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '14, Edinburgh, United Kingdom - June 09 - 11, 2014*, Michael F. P. O'Boyle and Keshav Pingali (Eds.). ACM, 145–155. doi:10.1145/2594291.2594304
- [9] Paolo G. Giarrusso, Yann Régis-Gianas, and Philipp Schuster. 2019. Incremental λ -Calculus in Cache-Transfer Style. In *Programming Languages and Systems - 28th European Symposium on Programming, ESOP 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings (Lecture Notes in Computer Science, Vol. 11423)*, Luís Caires (Ed.). Springer, 553–580. doi:10.1007/978-3-030-17184-1_20
- [10] Ashish Gupta and Inderpal Singh Mumick. 1995. Maintenance of Materialized Views: Problems, Techniques, and Applications. *IEEE Data Eng. Bull.* 18, 2 (1995), 3–18. <http://sites.computer.org/debull/95JUN-CD.pdf>
- [11] Matthew A. Hammer, Yit Phang Khoo, Michael Hicks, and Jeffrey S. Foster. 2014. Adaption: composable, demand-driven incremental computation. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '14, Edinburgh, United Kingdom - June 09 - 11, 2014*, Michael F. P. O'Boyle and Keshav Pingali (Eds.). ACM, 156–166. doi:10.1145/2594291.2594324
- [12] Jacob Holm, Kristian de Lichtenberg, and Mikkel Thorup. 2001. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *J. ACM* 48, 4 (2001), 723–760. doi:10.1145/502090.502095
- [13] Gabriël Konat, Sebastian Erdweg, and Eelco Visser. 2018. Scalable incremental building with dynamic task dependencies. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering, ASE 2018, Montpellier, France, September 3-7, 2018*, Marianne Huchard, Christian Kästner, and Gordon Fraser (Eds.). ACM, 76–86. doi:10.1145/3238147.3238196
- [14] Prashant Kumar, André Pacak, and Sebastian Erdweg. 2025. Incremental Computing by Differential Execution. In *39th European Conference on Object-Oriented Programming, ECOOP 2025, June 30 to July 2, 2025, Bergen, Norway (LIPICs, Vol. 333)*, Jonathan Aldrich and Alexandra Silva (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 20:1–20:24. doi:10.4230/LIPICs.ECOOP.2025.20
- [15] Yanhong A. Liu. 2024. Incremental Computation: What Is the Essence? (Invited Contribution). In *Proceedings of the 2024 ACM SIGPLAN International Workshop on Partial Evaluation and Program Manipulation, PEPM 2024, London, UK, 16 January 2024*, Gabriele Keller and Meng Wang (Eds.). ACM, 39–52. doi:10.1145/3635800.3637447
- [16] Yanhong A. Liu, Scott D. Stoller, and Tim Teitelbaum. 1998. Static Caching for Incremental Computation. *ACM Trans. Program. Lang. Syst.* 20, 3 (1998), 546–585. doi:10.1145/291889.291895
- [17] Frank McSherry, Derek Gordon Murray, Rebecca Isaacs, and Michael Isard. 2013. Differential Dataflow. In *Sixth Biennial Conference on Innovative Data Systems Research, CIDR 2013, Asilomar, CA, USA, January 6-9, 2013, Online Proceedings*. www.cidrdb.org. http://cidrdb.org/cidr2013/Papers/CIDR13_Paper111.pdf
- [18] Boris Motik, Yavor Nenov, Robert Edgar Felix Piro, and Ian Horrocks. 2015. Incremental Update of Datalog Materialisation: the Backward/Forward Algorithm. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*, Blai Bonet and Sven Koenig (Eds.). AAAI Press, 1560–1568. doi:10.1609/AAAI.V29I1.9409
- [19] G. Ramalingam and Thomas W. Reps. 1993. A Categorized Bibliography on Incremental Computation. In *Conference Record of the Twentieth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Charleston, South Carolina, USA, January 1993*, Mary S. Van Deusen and Bernard Lang (Eds.). ACM Press, 502–510. doi:10.1145/158511.158710
- [20] Tamás Szabó, Gábor Bergmann, Sebastian Erdweg, and Markus Voelter. 2018. Incrementalizing lattice-based program analyses in Datalog. *Proc. ACM Program. Lang.* 2, OOPSLA (2018), 139:1–139:29. doi:10.1145/3276509
- [21] Tamás Szabó, Sebastian Erdweg, and Gábor Bergmann. 2021. Incremental whole-program analysis in Datalog with lattices. In *PLDI '21: 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation, Virtual Event, Canada, June 20-25, 2021*, Stephen N. Freund and Eran Yahav (Eds.). ACM, 1–15. doi:10.1145/3453483.3454026
- [22] Runqing Xu and Sebastian Erdweg. 2026. Stateful Differential Operators for Incremental Computing. *Proc. ACM Program. Lang.* 10, POPL, Article 86 (Jan. 2026), 29 pages. doi:10.1145/3776728

Received 2025-10-10; accepted 2026-02-17